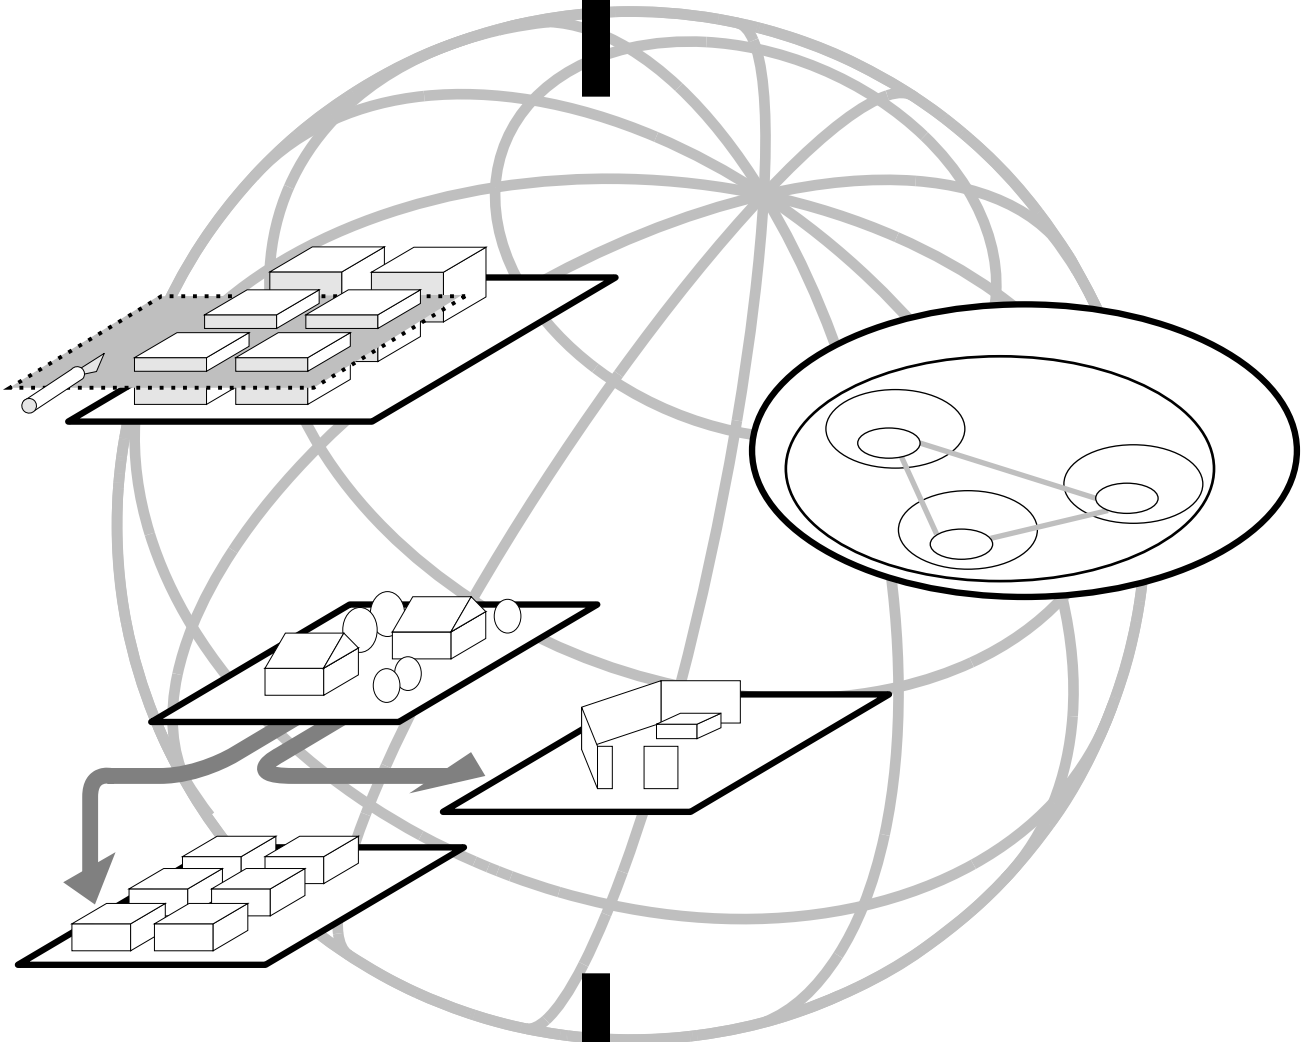


From WebSpace to Cyberspace



Kevin Hughes

Enterprise Integration Technologies

July 1995

From Webspaces to Cyberspace

Version 1.0: December 1994

Version 1.1: July 1995

Copyright ©1995 by Kevin Hughes

The opinions stated in this document are solely those of the author and do not necessarily represent the views of Enterprise Integration Technologies.

This document as a whole may be redistributed freely in any format for non-commercial purposes only.

Comments, questions, corrections, and suggestions relating to this document are welcomed and can be sent to kevinh@eit.com.

Trademarked names are used throughout this document; the trademark symbols have been omitted for editorial convenience with no intention of trademark infringement. Where such omissions exist the trademarked name has been printed with initial capitals.

About the Author

Kevin Hughes designs hypermedia products for EIT and is their webmaster. He has written *Entering the World-Wide Web: A Guide to Cyberspace*, an introduction to the World-Wide Web that has been used as training material in numerous companies and universities, and is a member of the World-Wide Web Hall of Fame.

Enterprise Integration Technologies
800 El Camino Real
Menlo Park, CA • 94025

Lobby: (415) 617-8000

Fax: (415) 617-8019

World-Wide Web: <http://www.eit.com/>

Thus science may implement the ways in which man produces, stores, and consults the record of the race.

Vannevar Bush
As We May Think
Atlantic Monthly, July 1945

The trouble with the future is that it usually arrives when you least expect it.

Arnold H. Glasow

Foreword and Preface

May 1993 was a quiet month, and it was business as usual on the Internet. There was email, news, and plenty of other services available online with discussion and software; the community was steeped in things like telnet, FTP, UNIX, and X; in many respects nothing had changed much for years. If my workstation had suddenly time warped to 1988, I would have noticed few differences. If I wanted to read a book on the Internet, I would have made a trip to the local university bookstore and purchased a good UNIX book. If I wanted to tell my friends about the electronic hobby I had immersed myself in, I would have first asked, "Have you heard of the word 'Internet' before?" There was little hope that I could explain things in any coherent manner if they had not first heard about, or read about, the Word.

If they knew, by some rare chance, they would inevitably reply, "But there's nothing there." They were right. Everything was there for me, for my classmates that pulled allnighters playing xtrek, for the slackers in the computer lab, for the occasional university student that wanted to learn about email. But there was little for everyone else, and I couldn't email just anyone else.

1993 was the year of the great Internet explosion, when the charts following net usage on the North American backbones went crazy. The numbers are still increasing. If I want to read a book on the Internet, all I have to do is find the Internet section of the bookstore. Telling my friends about what I do would serve no purpose since they already know; my relatives say, "The Internet sounds nice, but what would I do with all that information?"

The World-Wide Web is not a phenomenon. It's one step in a logical evolution. The World-Wide Web is not just a cause. It's a symptom. It's a reflection of how thoroughly my generation has been saturated with media, and a reflection of how badly previous generations are adjusting to that fact. The state of the Internet is an indicator of how we choose to act and react and of the ways in which we inform, learn, and dominate over others.

Cyberspace, a greatly misunderstood medium, is now upon us. There is a high probability that you have heard of this Word as well. Whatever meanings it has taken on in the past, there is one definition that influences and changes everything: cyberspace is *an interconnected computer-mediated environment in which all prior media are represented*. This document explores what this means for all of us: Where did it come from? Why? How can we work in this, and what can we do? How can we create and communicate in this ultimate medium? This is a beginner's guide, a primer for the Age of the Internet.

This document was first written in December 1994 as an internal white paper. It has since been greatly expanded and updated. The main sections are:

- **The State of the Web** - A brief history and opening.
- **Collaborative Systems** - An overview and analysis of current and future collaborative online media.
- **Cyberspace** - What it means and what can be done with it.
- **Creating Cyberspace** - Issues related to the development of systems that can be used to create and explore virtual environments within cyberspace.
- **Living In Cyberspace** - Issues related to building future human-computer environments for information visualization and exploration.
- **Tomorrow's Internet** - What the online future will look like. Topics related to business, media, online freedom, and what cyberspace needs.

Towards the end is a reference list of works, a timeline, and an index.

I would like to apologize for putting the initial version of this work into a paper-based format; I do so for the convenience of those that have traditional printers. Hopefully one day I will not have to think about measurements or double-sided layout when writing.

I thank Marty Tenenbaum and William Wong at EIT for supplying resources and supporting my vision. Thanks go also to Mark Pesce, Brian Behlendorf, Tony Parisi, and many others for spreading the cyberspace meme which is now lodging itself firmly in the 21st century.

This work is dedicated to all of my (four) parents and my teachers, doctors, and friends, not necessarily in that order or role. Although not everyone cared about the message, the important thing is that they listened.

Kevin Hughes
Menlo Park, California
July 1995

Table of Contents

Foreword and Preface	5
Table of Contents	7
1.0 The State of the Web	9
1.1 New Ways of Communication	13
1.2 Collaborative Systems	13
2.0 Cyberspace	35
2.1 A New Definition	36
2.2 Encapsulation	37
2.3 Notes on Successful Terminology	41
2.4 Myths	43
2.5 Why Virtual Environments?	46
2.6 Applications	47
3.0 Creating Cyberspace	65
3.1 VRML and HTML	65
3.2 New Definitions	65
3.3 Tools	66
3.4 MOO/WOO Thoughts	68
3.5 Coordinate Systems	69
3.6 Editing	70
3.7 Objects	71
3.8 Constructs	74
3.9 Links	80
3.10 Sensing Proximity and Location	80
3.11 Moving and Viewing	81
3.12 World Issues	92
3.13 Other Issues	94
3.14 A Prototype VRML/HTML System	94
3.15 Cyberspace Guidelines	99
4.0 Living In Cyberspace	101
4.1 Next Generation Environments	101
4.2 Vision	114
4.3 Input/Output Devices	120
4.4 Possibilities	127
4.5 Designing Worlds	144
4.6 Common Problems	158
4.7 The Open Interface	160
5.0 Tomorrow's Internet	161
5.1 The Future of the Web	161
5.2 New Business	165

5.3	The Medium is Still the Message	173
5.4	New Society	178
5.5	What Does Cyberspace Need?	193
	A History of Cyberspace	197
	References	239
	Timeline References	243
	Other Media	247
	Index	249

Everything good in the computer field has come about because someone said, "I want this, and I will build it for myself."

Ted Nelson
Computer Lib

1

The State of the Web

It has been over two years since the explosive growth of the World-Wide Web began. NCSA's *What's New* page, one of the most popular Web pages on the Internet, has evolved from what was once a meager listing of a few dozen sites to a corporate-sponsored page that is now updated three times weekly. From *The Economist* to *Wired*, from the *San Francisco Chronicle* to *Time*, from one person garage startups to billion-dollar multinational companies, news of the Web has spread faster than a journalist can type "information superhighway".

Some say that we are in the midst of a second revolution - while the first brought the word computer into daily conversation, the second has sparked the commercialization of the Internet, it has made online space a trendy place to hang out, it has done nothing less than tell people that one no longer has to be an undergraduate computer science student to enjoy surfing the net on weekends. It is bringing the world together online and its impact will be felt as long as anyone has something to say to anyone else.

From 100 sites to well over 20,000 sites in the last 24 months - the Web, the fastest growing service on the Internet today, has sparked a renaissance among the net life. This is, of course, despite the fact that the World-Wide Web project was created within the European physics community as a text-only system for allowing researchers to view scattered pieces of information in one place. And this is despite the fact that for four years the Web was mostly used by scientific institutions and was mostly ignored by everyone else.

Geneva, 1989: Tim Berners-Lee at CERN makes a software program called Enquire to organize his daily life: he has too much information, and he needs to be able to organize it all. He ends up liking it so much that he decides to create a larger system to collect physics information and works on developing a simple program to navigate through all of this data using hypertext links. He calls his effort the World-Wide Web project. For a couple of years as it is tinkered on and refined, his program is shown off at technical gatherings and wherever Berners-Lee can find a soap box to stand on. Some people whose jobs center around information, such as publishers, begin to do some simple experimentation by working on programs that can navigate through the Web, but for the most part it remains a toy for researchers and a curiosity at hypertext conferences.

For a while net people evangelize the Web, seeing the possibilities that could arise due to the very nature of its architecture, and those that can propose and create extensions to the system, allowing people to transmit not just text but other forms of media as well. After this simple change, administrators begin to create lists and warehouses of texts and images, because it seems like a good way to keep these things organized. Indeed, some users begin to make lists of information that are grouped very much like books, with introductory material on opening pages and more specific data in separate sections. The ability to link among all of these sections makes it possible to organize information within tree-like structures, with main bodies of thought making up the trunk and related files growing from them like branches.

It seems like a pretty obvious thing to do, since many similar systems have already been made on CD-ROMs or with HyperCard, but the difference is that any information created and put on the Web can be transmitted to just about anyone on the net with the right software. People can see your own creations, no matter if they are in the next room or in another country. Exhibits and virtual galleries begin to pop up online, and organizations that need to provide a variety of material to many users, such as universities, begin to start using the Web to create inter-departmental information systems that bring the exploration and viewing of what was once dull data into a new light. Some organizations that are used to allowing users to offer their own information via FTP begin to allow people to display short biographies and personal information on the Web. Hundreds of campus undergraduates take advantage of this ability, and some people begin to store and administer their own collections of files locally. By now the few manuals available on creating Web-compatible media are gaining in popularity, and as homemade content grows, so does the capability of users to learn how to create content from others.

As the collective experience of these authors enlarges, conventions revolving around the creation of information on the Web come into being. In addition to the traditional "postmaster" email address that UNIX boxes have, and the more recent "gophermaster" alias that many system administrators with a Gopher site sport, a new term comes into being - webmaster. Every new Web site sparks a dozen new ideas in the minds of aspiring webmasters.

Once users understand that it is possible to show thousands of people what they want to say, and that they can embellish these thoughts with all sorts of cool media clips (home movies, the sound of their dog, the latest CD from the store down the street), and that they can actually move from file to file in this space without having to move so much as their digit finger, it's no surprise that somebody would begin to use the system for personal expression. And if you can make all sorts of media available to hundreds, maybe thousands, of people, why, that sounds a lot like what advertisers do, doesn't it? It takes a very short conceptual leap to figure out what any business could do with such a system, and some large corporations decide to take the plunge

and create Web sites. This gives the signal to scores of smaller businesses and individuals to create many of the trappings of a marketplace economy. But out of this mess of users and providers that subsequently arises comes the need for organization, security, and all the things that groups of politicians start to think about when they form governments.

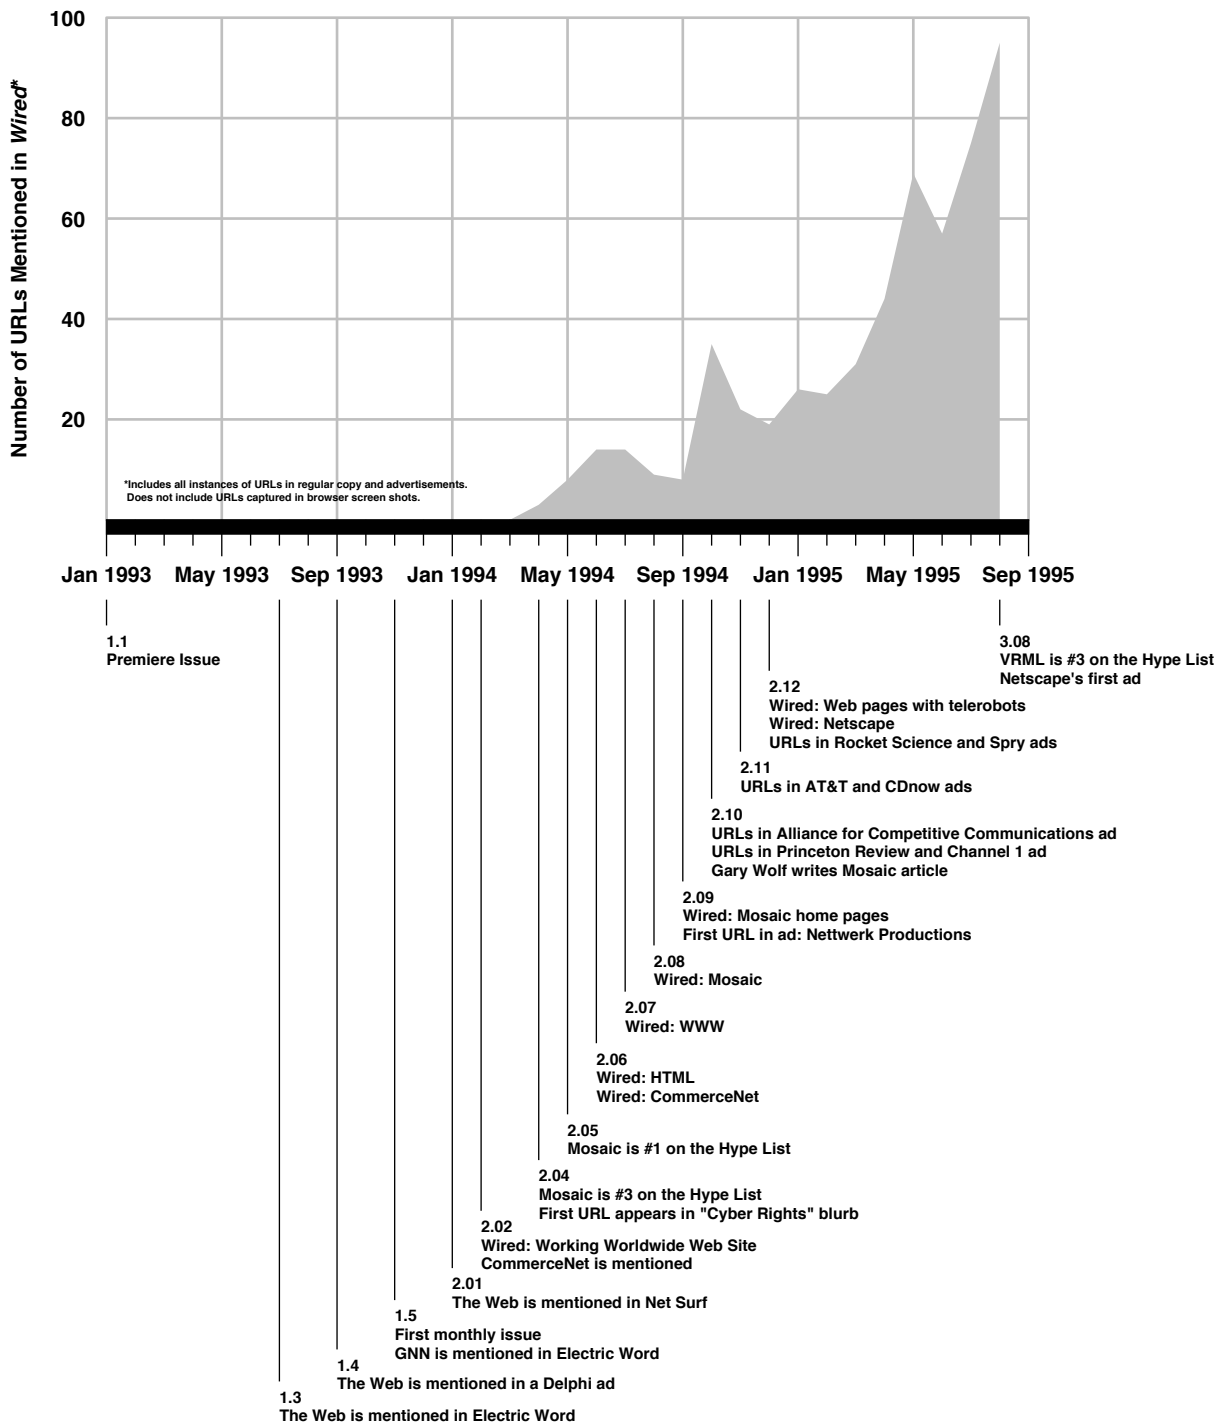
At this point the Web moves into the awareness of the popular media, which begins pointing out that everyone must have it. And this potential demand attracts packs of venture capitalists and creates new job titles, strange business models, and scores of books. And this is the point at which we are.

Today we branch off the links of giants. Hypertext is a 50-year-old idea and a 30-year-old technology. Hypertext is a word that was consistently miscapitalized by Berners-Lee. Hypertext used to be the kind of thing that was only discussed among arcane literary groups and underpaid university researchers. It was only until Berners-Lee created the first NeXT-based Web browser at CERN, which sparked efforts at O'Reilly and Associates, a computer book publishing company, which sparked efforts at NCSA, which sparked efforts by everyone else, did hypertext truly begin to reach the masses. The point is, is that we can create technologies from any idea, but before anything happens something must be built and used by someone first. Demo or die.

Look at the development of most successful ideas and they will seem obvious and logical. As you walk down the street of progress, new avenues appear. But a minor change in plans today will create a major change in the way things are done tomorrow, and it must be possible to experiment with new ideas early on before we are saddled with the results later. We can regard the World-Wide Web as one big experiment in creating a global multiuser collaborative information system. It's possible to see where some decisions that were made early on have affected things today, and we can apply what has been learned towards the development of such systems in the future. For this reason, many concepts put forth in this work have been created by looking at traditional World-Wide Web concepts and applying them in

a new environment. With this in mind, let's begin to talk about the next generation of tools, the next generation of users, and cyberspace.

1: One indicator of the evolution of the Web: the number of URLs found in *Wired* magazine.



1.1 New Ways of Communication

People want to communicate. The Internet grew 25 years ago with the birth of networking technology - communicating machines - and the growth of electronic mail - communicating individuals. Look at the most used services on systems such as America Online and the net itself, and chat services such as email and IRC (Internet Relay Chat) are guaranteed to be very near the top of the list, if not at the top. Whatever sort of information systems arise in the future, they must be able in some way to address communication with other people and programs. With communication we can help change, correct, and create data. With communication, we can create societies, and this is one of the major reasons for talking to people in the first place.

The Web is primarily a read-only medium - most of it appears static, and feedback is limited to the ability that online forms give. There are no real-time asynchronous capabilities built into the Web - when actions are performed, the results often don't appear until some seconds after. Collaborative systems are limited to the traditional type of low speed systems already available - mailing lists, news, and email. The original NeXT browser had an editor built in, so that users could make changes to documents without needing to know their HTML internals. The first versions of Mosaic for X Windows had an annotation capability built in, so that users could put annotations on pages and respond to other people's comments. These features, which were abandoned due to lack of time, manpower, and/or funds, are being reimplemented today because of sheer demand through jerry-rigged solutions and if fully developed could have created a more productive and progressive sense of society that the Web now lacks.

As Mark Pesce likes to quote, "There is no 'there' there." Communities make it possible for people to feel safe in and be familiar with environments. They help bring together like-minded individuals, making it possible to learn from others and create movements and activity on a larger scale. Communities create change and revolution and can help make sense of the overall organization of things. These are some of the reasons why institutions such as the town meeting and the local pub are so successful. Until it is possible to create online communities, much in the same way that email, Usenet, and user groups do, we cannot hope to grow or communicate. Using the Internet without a sense of society is like being locked in a room with the Library of Congress.

1.2 Collaborative Systems

The Elements

In thinking about new collaborative systems, ones that do offer a "there" there, we can think about what makes communities successful and the different kinds systems that exist today. Here are some of the things that are needed to make successful online communities.

- **There must be some sort of persistent storage.** When I make a change in the system, it must be stored in such a way so that other members of the system can see the change. Examples of this may be the building of objects in a system, an annotation or correction in a file, a transcript of a conversation, or the creation of an article on a bulletin board system. This point and others below raise issues concerning who has the permission to make and view changes and how changes in the system are created, deleted, and maintained.
- **I must be able to give and receive things directly to and from members.** While members of a system with persistent storage can receive and send information indirectly, the ability to transfer data directly can help make communication more efficient and may allow members to easily share data privately.
- **Members must be able to gather at a particular place.** There must be some sort of address or virtual space at which members can gather. For example, this place may be an IRC channel, a Web site, or a CU-SeeMe reflector. The ability to have a predefined meeting place can help keep unwanted people out and can bring a sense of familiarity to the group.
- **There must be interactivity.** That is, members must be able to communicate with others in a reasonably asynchronous way with as little response time as possible. Conversation in the physical world is often full of interruptions, simultaneous talking, and other voices coming and going. The spaces between words and style of conversation can convey much, and systems that preserve these qualities or provide methods to mimic them make interaction more humane.
- **Members must be able to experience, mimic, and follow other members.** A large part of learning consists of listening to and experiencing other people. The ability to learn from others in real time, not just from previously stored information, allows education to be up to date and highly interactive. Learning by doing is the best way for people to learn, and collaborative systems should not lose sight of that fact.
- **Members must be able to enter and leave groups at any time.** Forced communities are not true communities, and for people to learn and communicate freely they must be able to assemble and disband at will.

These properties should allow the creation of environments in which people can effectively communicate ideas and learn from each other. Systems that take these ideas to heart will likely be popular and used to their maximum potential. Usenet, for instance, has been instrumental in many cases in allowing local communities to affect business and government. Although it has also brought about the creation of unwelcome communities and online wars at times, this is to be expected as long as humans are part of the equation.

Systems that make it easy for people to create a culture through communities are likely to be well used, but there are other factors as well. We can look at the success of the World-Wide Web and other services for answers. Why has the Web been so popular?

- **It has made it easy for people to create and distribute media.** Anyone can make their voice heard; there are mechanisms in place (public announcement places, as an example) for them to do so. The ability to create and send data which can be easily navigated and viewed has allowed people to communicate all sorts of ideas; the more kinds of media that can be created and distributed in a system, the more effective learning and communication (whether through conversation or personal expression) will be in it. The creation of objects in the system should be as easy and intuitive as possible.
- **It is networked.** That is, it allows many users to communicate with each other remotely. The nature of computers today is such that only one person can use one computer at a time to send and receive information; therefore any collaborative online system must allow multiple users on multiple terminals to converse. It helps greatly if the networkability and architecture of the system is scalable and consistent, so that nearby users can be treated conceptually the same as remote ones.
- **It has created a uniform addressing scheme for many disparate services.** There is a consistent way of treating services native to and known by the Web, and it is such that new services that need to be integrated can be done easily following the same method. The Web has done the equivalent of homogenizing the world's phone numbers.
- **It is fairly consistent in its rules.** The URL scheme and navigation methods used by most browsers today are such that users need only to know a minimum amount of information (URL creation rules, for instance) in order to navigate the entire system, and users that need to formulate new rules can do so fairly successfully by guessing. The design of any system should combine features and use a

consistent method in its interactions with users, so that very little information about using the system has to be remembered but instead can be inferred by extending its rules.

Many of these features mentioned in this section can be summarized into three major points:

- Things aren't distant
- Things are persistent
- Things are consistent

Manipulation and communication methods should make use of real-world knowledge (things aren't distant)

Real-world knowledge is incorporated into the system - actions in the system are similar to actions in the physical world - so that users aren't forced to learn new methods. This is not necessary, but it can help people process information in ways that are familiar to them.

Changes to places and objects should be observable and kept for indefinite periods (things are persistent)

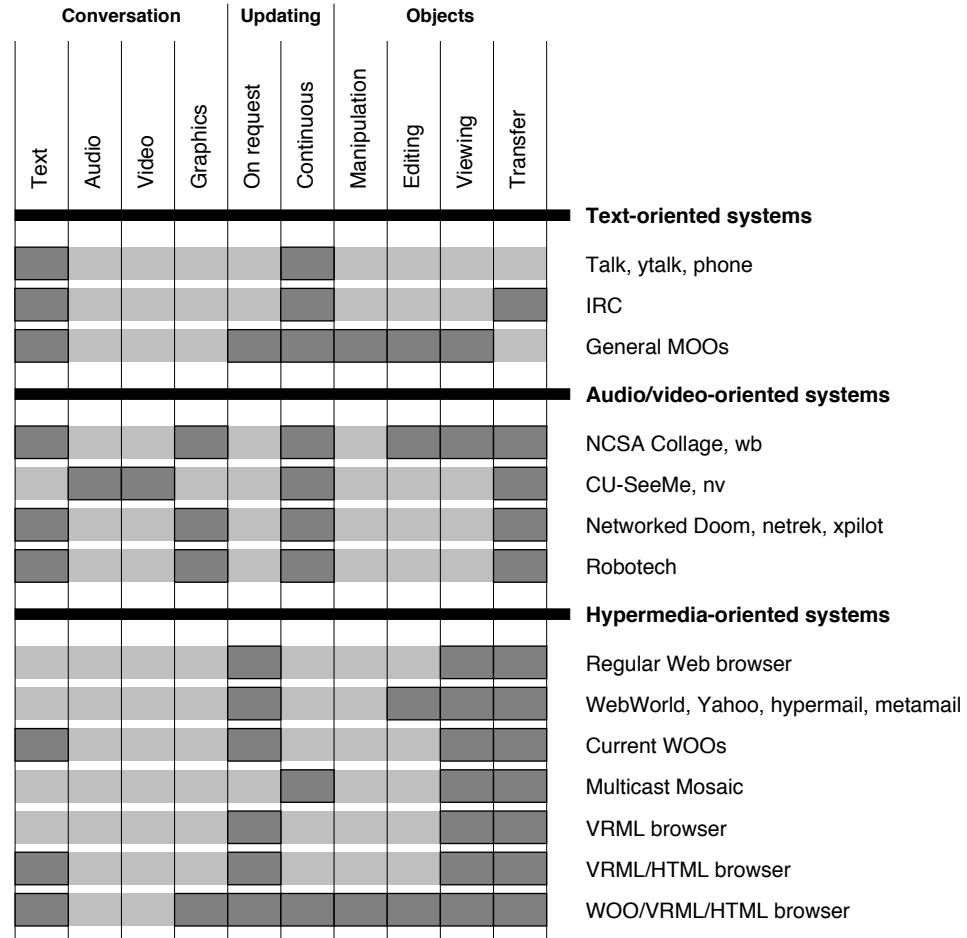
People can see what users doing and the changes they make. Users can decide to make something and keep it in their house as long as they live, or they can tell it self-destruct after a while. They might be able to give something a life of its own. All of this assumes that members can create things; they must be able to be authors as well as viewers.

Everything in the system should have common ways of interacting with it, so users need only remember a few basic things to do many complex things (things are consistent)

The metaphors the interface uses are consistent and not misleading. Functions are combined. Users only have to know one thing, but under different contexts they can do many things. They can select a link, an image, a button, an icon, a menu bar, a service, a program, and depending on what is selected they can turn a page in a book, visit Europe, or watch a movie, just by performing one action.

The State of Collaboration

A well-thought out cyberspatial system can be created which has all of the aforementioned attributes, but today a number of services with varying features are used. Mentioned are a number of networkable, interactive systems:



2: This chart displays the features of some of the collaborative systems that are used today.

Conversation:

This refers to real time (or very close to real time) conversation, either one way or both ways at once. Programs such as talk, IRC, and MOOs can provide this kind of collaboration. With Collage and other whiteboard programs one converses using graphics. To converse using audio and video, one can use nv. I categorize “graphics” as bitmaps and animation that doesn’t represent live video, so games such as xpilot fall into that category.

Updating:

Continuous updating means that one can be using the system and data representing changes in the system are sent to you in an unsolicited manner. Members don’t specifically have to ask for new information. For example, in Doom the computer shows the objects that one gets close to automatically; it doesn’t have to be

asked. This type of communication may be called asynchronous, where multiple channels of information are sent and received without waiting for particular events to happen first.

Updates on request means one has to ask to see what changes in the system have been made. On the Web, pages must be reloaded to see any changes; most Web-based systems are like this. This can be called synchronous communication, where something must be done before sending and receiving data.

Objects:

Objects refers to media other than text and virtual structures.

If a user can pick up objects and move them around, that is, do things to something another user has created, then I say that the objects can be manipulated. If users can change, create, or delete objects, then the objects can be edited.

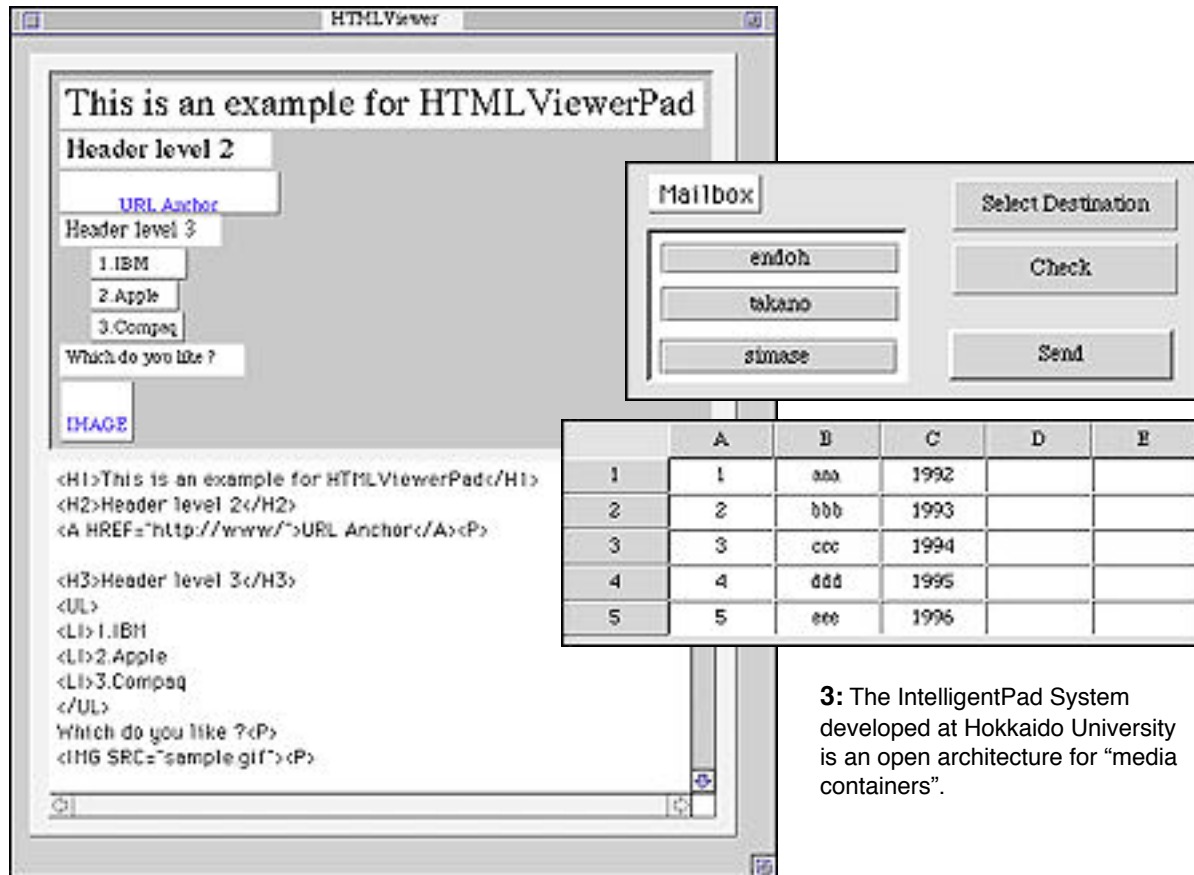
To be able to transfer objects, users must be able to package up different kinds of objects that aren't native to the system and send them to other users in the system. With IRC it is possible to send someone a GIF, and one can FTP nearly anything using the World-Wide Web. Transferability is a measurement of how extensible the system is, how easily one can use other kinds of media to do what they want.

- Collage is available at <ftp://ftp.ncsa.uiuc.edu/Visualization/Collage/>
- CU-SeeMe is available from <ftp://gated.cornell.edu/pub/video/>
- Information on IRC is at <http://edb518ea.edb.utexas.edu/html/irc.html>
- XPilot information is at <http://tklab3.cs.uit.no/XPilot/>
- Doom information is at <http://www.idsoftware.com/>
- MBONE software pointers and other MBONE information is at <http://www.eit.com/techinfo/mbone/>

Modular Browsers

We are moving into the post-Web era, in which the idea of having one monolithic Web-browsing application that does everything will take a much less prominent position than it does today. With the development of modular plug-in architectures, drag-and-drop capabilities, and interapplication communication protocols it is best to think about the next generation of Web browsers as being plug-ins in an application or extensions in an operating system that can interact with other tiny "applet" programs to get the job done. It will be possible to read HTML in many applications (take Microsoft Word as one example), and various capabilities will be combined to create new tools for computer users. Imagine a networked HTML-reading secure spreadsheet program, Frame-Maker with URLs, or Adobe Acrobat running CGI programs.

Future references to Web browser programs do not necessarily mean that they refer to one application as we know it today - keep in mind that they may refer to sets of programs or modules working together. The day of the Web browser as we now know it is fading fast.



3: The IntelligentPad System developed at Hokkaido University is an open architecture for "media containers".

- Information about the IntelligentPad System, a modular application framework, is at <http://www.huee.hokudai.ac.jp/>
- Adobe Acrobat information is at <http://www.adobe.com/Acrobat/Acrobat0.html>
- Information on InternetLink, a free plug-in that allows Acrobat PDF files to link to TCP/IP programs via URLs, is under <gopher://boombbox.micro.umn.edu/11/gopher/Macintosh-TurboGopher/>

MOOs *Multi-User Dungeon - Object Oriented* systems allow text-based interactive communications between members and give them the ability to create and edit virtual objects in a text-oriented manner. Many MOOs provide commands that allow users to interact with the environment and others in different ways. MOOs have most of the capabilities of text-only chat systems (there can be multi-user conversations, members can use handles, system com-

mands are provided) with the exception that persistent environments can be created by one or more users and given properties so that they may be interacted with.

MOO systems feature many of the necessary traits that foster the development of online communities - indeed, there is a MOO culture and there are MOO conventions; classes and simulations have been held in MOOs; people have even been married in them. What makes text-based systems like this so valuable is the nature of text-based media: the more primitive the media, the more one's brain has to fill in the gaps in communication. This is part of what we call imagination, and it enables creativity and makes learning and communication engaging. Books fire the imagination, and in MOO systems many people have written mini-novels in the form of fantastic virtual environments. Users of these environments are more like readers of interactive stories, and many contribute to the action and scenery when they feel like it.

Because of the creativity that text can foster, it is important not to abandon it as a medium in developing future systems. There must always be a way to manipulate and communicate using text, since:

- Text is a major form of communication and many people are familiar with it
- Text makes up a large part of Internet communications today
- Text makes it possible for users with low-end systems to be members
- Text is easily archived and manipulated

Should networked real-time digital audio storage and manipulation become popular, it may replace text-only communications in many instances. Until that day, however, text will remain a vital part of online life.

Traditional MOO systems are more than adequate for providing virtual settings, but the information navigation methods they use are more command-line oriented than mouse driven. Text-only MOOs in their current form are not well suited to the display and transfer of any kind of media other than text. Many groups have already begun work to fulfill the need to enrich MOO information space by creating different hybrids consisting of parts of MOO and Web-like systems.

- MUD and MOO resources can be found at http://www.yahoo.com/yahoo/Entertainment/Games/MUDs_MUSHes_MUSEs_MOOs_etc_/

Web and MOO Work

We can create new systems by looking at the more common forms of media and formats used in communications on the Internet today:

- Text, hypertext (email, telnet, Gopher)

- Audio, video (Sun audio, QuickTime, nv)
- Two-dimensional graphics (GIF, JPEG)
- Three-dimensional graphics (VRML)
- Hypermedia (HTML)

...and combining them with the different methods of communication that existing systems (collaborative or not) offer. In this way we can produce new tools, new forms of entertainment, and new ways of creating societies. Here are some examples:

- Merging hypermedia and MOO systems (WOO)
- Merging three-dimensional graphics and text (3D chat)
- Merging hypermedia and page layout systems (HyperCard, Adobe Acrobat)
- Merging hypermedia and three-dimensional graphics (prototype VRML)
- Merging hypermedia, three-dimensional graphics, and MOO systems (prototype cyberspace)

Today many people have made new systems by mixing together only the more common elements of current services; out of this has grown a new class of programs that are gaining popularity.

4: A visualization of what 3D gopher-space might look like, which was presented at the 1994 Gopher Conference. Users are in the center of an array of options. This scheme is already used in TurboGopherVR.



- One MOO with a graphic interface is exMOO, at <http://www.dl.ac.uk/CBMT/exMOO/HOME.html>
- Information on 3D gopherspace visualization is under gopher://boombox.micro.umn.edu/11/gopher/Gopher_Conference_94/Presentations
- TurboGopherVR can be found at <gopher://boombox.micro.umn.edu/11/gopher/Macintosh-TurboGopher/TurboGopherVR>

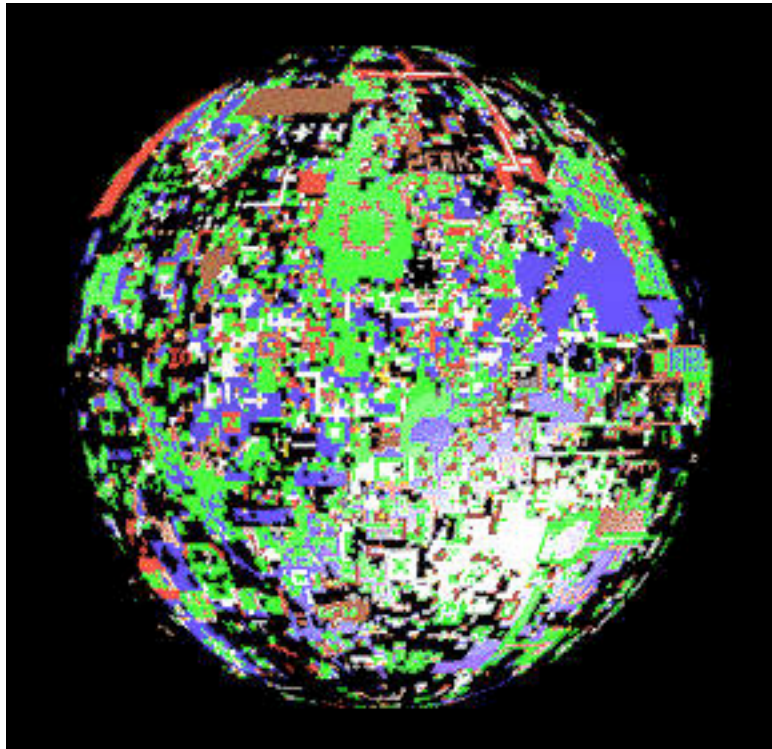
WebWorld was a system that was in every way much like a typical Web site, except for the fact that all interactions took place in a virtual world, one through which users navigated using a simulated three-dimensional rendering of the landscape. In WebWorld, users created different types of basic structures, such as land, sea, buildings, pyramids, houses, trees, and pavement. In its first five-months, users built more than 125,000 objects with

over 4,500 links on the system, creating such virtual areas as technology parks, structures that spell names, and shopping malls.

WebWorld worked with a simple database and server logs to determine where users currently were in the system and kept track of the busiest structures and links. Automatic robot programs bulldozed over bad links and took care of the virtual landscape. WebWorld showed the potential for systems that allow the creation of three-dimensional environments; it was very popular despite the fact that users could not directly interact with others and that objects had no real properties other than the ability to link to other parts of webspaces. Each movement in WebWorld entailed the separate rendering of the landscape at the current coordinates and it took many minutes to do any sort of extensive traveling. But it points out the fact that people are willing to take the time to create virtual worlds if creation and navigation is easy.

Simple hypertext-only systems that follow WebWorld's model are programs such as hypermail and hypernews. Instead of objects, users post articles and navigate through threads of ideas. Similar services such as Usenet have been very successful, but these systems are best suited to means of communication that must be archived and forums in which the ability to review and navigate information from others is important. One system called WAXweb appears to follow this model; the Web site consists of hundreds of images and pieces of text culled from the movie "WAX or the Discovery of Television Among the Bees". Users can add to the site using links and pictures, thereby growing new conceptual spaces in the system. Many meta-indexes now on the Web give users the ability to add to the database and provide comments through the use of forms; one of the first and most fully-featured is the Internet Movie Database (originally at Cardiff University), which allows users to enter movie information, provide reviews, and vote on other movies. One can see the how other movies have been rated and update information for any entry.


5: This is the entire contents of WebWorld mapped onto a sphere. If every object represented a Web server, this picture would represent about half of known webspace. About 45 new objects were created per day in WebWorld.



- Started in August 1994, WebWorld ceased operation in early 1995 for unknown reasons. Questions about WebWorld may be directed to webmaster@peregrine.com.
- WAXweb is at <http://bug.village.virginia.edu/>
- The main site for the Internet Movie Database is at <http://www.cm.cf.ac.uk/Movies/>

The **Yahoo** list, one of the most popular subject-oriented meta lists on the Web today, averaged about 50,000 requests daily in late 1994 and is built from user contributions and results from Web robots. New entries are flagged with a small "New!" picture and are updated and verified regularly. There are well over 30,000 links currently represented in the Yahoo system. Yahoo takes many of its ideas from the first user-customizable hotlist that was developed at NCSA by Eric Bina, one of the original Mosaic developers. In the early days of Mosaic's first official release, this list grew quickly with users contributing entries with links and graphics in them.

- Yahoo is at <http://www.yahoo.com/>



6: The Ping interactive landscape. Each triangle represents a link to a picture, sound, movie, HTML page, or 3D object.

The **Ping** project is one of the more intriguing examples of a user-built site. By selecting a location on a two-dimensional landscape (represented by an imagemap), users can add objects by specifying links to their media, much in a similar way to how WebWorld is built. This landscape is then subsequently generated in three

dimensions in real time by an SGI Reality Engine. As users contribute media to the system, the camera is directed towards the object as it appears in the three-dimensional landscape.

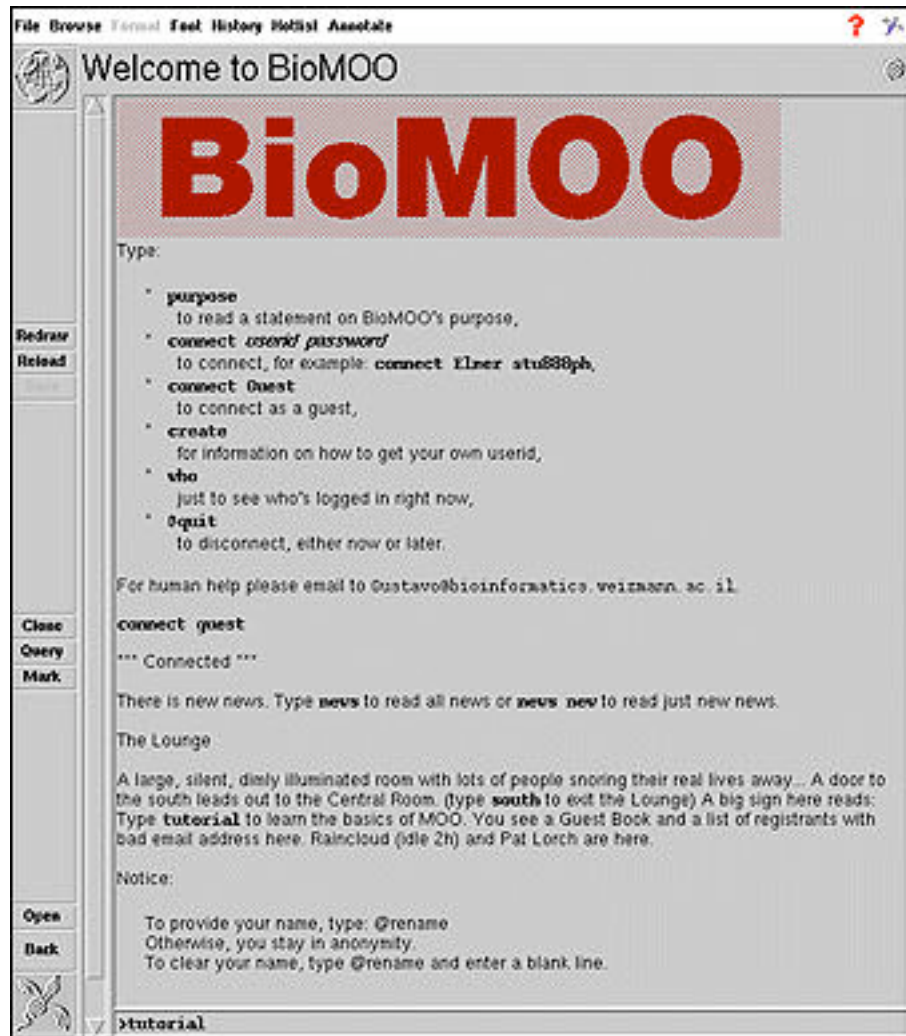
- Ping is at <http://www.artcom.de/ping/mapper/>

These are all systems that allow users to add to them incrementally but do not support the direct transfer of information among users or the kind of object creation that have made MOOs so popular. Many of these systems have been incorrectly referred to as hybrid Web/MOO systems; rather, they are simply nothing more than a particular class of collaborative service in which users can build things (sometimes cooperatively), but very little real-time interaction takes place if at all.

There are a number of MOO systems that have integrated the ability to link to Web pages - for instance, a player may want to look inside a room or view an object, but instead of getting text descriptions of these things a Web page can be brought up inside a separate browser program and viewed. These are referred to as *WOO* (Web plus MOO) systems. In addition to traditional MOO systems, there are others that use custom GUIs or other methods to display 3D images and other types of graphic information.

Telnet-based WOO systems take the basic capabilities of a telnet-reachable MOO server and add the ability to view Web pages that represent areas of the MOO; one may even be able to explore the MOO via the Web using links, but one cannot use the system to its full potential without making use of a telnet session. The BioMOO project developed at the Weizmann Institute is an example of such a system. HTTP-based WOO systems take the basic capabilities of a Web browser and integrate MOO or chat capabilities into them, so that all actions and commands can be input through the browser.

7: The BioMOO WOO, shown in the Phoenix World-Wide Web client. Phoenix is designed to act as a front end to MOO systems and as a Web browser.



- BioMOO is at <http://bioinformatics.weizmann.ac.il:70/1s/biomoo/>
- Information on Phoenix is at <http://www.bsd.uchicago.edu/~l-newberg/phoenix.html>
- The Sprawl is at <http://sensemedia.net/sprawl/11>
- htMUD is at <http://www.elf.com/~phi/htmud.html>
- A Web/MOO Integration paper is at http://www.bsd.uchicago.edu/Staff/Web_Notes/MOO-overview.html
- General WOO information and other systems can be reached from <http://sensemedia.net/about/>
- A page on WWW and MOO research can be found under <http://www.maths.tcd.ie/pub/mud/moo-www/directory.html>

Of the HTTP-based collaborative systems, the **Sociable Web** project, headed by Judith Donath and Niel Robertson at MIT's Media Lab, is nicely innovative. The Sociable Web consists of a modified Web browser (in particular, Mosaic for X Windows) and a modified server. The Web browser looks and acts like most any graphical Web browser with the exception of an extra window that contains a text entry space, a text display space, and a space that can list icons and names. Much in the way that MOOs allow people to create rooms and virtual gathering places in which to meet, the Sociable Web uses Web sites as gathering places.

Every Web server can act as a chat channel, an environment at which people can gather. When I enter a URL and go to a Web site, I can see a list of people currently at that site as a list of names and/or icons representing them. Just like IRC or a MOO, I can type to all or any of these people and they can respond, and our conversation will be recorded in the text display area. By creating Web pages at sites, virtual environments can be created, and I can type HTML to other people, which is rendered correctly within the text display area.

The Sociable Web is an elegant solution in that it takes advantage of the capabilities of the World-Wide Web and integrates them fairly seamlessly with a traditional text-only method of conversation. Current Web and MOO solutions often put the bulk of the computing resources on the server, using programs such as telnet as the main interface and the Web as nothing more than a display device that handles media that telnet can't. The Sociable Web takes the other position of beginning with the intuitive environment of the common Web client and building conversational capabilities on top of that. This is the more forward way of thinking about collaborative systems and represents the beginning of what will be a new family of services.

Ubique's product, **Sesame**, provides a slightly different solution in that users are displayed as icons within the browser's display area. If a person wishes to take others on a tour, he or she can roll out an icon of a magic carpet onto which other members can drag their icons. Wherever the carpet's owner goes, the other members can then go as well.

- "The Sociable Web", "Casual Collaboration", "The Illustrated Conversation", and "Visual Who" are papers available at <http://big-sleep.media.mit.edu/Judith/home.html>
- Ubique is at <http://www.ubique.com/>
- Information about the MOO BOF at the Second International World-Wide Web Conference in Chicago is at <http://union.ncsa.uiuc.edu/HyperNews/get/bof/MOO.html>

A Future WOO System

The Sociable Web project is well done, but it represents only a start. We can build many more features into such a system to make it even more collaborative, efficient, and useful. The following are features which can be built into a second-generation well-integrated WOO:

First, such systems should act in every respect like a normal Web browser and server. They may be compiled without WOO capabilities, in which case no conversation at the server is possible. This idea makes it easy for users who already know how to use such browsers, and makes it easy to envision adding WOO capability to normal browsers and servers using add-on kits, patches, upgrades, and so on. Normal users who wish to use WOO services should be able to call up or hide the chat windows through a Window menu and the default action should be that one is not a part of conversation at the server but can observe ongoing conversation by opening the chat window.

Conversation must be particular to pages. One problem with the Sociable Web model is that a particular conversation applies to everyone using the same server, no matter where in the file space other members may be. It does not seem like an easy problem for servers to keep track of conversations that are particular to certain pages, but the ability to do so can open up many new possibilities for these systems. Server administrators should have the option of limiting groups to pages or whole sites.

Members can join and leave server chat at will. There should be some sort of widget within the chat window that allows users to join and leave the conversation. One should be able to easily tell at a glance what their status is by glancing at this element. If a member joins or leaves, their chat window is still displayed so conversation by others can still be observed, but the server will announce that they have joined or left the discussion.

Conversations can take full advantage of the capabilities of HTML. The ability to let people converse in HTML opens up an incredible wealth of possibilities. Conversations then can be displayed as HTML documents and can be treated and addressed as such. A WYSIWYG method of writing HTML may be developed (for instance, a user may select text and use a mouse button or key combination to bring up a pop-up menu of styles, attributes, and links) or users may type raw HTML to each other. The list of server members and icons can have links in them so that users can select someone's icon and have their resume or personal information appear. People can include images and sounds in conversation and include references to documents and services in a meaningful way. Conversation can be formatted; this gives people the ability to create new conventions in HTML speech much like the smiley in plain text. Imagine being able to include a picture of your smiling face, or a sound of you laughing. Users can cut and paste HTML documents and references to them into conversations; users can create graphic labels for

themselves to denote the information that they have included. Horizontal page breaks can denote different users and colored sections can be used to denote new threads of conversation.

Users can customize their persona and style of conversation and create macros. By specifying preferences to browsers, it should be possible to change entering and leaving announcements for your persona, your icon, and many other variables. It is possible to make a macro (and/or variable substitution) scheme available to users so that common phrases and ideas can be easily typed. It may be possible to reference CGI or other types of programs so it may be possible to create bots, or computer-driven robot members which can act as watchdogs for groups, interactive programs, and so on. IRC in its current form implements all of these features; it is worthwhile to study them and see how they may work in a WOO environment.

Conversations are archived by the server and can be saved, mailed, and browsed by users. Because conversations are HTML documents, they can be viewed and manipulated via CGI programs on the server side. For instance, I may enter a WOO server and go to a conversation archive page, where I can choose to view the conversation for the last week, last day, or last hour. I may want a transcript with conversation from certain users filtered out, or conversation that follows only a particular thread, or I may request a PostScript or text-only version of these conversations. Using this kind of interface it is possible to manipulate traditional access and error logs, something which many people have written CGI programs to do today.

Servers can restrict who can converse and who can read conversations. Through some method of access control on the server side, it should be possible to restrict who can read conversations and who can converse, much in the same way that access control today allows administrators to decide who has access to the server or not. This can help foster the development of private groups and the development of public areas which may display conversations but not let anyone contribute to them.

Members can choose to hold private conversations and follow others. The Sociable Web project offers a list of members' icons to denote who is active in the conversation, but it seems that there need to be at least two other options for each member: the ability to chat privately with them and the ability to follow them. Next to a member's icon there may be two toggle buttons: one marked Private and the other marked Follow. When I select a member's Private toggle, this means that anything I type will be sent to this member only until I hit the Private toggle again. If I wish to speak to a group, I can select more than one member's Private toggle to do so.

Any collaborative system should be able to give a member the chance to play leader. In traditional shared whiteboard systems, users may drop and pick up "chalk" tokens which indicate who has the capability to draw at any time. WOO systems are a little

different in that users do not choose to lead; rather, they choose to follow others. People may choose to wander around in web-space and leave a conversation at any time while others may choose to stay. In whiteboard systems there is usually no way of ignoring input from the leader who possesses the chalk. This in a way becomes an interesting form of coercion; it is better to give users a choice in these matters.

When I want to follow somebody I select their Follow toggle. If they have indicated in some way that they do not want to be followed automatically, then this button will be unusable (users may be able to type an "invite" command that enables others to follow them). I can only follow one person at a time - when I do, the actions of this person as they roam around in webspace are mirrored exactly as they are performed in my Web browser. The extent of this mirroring is up to developers - certainly any URL, scrolling, and highlighting should be mirrored, while other systems may choose to mirror the actions of external viewers via multicast or other such methods. When a user opens a new window, this action may be reflected in the conversation: "John Smith opens a new window to MIT's Media Lab", where the descriptive text is a link to the new window. Incidentally, by having multiple windows open, users can keep track of and participate in many discussions at once. When a user leaves a server, this action may also be reflected: "John Smith links to Wired. John Smith has left the server", where a link is provided so that users can follow others manually if they so choose. There is a common problem that arises in multicast applications; this is what one may call the "wandering riot" problem: If you are leading a virtual tour group on the Web to a particular server, this server will get hit once when you arrive and at least once for every member of the group that is following you. This slows down network speed considerably and introduces an increasing lag time for group members. One possible solution to the problem is for the leader to receive the initial data via HTTP and then have the system pass the data among the members using a more efficient protocol.

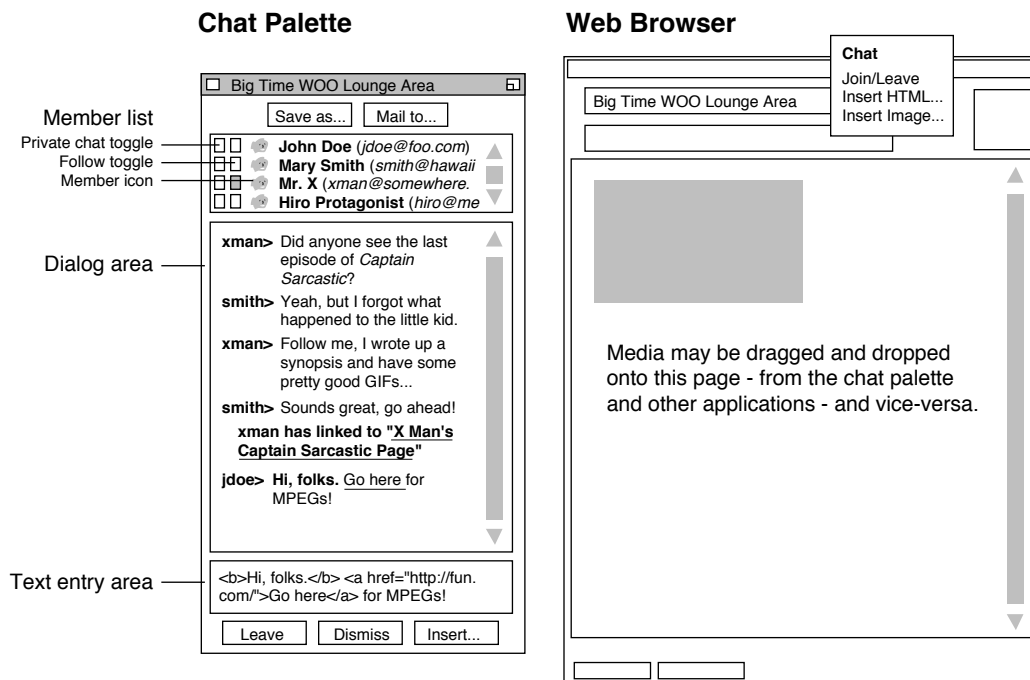
Real-time protocols and MBONE-related developments can be used. Much research can and should be done into seeing how asynchronous communication can be integrated into networked hypermedia systems. WOO systems may very well give users the ability to transmit HTML as well as teleconference, and solutions that have been proposed for use in the MBONE community can be applied to any manner of collaborative services today.

Users can have the ability to edit pages. We cannot call this system object-oriented if it is impossible to create, edit, and manipulate objects. It seems natural to couple chat abilities with a Web browser and WYSIWYG HTML editor, so that users can work on the creation of HTML documents and build environments that fit them. One way of editing Web pages is to be able to drag and drop things from the chat window to the HTML page and vice versa, so elements of conversation and pieces of offered media can be integrated. Likewise, things on Web pages can be incor-

porated into the conversation, making it easy to point out details, quotes, or pieces of media. Audio and video capture abilities can be integrated into this capability to allow a higher degree of interaction. For users without drag and drop capabilities, an Insert button in the chat window and Insert menu commands in the browser window may be provided to give users the ability to put different types of media into the conversation and HTML pages.

The idea is to allow people to create their own Web worlds together, represented by different servers, like islands in the network. Instead of telnetting to a MOO server and typing "go" and "look" commands to walk to and look at different areas, the new model allows people to link to a Web server and link to other pages and experience the media they hold. This new model is backwards-compatible with traditional MOO systems; anyone can create a Web server that holds nothing but text. But it is possible to extend things much further.

One of the real strengths and revolutionizing capabilities of such a WOO system as is described here is that much in the same way that the Web has made it very easy for people to create and distribute hypermedia, this system makes it very easy for anyone to create communities and capture their conversation. Where setting up an IRC or MOO server is not the easiest task, installing a WOO server can be made trivial and even automatic, and it opens up a number of future scenarios which are discussed in the next section.



8: Layout for a prototype WOO system. A chat menu may be provided in the Web browser, or the chat palette may be controlled by a separate add-on module or application. HTML can be typed in and media can be dragged and dropped into any area.

Visions Collaborative notebooks. WOO systems are ideal for projects in which changes and ideas are shared and discussed in groups. Imagine a scientist, Bernie, using such a system for designing a gear. Two colleagues from Italy and his coworker Jim converse on the Web server and offer small pictures of their proposed gears. Bernie drags one proposal onto a page; Jim writes a caption below it. As different pictures accumulate on the page, Bernie decides he likes some more than others and sorts the images in order of preference. Jim offers a link to a movie regarding gear structures that he recently found at a NASA database; everyone views it for a moment and writes their commentary on the page. Jim asks someone to follow him to the NASA site to help him with searching other material; they leave and come back with references to three other papers that may help them with particular problems. The scientists take a vote on the proposals and decide on one. Bernie makes a new page, labels it “Gear Proposal Page 2.0” and puts the image in it. The scientists make links to the related material they found underneath it and Bernie puts a GIF of his signature at the bottom. Whatever Bernie types into the chat area is displayed to the others as part of the conversation, unless it is preceded by a special character or sequence of characters that indicate to the server that it is a command. At this server, the character happens to be the @ symbol (there is no command entry box in his browser). So he types “@date” in the chat area, which generates the current date and time and Bernie drags this onto the page. The group splits up and Bernie appends a separate PDF version of the page to the project archives.

Virtual environments. DarkStar is a warrior knight visiting his favorite WOO - it has a home page with a huge castle gate on it; he supplies his login and password to the server and selects a link that lets him in. He enters the castle, which is portrayed by nicely rendered images of the interior that have been contributed by users, and a number of arrow icons give him navigation choices. He reads the text below, which offers a link to the castle map, which he goes to and prints out separately so he can use it for reference. He selects the local pub on the online map and is transported to the Castle Pub page, which has an image of the pub as if one were just outside the front door. DarkStar mouses down on the door, enters the lobby, and notices that a number of friends are here - his buddy Goliath is gambling.

DarkStar types, “What’s up, Goliath? I made a new sword...” He drags a picture of the sword from his tool box of images and drops it on the screen so everyone can see. He takes it back before anyone can grab it. Goliath replies, “I haven’t had so much luck at backgammon - can you loan me a few credits?” DarkStar types yes and then “@give Goliath 10 credits.” This tells the server to look at its database of players and their inventories and transfers 10 credits from DarkStar’s inventory to Goliath’s. Now when DarkStar types “@inventory”, the server tells him he has a nice suit of armor, some assorted items, and is 10 credits poorer. Goliath types thanks and goes to the backgammon page, where he and a

very good player have been betting on games for some time. There are a number of spectators that type idle talk to each other; they don't say anything of importance. The backgammon board is made of a table of squares of uniform size, and when a player wishes to make a move they drag a square with the appropriately colored piece onto an existing square, which is replaced by that image. There are many other games on other pages; but backgammon seems to be quite popular here.

Tours and lessons. A number of students gather at the Virtual Academy's home page, where they await their professor. Dr. Hoffman is a noted art historian and websurfer. He shows up and asks the students to follow him to an art site in France, where he will begin talking about impressionism and showing many examples.

The students follow him to a page with a number of pictures of artwork. He drags two into the conversation and talks about the particular colors that the artists used. He goes to a page that has a picture of a color wheel and talks about why the particular colors that were chosen work so well; they are complimentary colors. He asks to students to go off and find some good examples of artwork with similar color schemes; he gives a list of about half a dozen links to sites at which they can start searching. After about ten minutes or so they regroup and each one drags their offering onto the page. Dr. Hoffman talks about how each picture uses colors to achieve different effects and pastes some of his favorite quotes from artists regarding color onto the page. After the students have left, Dr. Hoffman decides to keep the lesson transcript in his folder of tours so he can remember it for when he writes an online textbook that he's working on.

Entertainment. *MediaDude* magazine, one of the hippest pre-teen digital publications, is hosting a show tonight, one that features a presentation by computer expert Phil Fargo and a concert by the DEC rock group Severe Tire Damage, and Max really wants to go. He gets to MediaDude's server and selects an image of their auditorium. The show is \$5; Max enters his Dad's Visa and gets into the page. Fargo is discussing the finer points about modem cards and he's putting up a number of slides as he continues. Max drags and copies the slides in a directory on his hard drive - he's not stupid. He sees a mistake on the slides and would like to put a note on the page correcting the error, but he doesn't have permission to do it and it's not such a big deal anyway. Since he was some minutes late, Max asks the server for the first five minutes of the talk and briefly reviews them.

After the lecture, there's a question and answer period, and after that the Severe Tire Damage logo appears on the page. They're about to play! Max connects to their CU-SeeMe reflector and the audio is a little choppy; he doesn't care about video since he's going over 14.4 from home and the band is putting up live color shots on the page every two minutes or so. Max downloads the band's screen saver and anxiously awaits the day when he can get the \$100 28.8 modem Radio Shack just advertised so he can

ask the band dumb questions and see their reactions (their interviews are always funny), but he has quite a few papers to deliver before his next paycheck.

Marketplaces. Jane wants to buy a VCR; she was thinking about going to Eworld, but she knows that they're a little too cheap and they don't offer the greatest customer support. Instead she links to the BargainBox Systems server, gets through the home page, and goes to the VCR section. She browses through a few pages' worth of machines, and decides on a Sony that looks good and is not too expensive. She wants to know the particulars about the remote, though, since she really needs a VCR with a remote, so she hits an icon marked "Service" which pings a salesman and brings him to the page she's on. "I was wondering about the remote that came with this" she types, "Do you have some specs on it?"

The salesman types "absolutely" and gives her a link to the specifications page and drags up a picture of the remote. "As you can see, it's quite fully featured. It's got a rubber grip for better handling ability too." Jane says, "It looks good. I'll take it!" She is sold on the Sony. She follows the salesman to a newly generated purchase form and he enters the model number of the Sony in the appropriate field. "That's it," he types, "if you need any more assistance, just let me know." He exits the server, leaving Jane to enter in her billing information, card number, and other mundane stuff. She knows that if there's a problem, she can always visit the troubleshooting page, the online manual, or the Sony VCR User Group (SVUG) mailing list.

Cyberspace is the Shadowland of human technos; at the end of our striving we will find the perfect, still mirror in which we see only ourselves.

2

Mark Pesce

Understanding Media: The End of Man

Cyberspace

Over the last decade the word *cyberspace* (cspace, pronounced *see-space*) has been used to refer to many different things, from the Internet to webspace to “wherever you go when you make a phone call”, and it has lost much of its original meaning due to media-driven misinformation. The original sense in which the word was used is probably fairly close to this definition: cyberspace is a globally networked computer-sustained real-time interactive environment in rendered three-dimensional form in which users can collaborate and view and manipulate information.

Typical visions of cyberspace consist of a collaborative system that embodies all of the successful features of the World-Wide Web and other community-friendly services. It has been proved many times over via systems such as traditional MOOs and the Web that people are ready and willing to learn, collaborate, and build in cyberspace. The bulk of the existing work now lies in the hands of developers and researchers, who must build a solid framework on which cspace can evolve. One of the best descriptions of cyberspace can be found in Michael Benedikt’s introduction to *Cyberspace: First Steps*:

Cyberspace: A word from the pen of William Gibson, science fiction writer, circa 1984. An unhappy word, perhaps, if it remains tied to the desperate, dystopic vision of the near future found in the pages of *Neuromancer* (1984) and *Count Zero* (1987) - visions of corporate hegemony and urban decay, of neural implants, of a life in paranoia and pain - but a word, in fact, that gives a name to a new stage, a new and irresistible development in the elaboration of human culture and business under the sign of technology.

Cyberspace: A new universe, a parallel universe created and sustained by the world’s computers and communications lines. A world in which the global traffic of knowledge, secrets, measurements, indicators, entertainments, and alter-human agency takes on form: sights, sounds, presences never seen on the surface of the earth blossoming in a vast electronic night.

Cyberspace: Accessed through any computer linked into the system; a place, one place, limitless; entered equally from a basement in Vancouver, a boat in Port-au-Prince, a cab in New York, a garage in Texas City, an apartment in Rome, an office in Hong Kong, a bar in Kyoto, a cafe in Kinshasa, a laboratory on the Moon.

Cyberspace: The tablet become a page become a screen become a world, a virtual world. Everywhere and nowhere, a place where nothing is forgotten and yet everything changes.

For the last few years, with the demise of VPL and related virtual reality ventures, cyberspace has been the object of jokes and dismissed as a useless venture, with people writing articles on the current "Post-VR" age and placing such technologies on their hype lists. But it has made something of a comeback now that global networking and interactive multimedia have been gradually permeating the public consciousness. People may still not grasp the idea of what it is exactly, but they have a feeling how big it is, they understand that it now works, and they believe that to have an email address is to be "in the know."

The time is ripe and the moment is now to begin thinking seriously about cspace systems again: how they can be implemented, what new kinds of business may arise from them, and what sorts of architectures, communities, and cultures will develop. And we can apply these ideas to so many real-world problems, it is only a matter of time before they are taken and absorbed by the global community.

In the same way that the Web encompasses Gopher and FTP, a true cspace system can encompass many different protocols: MBONE-related protocols, HTTP, NNTP, and so on. When there is a basic language that can represent many systems, we can create new tools within the meta-system. For instance, I can create a program with a graphic interface using X Windows on a UNIX workstation or a mail reader with an HTML interface on the Web, but I can create a window manager, an operating system, a spreadsheet, a compiler, and a videoconferencing suite with (as an example) a VRML interface in cyberspace. Some people are talking today about being able to build Web browsers in which text and other elements are represented by simple vector graphics, but it's possible to go much farther than that. We have the power to make new structures, new environments, and new services.

2.1 A New Definition

Use the word "cybernetics," Norbert, because nobody knows what it means. This will always put you at an advantage in arguments."

Claude Shannon

In a letter to Norbert Wiener, 1940's

The word *cyberspace* descends from *cybernetics*, a phrase coined in 1961 by mathematician Norbert Wiener in his seminal book *Cybernetics; or, Control and Communication in the Animal and the Machine*. Wiener chose the Greek word for "steersman" - *kybernetes*, or navigator - to name this new science. The prefix *cyber* means "rudder" or "governor"; it is a thing that controls.

Cybernetics deals with the study of systems that make use of feedback and with the study of how people construct models of these systems. It is a general field whose scope is not very well defined, but the key point to remember is that it deals with control and feedback. The navigator of a ship takes control of the system

and receives feedback from conditions such as the wind speed and direction, the ship's speed and direction, and so on. The rudder controls the direction of travel. Anything "cyber" implies interactivity and control of some sort; cyberspace does not deserve the name if it is a one-way medium; if one has no control over it. Passive media is not cyber- anything.

- Cybernetic resources can be found at <http://schubert.media.mit.edu/Cybernetics.html>

Jaron Lanier came up with the phrase *virtual reality* (VR) in the late 1980's, probably borrowing from Myron Krueger's mid-70's term *artificial reality*, to characterize the work that VPL was doing then; this amounted to creating immersive computer-assisted environments utilizing three-dimensional graphics technology. In one sense the goal of VR is to create an artificial environment that appears as real as possible, whether or not the environment is a fantasy setting or a real-world duplicate.

The popular media has brought about many meanings of what VR actually stands for and in many cases has made incorrect assumptions and predictions of what the technology can do. Because of this fact, a better term to use today when referring to a VR application is *virtual environment system* (VES).

Those who create virtual environment systems work within a broad field - a virtual environment system may be one that tries to mirror reality, or it may be any type of computer-assisted interface that makes use of three-dimensional graphics or display technology that imparts real-world knowledge. Virtual environments are basically computer simulations that convey information to and from our various senses.

Gibson originally portrayed cyberspace as a world-wide networked virtual environment in which people could move, talk, manipulate objects, and navigate via direct connections to the brain. Today cyberspace more commonly refers to the networked, multi-participatory virtual environments that are being built today on the Internet, without any regard to how people are hooked to the system. This is what I call the first sense of cyberspace.

2.2 Encapsulation

Marshall McLuhan wrote that new media tend to encapsulate older media. That is, old media can be represented within the space of new media. The alphabet as a medium was encapsulated by print; radio as a medium was encapsulated by television. I can send a radio broadcast via television and could even use my television solely as a radio if I felt like it (although it may not always be practical).

What many people have failed to understand is that the computer as a medium has the potential to encapsulate all prior media. Therefore, the interface and working environment of the future will be able to handle or simulate all of the various forms of media. It will probably be networked, and it will probably incorporate three-dimensional graphics to some degree. The medium of

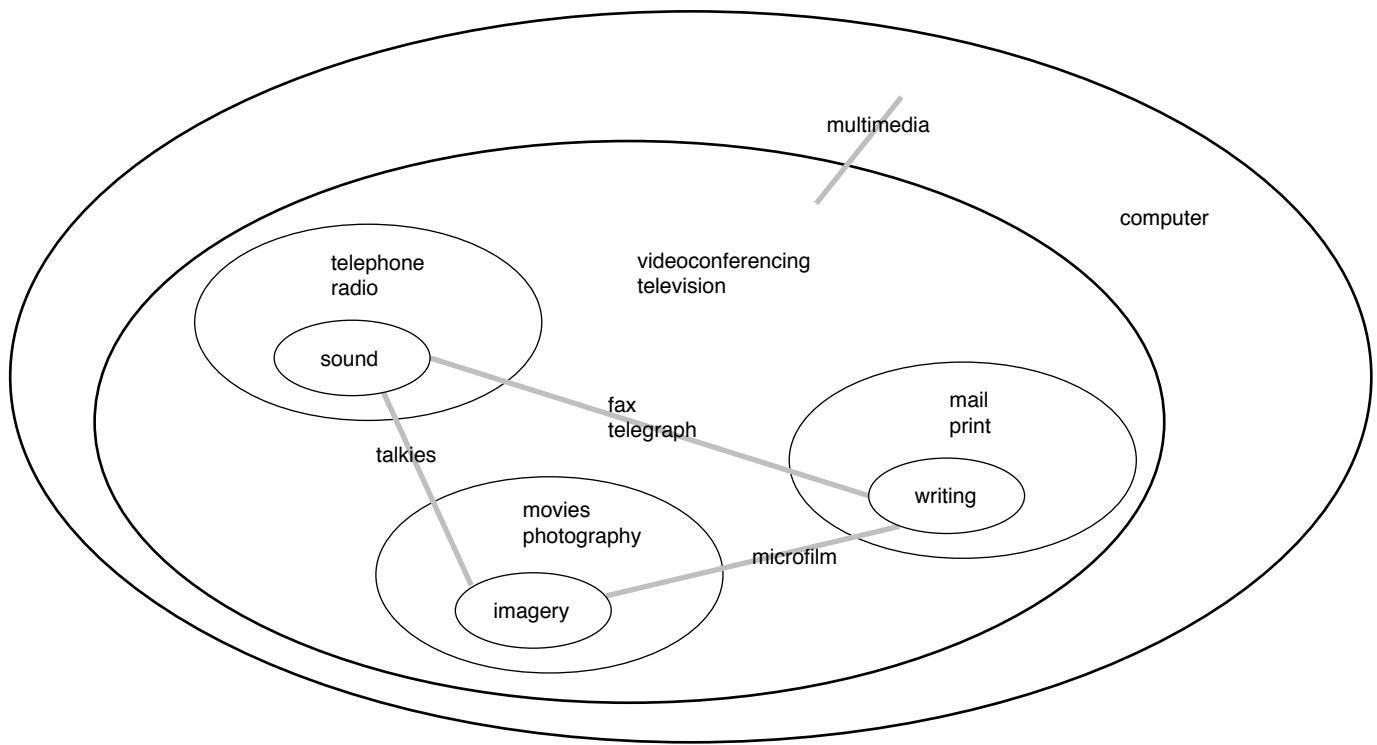
tomorrow is essentially cyberspace (in the first sense) coupled with every medium known to humankind. It is a top-level, holistic view. Thus the second sense of cyberspace is *an interconnected computer-mediated environment in which all prior media are represented*.

Self-contained networks that contain only one form of media or a certain subset of all media, such as the phone system or cable television system, are not a part of cyberspace using this definition. Interconnected in this sense means always connected, or the ability to be continually connected. The Internet is not a subset of cyberspace. Rather, cyberspace uses the Internet as its main method of transmission, just as television is carried on the cable system.


There is not necessarily one cyberspace either - people may create their own networks, their own particular cyberspaces. It is not easy to create your own television network, but it will be trivial to make separate cyberspaces consisting of particular people and environments that others can connect to. The sum total of all these cyberspaces is cyberspace. It grows, branches and divides as fractals do.

In this new environment, the use of the word "multimedia" should cease to exist. Multimedia, an overused word, has put the emphasis on the separation of media, the *multiness* of media. We will see environments that have no concept of different media as separate entities, ones that will see all media as a whole, allowing everything to be created and manipulated in the same fashion. In the end, it will all be simply bits or a similar uniform format. Before the age of the personal computer, multimedia works referred to non-electronic media - art works made of multiple materials. Why we would choose such a word as a basis for the integrated digital systems of the near future is mind-boggling.

Below is an encapsulation diagram of cyberspace that shows the encapsulation of current and prior media and the various media that have formed among them.



9a: Encapsulation diagram of cyberspace.



9b: Simplified encapsulation diagram of Gibson's original view of cyberspace.

The faint lines interconnecting the different media represent what McLuhan called *cross-pollenization* or *hybridization* - two or more media combining to create new forms. These are areas in which new hybrid media can arise.

For many media there is a broadcast as well as a one-to-one interactive version. Print (in this definition) is broadcast; mail is the interactive form. Fax is essentially two-way telegraphy, radio combined with print.

For the purposes of this diagram, there are three basic media: sound, imagery, and writing. This is how many computer systems view the world today, although writing is really a constrained type

of imagery. Writing systems are simply subsets of imagery, but many forms of media today tend to treat imagery and text as completely separate entities.

This diagram succinctly points out why such hybrid media as fax and interactive television are backwards; they will already be encapsulated by the medium of the computer, which has the rare capability of being able to transmit prior media in a broadcast as well as in an interactive (one-to-one) fashion. Forcing new media to do the work of the old is “rear-view mirror thinking” - another McLuhan term - and is most commonly practiced by those entrenched in the former media. The most suitable thing to do is to create fresh forms of communication that are appropriate for the new media so that they can be used to their maximum potential. Using the television as a radio or the computer as a fax machine is not an efficient use of the medium. Going the other way, forcing old media to do the work of the new, is backwards thinking and makes little sense. It may be due to people misunderstanding the new media, thinking that they can use the old media more effectively. It is not the easiest thing to make a movie out of photographs or a computer out of a television.

As a medium encapsulates previous forms, it also informs us of prior media. Thus we see television advertisements for books and albums and Web sites displaying advertisements for television and movies. This is one of the sure signs of encapsulation, when a medium is deemed worthy enough to notify people of prior media.

In the past communications have been categorized by the devices they are transmitted on or through, hence phrases such as “record album”, “TV show”, and “compact disc”. It is not forward thinking to name the media after the physical form, as in the case of CD-ROM or CD-I. The practice lends itself to obsolescence. When the computer homogenizes these formats into a single digital signal, new terms are bound to appear. Certainly all of these forms will fuse together and mutate in interesting ways, and may often end up being indistinguishable from one another. Cyberspace will serve as a medium for them all.

It is possible to come up with another term to refer to the computer-mediated environment that encapsulates all media; some have argued that the World-Wide Web is appropriate. But many people today refer to such environments as “space”, something to *go* to; many use the word “cyberspace” in that sense already, and the World-Wide Web has become mentally attached in the public mind to a specific kind of environment using a specific protocol. Perhaps in time the terminology will change, and it could be that a new word is needed that properly describes the concept.

We might then refer to a *metaspace*, a word that may be more appropriate. It implies nothing about control, which is more correct, and it properly conveys a sense of encapsulation as well as a sense of place.

2.3 Notes on Successful Terminology

When a new medium is created, two or more media combine to form a hybrid, or new fields relating to new media appear, there is always the question of deciding a name that can be used to refer to these things. Here are a few guidelines people can follow to create lasting and successful terminology.

- Refer to the known
- Keep it short
- Allow for abbreviations
- Make it easily pronounceable

These guidelines arise from the fact that people will probably be referring to the word often in the course of daily life. They must be able to form a quick mental picture, or at least understand a part of it, when they first encounter the word. This allows the idea to be transmitted and spread quickly.

The names of many new media contains Grecian or Latin parts - *television*, *cyberspace*. This is very much due to the tradition of scientific nomenclature, and it is fine if an appropriate suffix, prefix, or modifier can be found. However, a completely foreign name will not lend itself to being remembered: the name must refer to known elements that refer to the common.

The best way to do this is to incorporate an entire common word within the new name or phrase. This causes instant recognition; the known part also allows people to understand its relation to the unknown part. Thus the successes of the words *hyper~~text~~*, *Inter~~net~~*, *multimedia*. We understand the known words and they give light to the unknown. Failed words often do not incorporate words that are common: *fantics* (Ted Nelson's name for a kind of media design), *matrix* (used to refer to the Internet) and *blendo* (used to refer to multimedia) are good examples. Exceptions arise when the media are so completely different, there is little or no precedent; when the concept is easy to understand, when the public is particularly accepting, or when a great deal of marketing has been done to introduce the word: thus the words *synergy*, *radio* and *robot* have come into being.

The word or phrase should be short, ideally no more than two words. This is why phrases such as *virtual environment system* is unlikely to become popular, despite its accuracy, and why *radiophotography* is not something you are likely to recognize. One can also anticipate the abbreviations of the phrase as the new media becomes commonplace: witness how such words as *telly*, *tube*, *stereo*, and *phone* as abbreviations have become popular.

The name should also be easily pronounceable - it will sound right and float easily off the tongue. *Fax* is short and easy to pronounce, *kinetoscope* and *phonofilm* sound awkward.

One should stay away from creating sensational names for new media. They may become popular, but in the end attract the kind of attention and associations that all sensational things do. Partic-



ular examples are *virtual reality*, *cyber* as an indiscriminate prefix, and *teledildonics*. And care must be taken to choose names that describe the media somewhat accurately, lest we end up with such mental disasters as *horseless carriage* or *information superhighway*.

In summary, here are a few of the top-level terms in common use today:

- Internet (the net) - the global collective of interconnected computer networks, typically those that communicate using TCP/IP.
- World-Wide Web (the Web) - the initiative to create a universal, hypermedia-based method of access to information. It is also a client-server based service that runs over the Internet. Also called webspaces in another sense.
- Webspaces - generally referred to as the collection of all known Web servers and the information and services available through them.
- Virtual environment system (VES) - an interactive computer simulation that typically imparts information visually, aurally, and via force.
- Virtual reality (VR) - in one sense, a field that makes use of virtual environments that contain behavioral, visual, and other sensory similarities to the real world.
- Cyberspace (cspace) - the purist may regard cyberspace as incorrectly referring to the Internet, particular services available on the Internet, webspaces, or the domain of electronic communications in general. This work refers to two meanings of the word: 1) a world-wide networked,

multiparticipatory virtual environment, and 2) an interconnected computer-mediated environment in which all media are represented.

10: CutPlane, a full 3D desktop environment, is being made as part of the DesignSpace project at Stanford's Center for Design Research. The DesignSpace page is at <http://cdr.stanford.edu/html/DesignSpace/home.html>.



11: The Conceptual Design Space (CDS) project at Georgia Tech is a tool for creating designs in a virtual environment. Shown is part of the CDS 3D interface.

- Virtual reality resources are at <http://www.lightside.com/~dani/cgi/virtual-reality-index.html>
- More resources are at http://www.yahoo.com/yahoo/Entertainment/Virtual_Reality
- Much virtual environment research is being done at Georgia Tech's Graphics Visualization and Usability Center, <http://www.cc.gatech.edu/gvu/virtual/VirtualEnvironments.html>
- The Knowledge Base project at the University of Washington's Human Interface Technology Laboratory (HITL) has a great deal of pointers to current information and research, http://www.hitl.washington.edu/projects/knowledge_base/

2.4

Myths

There are I believe a number of myths surrounding virtual environment technology, partly due to its beginnings and partly due to misinformation spread through an unaware mass media. Here are a few of the most common ones:

- You need expensive and hard to get hardware and software.

For a long time VR research has been undertaken by scientists wearing dozens of thousands of dollars' worth of equipment connected to \$500,000 mainframes, and this image has been propa-

gated wildly among the public. Rendering technology appeared in upright video games in the 1980's, gaining some attention, and movies such as *Tron* introduced the concept as being something only an expensive special effects budget could afford, but at the time, this was mostly true.

Compelling games like *Stellar 7* (for the Apple II) introduced the possibility of having 3D graphics technology in software without having to spend a bundle; ten years after the game was introduced, Sega was selling \$100 Power Gloves for their game system and one could run shareware VR real-time world viewers (which incorporated color and shading) on a 386-based machine.

You do not need a *head-mounted display* (HMD) to believe you are interacting within a world. It's possible to read books and do this. HMDs can enhance the experience, but until such hardware technology becomes superior to on-screen (flat panel) display technology, one has to wait. The low-end of technology is continually getting higher; the text-only displays and 286 PCs that gather dust in classrooms today will be replaced by obsoleted graphic terminals and color Macintoshes tomorrow. The computer that cyberspace will not be able to reach is the hardware that can't network to anything else. Cyberspace is by definition communicative. Anything else must be called by another name.

We can no longer name hardware as being the bottleneck in the process of reaching the goals of cyberspace development. One major bottleneck is education - of children, businessmen, and politicians - and will continue to be for some time. The older generations must be educated so that they do not misunderstand what is happening and prevent or stifle the great changes that are in progress today. The education of the members of the coming generation is important as well - they may know how to use and experience the new media better than anybody else, but do they know what it means and what they can do with it? A \$5,000 computer is ultimately worth nothing if a teacher can't motivate others to use it to reach their potential. Schools must learn to recruit the best and train them - once the trainers realize that they are equipped with powerful tools, once they know what it can do and become familiar with it, only then can computers become part of the teaching naturally. Computers are the chalkboard of the next century and should be treated so.

- You need high-end graphics.


Digital imagery is mostly the art of illusion, and cyberspace is the biggest illusion of all. The idea is to provide a system that gives users a sense of place so that they may interact and feel comfortable with it. There is no need to make exquisitely detailed and rendered objects - as long as the world is consistent with its rules, it becomes possible for people to believe in it and learn from it.

The lessons of *Habitat* - a multiuser system for the Commodore 64 - and *SIMNET* - a military multiuser tank simulation environment that at first consisted of little more than vector graphics - showed that it is possible for people to believe in and interact

within a world even if it does not look realistic. Users in Habitat were represented more like blocky cartoon characters, yet discussions, meetings, and games all took place in this environment. For many, the system was engaging and compelling, and some would stay online for many hours at a time. The bottom line is that fancy graphics are not necessary - in the creation of today's systems the majority of the work should go into developing efficient protocols and languages to enable fast and believable interaction and into developing a framework that allows users to easily navigate through and build worlds that are consistent in their laws.

As microprocessors get faster and it becomes easier to write sophisticated graphics routines (via the help of libraries or programs), there is no doubt that future worlds will begin to look more and more realistic. It is reasonable to guess that 80 million textured, Gouraud-shaded, anti-aliased polygons per second, one definition of visual reality, will be achieved within a few decades. Today virtual environment imagery is in the middle of the realism spectrum - objects don't exactly look real, yet they don't look exactly fake either. In time, computer graphics will evolve so that one can make worlds look as hyperrealistic or as cartoonish, pixellated, and impressionistic as one likes. At this point, it is not a serious issue.

- Cyberspace will be stored in one place.



12: Cyberspace will most likely be distributed. It may consist of a collection of connected worlds that can be travelled through, each world being managed by its own server.

Many people have the idea that cyberspace will consist of one gigantic virtual area stored in one machine somewhere on the planet. This likely comes from the timesharing, mainframe-centric point of view of the early 60's and 70's, the time in which visions of computerized global media systems consisted of Asimov's UNIVAC stories and the first descriptions of Ted Nelson's Xanadu. It was thought that the digital library of the future would consist of a massive, state-run computer which people could connect to and explore around. Today, of course, the client-server model is now the popular paradigm, and the Web has shown that it is at least possible and desired to create globally distributed systems. Cspace will not be owned by any one organization or zoned by a central authority - rather, it seems that it will consist of many individual worlds and scenes linked together, any part of which may be stored on a separate machine somewhere in the world.

This is not to say that cspace will not be pieced out like real estate - it will be desirable for organizations to create directories of worlds to make sense of the Internet, much in the same way that the InterNIC has functioned. Individuals as well as access providers will have local control of their own worlds, making their own rules and doling out parts to others who wish to collaborate with them. It will probably be desirable to follow a number of general guidelines in doing so (so that cspace navigation does not become the equivalent of roaming around towns in the Old West,

each with its own sheriff), but it may take some degree of cooperation between individuals to create rich and varied world directories, unless the process can be automated.

- Xanadu information is at <http://sensemedia.net/xanadu/>

2.5 Why Virtual Environments?

Cyberspace will most likely use some sort of unifying environment as its virtual, software-based interface. Whether or not this framework utilizes three-dimensional graphics to a heavy amount is not necessarily important; what is important is that we be able to work within it fluidly and be able to suitably experience and manipulate the various media we encounter there. Here I make the case for a networked virtual environment that uses real-world knowledge and 3D to some degree in order to navigate through and visualize data. Why should we use virtual environment technology to build this framework?

Because the world is information, and we can navigate it.

There is so much information online today, it's hard to imagine the amount and size of things that will exist in digital storage tomorrow. We must be able to create new tools to visualize, search, and manipulate information. Virtual environments provide a way to combine the best features of real-world information navigation - memory of places and visual cues - with the best features of online navigation - fast searches and sorting and quick cross-referencing. By being able to manipulate data the same way we manipulate objects and ideas in real life, we can spend more time creating and understanding things and less time figuring out how to operate them. A well-designed environment can help us experience things in a larger context and pick out useful information from complex systems. Environments that mirror the physical world can show us exactly what is going on and keep us informed.

Because reality is experience, and we can learn from it.

Experience is by far the best teacher. Today, one has to learn an incredible amount of unique skills in order to use a computer effectively; typing is not a natural form of communication and one must learn what buttons do and how dialog boxes work. Because virtual environments can mirror the effects of reality, it is possible to make systems and objects that we can learn how to use as effectively as we can learn things in real life. By creating realities, it is possible to do away with many aspects of interfaces as we know them today - information becomes as close to the "real thing" as possible, and therefore we can come closer to understanding concepts than before. What all of this entails is that the best virtual designers in cspace may have a great ability to design objects in the real world as well as the virtual world. And users and developers must be aware of privacy issues in cspace because it offers the potential ability to extract information from, influence, and inform the mind in the same way reality does. To better understand these issues, we can think of television as a

primitive read-only cyberspace. To some extent, TV can influence and inform, but it can also trivialize and misinform. We must be careful in designing mass environments.

Because societies are human, and we can communicate in it.

People need to be with and communicate with others. At least, it is the most natural thing for people to do, as humans have been stuck with others since their beginning. Collaborative virtual environments provide gathering grounds for new communities and types of interactions, and they give people a voice like they have never had before. It has been written that Andy Warhol was far off the mark when he said that everyone will be famous for fifteen minutes - indeed, we may all be famous for days today, years tomorrow, and perhaps one day eternity. Groups and clubs give us purpose to do things and puts life into the machine. We can share experiences and visions and learn to understand the other person's point of view. We can attend concerts, act in plays, and attend classes with an international audience. As long as people have something to say to somebody else, they can say it online.

When it is possible to do so many things online, some people will choose to stay wired most of the time, and some businesses may even require employees to do so. Cyberspace can be dangerously addictive - many people today work in front of their computers around the clock. The late 20th century has witnessed the birth of the phrase "couch potato." But as computers become smaller and more ubiquitous, and display devices become better, faster, and more environmentally and biologically sturdy and safe, people can learn to use them in new ways, allowing them to do more and communicate more efficiently in the real world. There is no doubt that people will do more in the future and they will have more time to do it in. What they do with that time, though, is another matter. Computers have not given people more leisure time; today computers are the leisure time. Society must grow a set of mores that allow people to integrate computers into their lives in a healthy way as well as take care of their physical selves and the physical environment. Computer addiction and the idea that the online world will make reality more plastic is a misperceived argument against cyberspace; rather, it is up to society and cyberspace developers to make it a useful part of life and the real world one worth living in.

- | | | |
|-----|--------------|--|
| 2.6 | Applications | The following section encompasses a number of ideas related to current and potential applications of virtual environment technology, some of which are a decade old, others which are relatively recent. They are written with current technology in mind; any of these ideas and scenarios could be implemented with the proper tools. All it takes is support, understanding, and a little vision. |
| | Architecture | Architecture and urban planning are obvious fields in which virtual environments can be used. One can use environments for initial trials, before structures are built, using simulation technology; they can be used after the structures are already built, guid- |

ing people to their destinations or providing control over systems as well as helping with interior design; and they can be used after the structures are gone, providing historical information.

CAD systems have been used in construction for many years now; by using environments one can complement the benefits of CAD by allowing walkthroughs, acoustical evaluation, structural damage simulation, and so on. Environments been considered in helping to rebuild areas such as South Central L.A. and East Berlin.

Architectural education and history can benefit greatly. One can see cities building and destroying themselves over time; construction companies can check records to find pipes and lines underground and dig to the appropriate depth, and virtual environments can be used in any manner of GIS applications.

Architects will be able to explore the outer boundaries of the field; they can treat worlds as ideas. In cyberspace one can make a Japanese garden or temple, spaces to contemplate upon. In doing so, architects can give us natural settings to contemplate our own creations, what it means to be human. We can make environments art, architecture that is not meant to be lived in but to be experienced. In the real world, it takes years to make buildings; in many cases their social impact is not felt for a very long time. Cspace will bring about the development of a kind of architecture that makes an immediate social statement, architecture that can come and go overnight and respond to current events. The public will eventually be able to see architects more as creative performers and less as structural engineers.

Art

Besides the obvious applications of virtual museums and art exhibits, the possibility exists to create new forms of art, virtual sculpture perhaps, or new forms of art that are like traditional computer animation but can interact with onlookers in ways that cannot be done in the real world. There has been a good deal of talk about the creation of 3D clip art; companies such as Virtus sell relatively cheap libraries while more professional datasets can cost hundreds of dollars per image or more, depending on the complexity.

We can imagine 3D art programs in which we can paint creations viewable in three or more dimensions, with strange geometry and physics never before experienced. One can imagine creating landscapes of new organisms that act and react to themselves and the viewers.

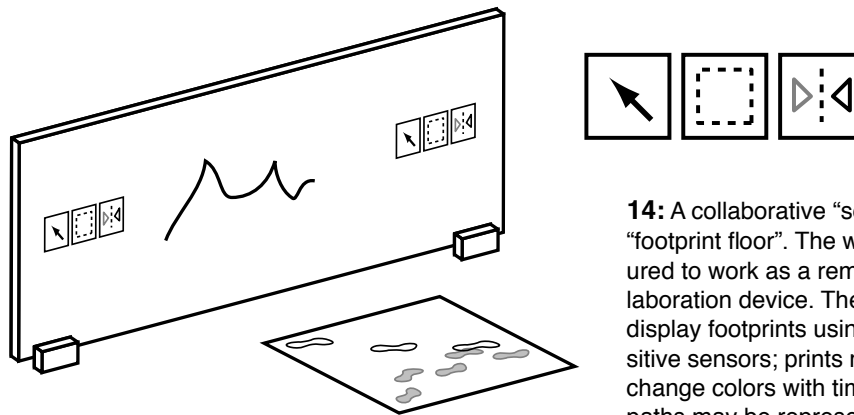
Exhibits, festivals, and cultural exhibits are all good candidates for cyberspace; although live video is not exactly widespread, there are many companies that may create animations, advertisements, and previews of movies and other media events. Applications to the real estate and (interior, exterior, landscape, industrial) design fields are natural. Anything which requires a "what-if" or "what does it look like" scenario can be fairly well satisfied with multiplatform rendering technology now available.




13: A 3D representation of Hilbert space by Nelson Minar. His 3D art gallery is one of many on the Internet.

- Virtually Reality, a 3D comic, is at <http://www.onramp.net/~scroger/vreality.html>
- Ray-traced MPEGs are at <http://www.santafe.edu/~nelson/mpeg/>
- 3D experimental art is at <http://www.uio.no/~mwatz/art/gallery/>

There is a great potential for art to be used as a social device in this medium, especially when coupled with unique display devices. We can take the “Media Spaces” and other collaborative video work done at Xerox PARC and the NTT Human Interface Laboratories as a basis for thinking about new types of installations that can be created using digital video coupled with virtual environments and large displays.



14: A collaborative “scribble wall” and “footprint floor”. The wall can be configured to work as a remote or local collaboration device. The floor may display footprints using pressure-sensitive sensors; prints may fade out and change colors with time. Contiguous paths may be represented in different ways.



15: A shot of Bar Code Hotel, a multiuser exhibit that allows people to control objects by swiping hand-held devices over bar codes. The installation is in a room whose surfaces are covered with codes. At <http://www.portola.com/PEOPLE/PERRY/BarCodeHotel/>

Using wall-sized displays, akin to Xerox’s LiveBoard, one can create a collaborative touch-sensitive drawing device. Used locally, the wall may be semi-transparent, allowing users on both sides to draw and see each other. Interface elements of some kind can be located in different areas of the display to allow participants to select and use different tools. Used remotely, a remote video display can be shown on one side, allowing users in different places to interact as if they were physically on both sides of the wall. In the case of public installations, certain tools may be restricted (such as an eraser tool) to allow and encourage different types of interaction.

Using this type of technology on floors gives rise to devices such as footprint trackers that may be able to give one a sense of social patterns and diversity among participants. In a museum, such an installation can track and record the movements of various onlookers, indicating popular works. “Ghost” footprint trails can be made by playing back previous recordings. Different displays can appear depending on where one goes, whether he or she runs or walks, and so on.

Such devices would work particularly well in social settings such as museums or clubs. In a club the floor may be covered with such devices, simulating things such as schools of Japanese koi that group around tables or nibble at the feet of those on the dance floor. Tabletops may consist of multiple displays, allowing people at the same table to converse, eat, and use their computers remotely. Devices such as trackpads or virtual keyboards can be used as I/O elements on the flat surface, or I/O ports may be provided elsewhere on the table so that people can bring in and connect their own equipment - head mounted displays, storage devices, art tablets, and so on.

Business Tourism may benefit particularly; if I want to go to Fiji, I can explore the place beforehand so I am familiar with the territory. I go to different travel agencies and experience multimedia exhibits (i.e. good looking environments) and see step-by-step with a first-person view all the different places I'll be going to and the things I'll need. I go to Fiji's tourist bureau and gather needed resources (a map, currency exchange rate calculations, I can glance up and see what the current weather in Fiji is like), visit a few restaurants there and make reservations. I tell someone at one cafe that I have a particular coupon from Compass Point Travel; she makes a note of it and will be expecting me.

We can imagine storefronts being even more garish and awkward than they are on the Web; most small businesses will have no concept of good 3D modelling or interface design. But one can visualize newsstands and shelves of information: pick up a software box on a shelf and go to a page describing it with links to reviews, demonstration versions and a handy purchase form. Enter in your information, the box opens, you reach in, and the real thing is transferred to your machine so you can install it (or perhaps an online installation kit guides you through the process).

Snow Crash introduced the interesting concept of carrying virtual business cards - hypercards - that represent data. When you take someone's card, this represents a transaction of data from their machine to yours. So when you insert your virtual credit card into a virtual ATM machine, you can do all the normal things you do with ATMs, and when a merchant takes your card, your account is debited accordingly. Right now these things can be done more easily with secure protocols and forms via the Web, but using the card metaphor may make transactions more natural.


Since stores in cyberspace don't have to be any particular size or dimension (they don't even have to look anything like a building), it will be tempting for marketing departments to go wild in this area. But it will be useful to mirror the real world, so that should users go to, say, a Sears, they will know how to get around, having taken a tour and printed out a list of things to buy sorted in a way so that only one pass around the store needs to be made. Airports may find it useful to put such technology in kiosks so travellers know exactly how to get to their gates, and using an online business directory customers would be able to pick out the

store in a crowded mall or off a freeway. Virtual environments can also be particularly good for auctions that may take remote bidders who would like to see and get as much information about the products as possible.

The temptation to sense or see something popular that happened somewhere else is a human trait newspapers and lawyers take advantage of every day: in a plane crash story, articles may show PostScript illustrations, but airlines may also offer animations generated from the plane's black box information so that online papers can make them available. In the courtroom, computer animation technology has been used many times, and we will likely see much use of networked hypermedia to show presentations with links going into more details, offering different scenarios, linking to precedents and transcripts, and so on.

During boardroom presentations and reports, the ability to travel to online sources and gather and visualize real-time statistics is convenient and can make quite an impact. A video of Starfire, Sun's conceptual future interface designed by Bruce Tognazzini (the original Apple interface guru) described exactly this scenario, in which a businesswoman had an edge over a competing sales proposal by being able to gather and visualize online data quickly. Starfire is basically a cyberspace system. One portion of the video portrayed Starfire mapped onto a very large, slightly concave, touch-sensitive display area.

16: A desktop view of a display device using Starfire, Sun's conceptual interface of the future. The wide wrap-around screen is touch sensitive; paper-based media can be scanned by lying them flat on the display's surface.



The ability to find business resources can be augmented - if I'm looking for a restaurant, I might go to a third-party online directory and search using certain criteria: I want Japanese, under \$10, within five blocks, three stars (out of five) or better. I see a map with the candidates glowing; the buildings may be larger in proportion to their rating or price, whatever I choose. I select one building and the map appears with a line going from my home to this restaurant so I can print it out if I'm the type that often gets lost. I go to user-contributed reviews of the place and see if it's worth it. The reviews look pretty good, as does the images of the daily special, so I read the menu and order in advance. I'm in a rush.

Computers It would be interesting to see complex programs represented as three dimensional objects; certainly computer visualization is a field in itself. Three-dimensional cellular automata and evolving, moving, self-modifying programs are not new ideas (Karl Sims' work at Thinking Machines is a good example). In the area of microchip design and collaborative CAD work, at least two dimensional image work is necessary; the ability to select areas of a design and zoom in or out and inspect in detail a design in three dimensions supports a greater sense of the finished product. It is possible to create more intuitive interfaces to machining tools, factory equipment, and any machine with real-world movement that may require remote operation.

- Karl Sims' virtual creatures are at <ftp://think.com/users/karl/welcome.html>

Computer administration can involve a lot of numbers: traffic analysis, server loads, CPU usages, packet loss statistics, outgoing files, times, dates, and so on. In this job, any tool that can make analysis of data easier, particularly over a networked system, is welcome. Charts generated on the fly from Web logs, maps with locations of visitors, and lines representing heavy interdepartmental transactions can help one study information flow and reroute areas to make things more efficient. Popular programs such as Purify and Quantify make use of two-dimensional tree graphs to analyze performance time and memory leakage in software. A program called Etherman shows all machines on a network as points on the circumference of a circle; lines of varying thicknesses between these points can represent packets or many other variables, and the program helps pinpoint and monitor problems well. In cspace new kinds of three-dimensional graphs can be created and manipulated to make troubleshooting even easier.

Economy As money becomes increasingly digital, the ability to keep track of the numbers will become increasingly important. Imagine being able to create economic simulations, in which graphs and structures mutate with time; these are ideal ways to teach basic concepts. A powerful tool for any stockbroker would be a virtual trading area - it is notoriously difficult in many countries to sell and buy stocks in person, but the ability to do so remotely even on a one-to-one basis with investors may make things easier. Being a stocks analyst, I may see multiple streams of stock quotes flowing around my environment, and if I'm interested in analyzing a particular stream, I can grab it and position it so it flows in front of my view, perhaps enlarge the symbols or activate a layer on top of these symbols that translates their meanings. I can manipulate snippets by using a scissors-like tool and drop them into my trading program so I can do business.

Education

Everything is interesting, until ruined for us.

Ted Nelson

Computer Lib

Edutainment is a huge business today, and cyberspace opens up a wealth of possibilities - education is quite high on the list of the possible cspace applications that people think of first, perhaps because they understand that it has the possibility to teach everyone anywhere in a way that can incorporate few of the restraints and failures of the traditional methods of teaching.

First, we need to understand why education today is not as optimal as it could be, and then see how we can use hypermedia collaborative systems to improve things. In the late 1960's a professor named Jerry Farber wrote a collection of essays and stories on education and titled it *The Student as Nigger*. He reasoned that the classic classroom model of learning and evaluation was restrictive, unhealthy, unfair, and ineffective; his book became quite popular among university populations and remained the center of much controversy for some years.

What is overlooked by many in reading the book's main essay, *The Student and Society*, is the fact that Farber saw and understood the potential of technology to fundamentally change the way that people learn. He stated that when it becomes possible to learn just as well at home with as good or better resources than are available at compulsory public schools, then the responsibility for education becomes less parental and federal and begins to lie more with the individual.

Today we can merge relatively new institutions (summer extensions, adult education, autonomous virtual colleges such as the Globewide Network Academy) and create new ways of teaching that do not force students to keep a particular train of thought in a particular place (facing forward, sitting down, keeping silent) at a particular time (between bells, on weekdays). As Farber wrote, learning is not something that is done to you. As work-study programs and new educational models have shown, experience is by far the best teacher, and online collaboration can at least offer something very close.

The fact that networked hypermedia makes it possible to learn at one's will and frees students by making exploration fast and fun is obvious; there are many indications of that already with the many successes of the Global Schoolhouse and other K-12 institutions that have connected and communicated online. What is not so obvious is the fact that it gives teachers the freedom to teach the way they want, where they want. They will have so many new ways in which to express and communicate their ideas, the very possibilities of what a learning environment may look like in cyberspace is mind-boggling.

For the next few passages I will quote Farber and present some ideas about education and how they may relate to cyberspace:

Consider how most classrooms are set up. Everyone is turned toward the teacher and away from his classmates. You can't see the faces of those in front of you; you have to twist your neck to see the persons behind you. Frequently, seats are bolted to the floor or fastened together in rigid rows. This classroom, like the grading system, isolates students from each other and makes them passive receptacles. All the action, it implies, is at the front of the room.

There is no reason to hold classes in anything resembling a room any more. Students may have a view on the lecturer, but they may be anywhere, facing anything, running a dozen other programs at once. Certain structures, such as administration centers, libraries, labs, and so forth should be preserved in that they provide a central place to store related resources so they can be found easily; in that sense a virtual schoolhouse may be nothing more than an informal gathering place. But wouldn't you rather learn at the bottom of a teeming coral reef, on the floor of the Notre Dame, in the center of a spinning galaxy? I believe that these environments will not be distracting but will instead be stimulating, encouraging, and engaging.

The scariest thing about a classroom is that it acts as a sort of psychological switch. You walk into a classroom; some things switch on in you and others switch off. All sorts of weird unreal things start to happen. Any teacher who has tried simply to be real in a classroom knows what I'm talking about. This is so hard to express...you walk in and everyone's face is a mask.

It is true that people tend to be more open when talking to computers, and people have shown to be open and more willing to contribute in online environments as well, particularly in highly interactive systems such as IRC and MOO. But even a shy student need not worry about not participating, because online it can be possible to send any kind of media to any person with every assurance of privacy, even when hundreds of people may be in the same room. Because of this, one must give up all pretense of holding a student's attention in a virtual class; if they do not want to listen, they should not attend. But crosstalk among students and onlookers should be encouraged, since if done privately it offers no distractions, and restrictive phrases such as "no talking in class" can become obsolete. Two of the all-important tenets of cyberspace are the freedom to learn and the freedom to communicate.

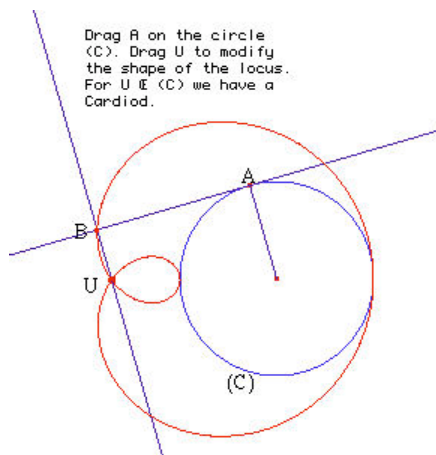
The freedom I talk about, incidentally, is not merely a matter of "academic" freedom. Schools are not just learning places but communities as well. Many schools are communities in the full sense of the word: people don't just go to them; they live in them. And, in the future, the distinction between "school" and "community" is likely to be much vaguer than it is now.

Many, many scenarios can be invented by simply applying the ability to search and retrieve incredible kinds and amounts of information online to lessons and by integrating the ability to communicate with other people, groups, and institutions. This kind of

creation is useful but we can also think of new ways to apply the technology. It can be possible, for instance, to run education servers, public online institutions that like libraries may offer a wealth of information (with online librarians, educators, and experts), like CD-ROMs may offer multimedia games, simulations, and exhibits, like a WOO may offer a way to collaborate and navigate with others, and like a textbook may be used as a tool for educators around the world to use separately or cooperatively.

Such a meta-university could be useful in many ways; a place to meet visitors and experts from other fields and recruiters from companies, a location in which students could come together and learn about each other as pen-pals do. Imagine a “global room” in which international students can gather. A teacher may ask the student to stand or highlight their country and give a tour or report about it. Now that many things are possible, we should think about making learning environments that are fun and intuitive that may be costly or difficult to create in the real world.

To say that learning is not something that is done to you implies the need for more than just autonomy. Within the school it means that everyone must have a voice in the decisions that affect him. This kind of arrangement - democracy - doesn't eliminate discord but it does put the responsibility for a school on all of the people in it.



17: Cabri, a relative to Ivan Sutherland's Sketchpad from 1965, is a geometry tool that allows people to create interactive geometric constructions. Similar software can be found at Swarthmore's Geometry Forum, <http://forum.swarthmore.edu/>

When the responsibility for learning is put on the teachers and students, we can think of new ways to encourage this. Students write oral reports not only to learn something but also to teach their peers; students who collaborate on building lego robots (in Southern California), create hypermedia tutorials (in Hawaii), and build CAD houses (via the Institute for Research on Learning) understand how to teach others as well as themselves. In a virtual environment, tool kits to create collaborative simulations can be made to teach by doing in areas like biology (controlling the behavior of different animals to understand flocking behavior and group dynamics), math (the creation of a virtual building, learning about physics, thermodynamics, and earthquakes), and art (building a collaborative gallery, dividing works by style, color, and period, giving tours to other students and classes, inviting guest docents from museums), to name a few. One can use the model of CommerceNet to create “EduNet”, which may be a non-profit consortium of educators whose goal is to stimulate education via the Internet.

At the Human Interface Technology Lab in Washington, they have found that kids love to explore in VR and push its limits, often coming up with scenarios that simply aren't possible with today's technology. After building and programming worlds, many report that it is hard but fun, that hard work definitely pays off well. Similar comments are heard from students learning with LEGOs.

There are plenty of obvious applications for educational simulators in a number of fields: physics, planetary exploration, archaeology, biology, and chemistry all can benefit greatly from better

visualization technologies. The point of these systems is not necessarily the amount of things that can be done with them but the fact that students can have the power to create them on their own and can find it fun and motivating to do so. Such systems make democracy easy; self-esteem and empowerment are side effects.

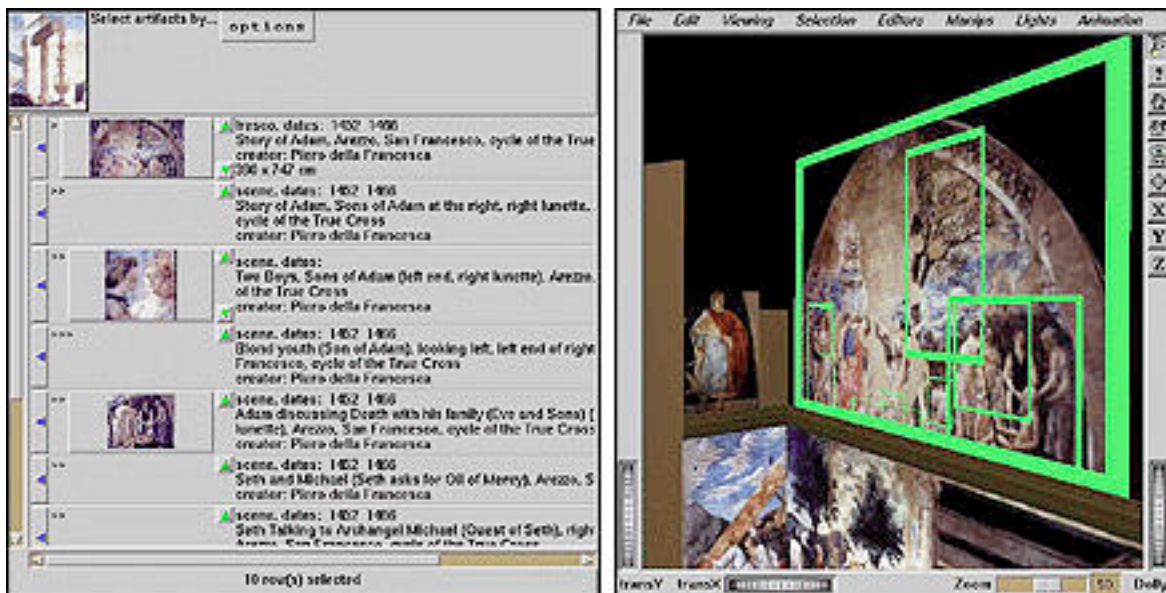
A well-rounded educational process should:

- Show that learning is an ongoing, active process
- Provide experience with creating knowledge and asking questions
- Encourage social learning
- Allow students to experience and appreciate different points of view
- Make learning relevant
- Encourage different ways of doing things

It is possible to use new technology to enhance the ability to do all of these things, and many aspects of virtual environments and hypermedia lend themselves easily to these goals. But technology by itself is not a cure for education: its operation requires an understanding of what it can and cannot do.

In summary: The idea is that the teacher becomes a guide. Students create their own textbooks and learn from them. There is no difference from authoring and reading, from creation and experience. Each of us has the potential to positively affect the

world, and we can do so throughout life. Everyone is different. Everything is interesting. These are the things that we should be taught.



18: ECIT, the Electronic Compendium of Images and Text, combines media databases and virtual environment technology that allow one to learn about the humanities. More information is at <http://mondrian.princeton.edu/art430/>

- Information on GNA and other online education resources are at http://www.yahoo.com/Education/On_Line_Teaching_and_Learning/
- *Reinventing Schools: The Technology is Now* from the National Academy of Science is at <http://xerxes.nas.edu/nap/online/techgap/welcome.html>

Entertainment

Enough has been written on the subject of entertainment and cyberspace; entertainment is one of the few areas in which VR technology has been put to profitable use. The first graphic adventure game displayed its screens using vector line graphics on an Apple II - the idea was to give one the sense of exploring and moving through and around a house. Today we are accustomed to 360 degree fields of vision, multiple environments, and highly textured objects. Video games are very much about giving the player the sense of being in another world; fantasy isn't always as compelling in the real world.

It is possible to visualize music in new ways, allowing artists to create programs that interact with audiences and display colors, images, and scenes. This is much more than MTV in that the experience can be immersive, or close to it, and cspace music videos now have three dimensions to take into account rather than two. Why not create movies with dissolving sets, or movies

taking place with scenes at different physical locations, with players half a world apart? There has been a great amount of discussion on the role that theatre can play in cyberspace, where there is less of an interface and many more ways than one to get an idea across to others. Much in the same way that ten years ago it was very easy to create a desktop-published page, and that today it is very easy to create Web exhibits and galleries, ten years from now it will be just as easy to create worlds and virtual plays, movies, and experiences.

Courses teaching others about playwriting, cinematic concepts, dancing, instrument playing, and acting can be made easy with the right tools, and all this can be shown live or canned, it can be replayed and saved, it can be shown to thousands of people from everywhere else. Learning can become fun, useful entertainment, and vice versa.

Environment and Nature

Environmental decay and erosion can be shown clearly; the small digital images that NASA scientists brought back from polar atmospheric runs years ago that depicted the growing hole in the ozone layer moved many people to action and ultimately brought about a fair deal of positive change. What if it were possible to report rainforest cutting in this fashion except with realistic detail and movement? What if it could be updated daily and viewed by the online world?

Simulations and depictions - migrations of birds, extinctions of species, layers of topsoil, snowfall and rainfall - if done well, they can educate, clarify, and encourage others to take action. Graphic earthquake and weather reports, storm warnings, fire damage, snowstorms - meteorology depends a great deal on visualization. It's the weatherman's only way of keeping the public's interest and their only way of keeping track of massive amounts of data.

Animal studies, locomotion studies, ecosystems and forest growth simulations can show people how entire systems can evolve and react. As a student I may see an average city and a redwood forest next to each other; as decades roll by I can see the apparent changes in tree height, skyscrapers, carbon dioxide buildup and where it occurs. I may study migrations of animals with elevation data, making it possible to compare the movements of schools of marine life with virtual fish controlled by flocking algorithms. I can ask myself, how do fish sense each other? and tune my simulation accordingly.

Those studying ecosystems can create simulations in which different virtual animals inhabit environments that they are suited to - as they migrate into different climates with different kinds of prey and predators, it is possible to see how they react. Brandon runs Forest World and tries to maintain a balanced system in its own right; Linda oversees Desert World and must keep things running smoothly. What happens when animals are displaced out of their natural habitats, either by force or by necessity? Is it possible to adjust both environments so a middle ground can be found?

Cspace is appropriate for studying “what-if” cases and distributing visual information to others. We can take the ability to create visual, “what-if” environmental scenarios and put it to use in research, education, and many other fields.

Government

The benefits that online participation can bring to local government have been well-discussed; we can explore scenarios which bring information to users and allow them to contribute in a meaningful way.

Imagine a walk-through model of the White House; as I pass into and through rooms I hear descriptions of the history and furniture associated with them. I can turn a “history dial” so I can see what the room looked like under the Kennedy administration or at the turn of the century. I can see outlines on the floor leading to the Press Room; I can literally follow the President’s footsteps as he prepares to make an announcement. I can examine artworks (which may lead me to the Smithsonian or another museum), or browse an archive of Presidential debates and follow along with transcripts annotated by various spin doctors and political analysts.

I view a site tracking the latest election as it happens - as votes come in, regions on a map gradually change colors as the display indicates possible outcomes. Growing numbers and percentages appear below the map along with links to a sample ballot and paragraphs about the candidates. I can view commercials from any of them (each one is allotted a fixed amount of time in cyberspace), listen to their speeches, and debate with others on the candidates’ viewpoints. I enter a private area and can vote at any moment; I decide to, and the choices appear with easily marked options. At any time I can request for more detailed information on any issue before I make my choice, but I must complete the ballot before I submit it. If I leave the private area (in which no one else can enter and in which all links are defined as being secure), then my unfinished ballot will be erased and I’ll have to redo it. I submit it and one of the regions on the map turns a different shade by a barely detectable amount (my monitor’s set to view thousands of colors).

Geography and Geophysics

Along with entertainment, geographic applications are a natural for cspace development. Maps, landscapes, real-world navigation aids, studies of continental drift and how populations migrate over time, these are all potential applications. Tools for cartographers; archaeologists who must record and annotate every discovery by orientation, location, and elevation data; historians who may view dynamic landscapes which move according to statistics and predictions; and people who simply want to know how to get from one place to another - all these things can be made.


Subway schedules, bus routes, traffic patterns, and train lines can be recreated in a dynamic form, so commuters can know when the next bus will come and where it will show up. Real-time tracking capabilities can let users know when an accident has

occurred or when a train is running late (and how long it may take to arrive). As personal GIS hardware gets cheaper each year; the ability to develop systems like these increases.

The BADGER (Bay Area Digital Georesource) project hopes to develop an online visualization system for geographic data of the San Francisco Bay Area. VR technology could easily be used in this project to allow easy navigation and visualization; it could be used to develop modules to incorporate the system on the World-Wide Web and could be used as a basis for developing filters and gateways for new data so Bay Area residents could use the system to browse many local services: online business, real estate information, traffic and weather information, environmental and census data, and so on.

One can create a navigable interface to the Earth's geography; this idea is not new (the Geosphere Project has had a similar goal for years). However, this interface can be made general so it may act as a directory to phone systems, Internet services, country information, and so on. This idea has been embodied in two dimensions on the Web at the Virtual Tourist site. In three dimensions, as I move over the world, I may be able to hear radio and net broadcasts, or see other media appear and disappear. I can watch the sun rise and set as I circle the Earth, or observe graphs following the flow of import and export of cars, oil, currency, infor-

mation. I can see networks connecting countries and satellites and tap into different streams of data. In this environment we have the possibility to see the big picture all at once.



19: The TerraVision Project allows navigation over the Earth using a globe-like input device that users may roll. By zooming into the surface, it is possible to view large buildings and other features, achieving a very high degree of resolution. The project's main page is at <http://www.artcom.de/projects/terra/>

- The Virtual Tourist is at <http://wings.buffalo.edu/world/>
- BADGER information is at <http://www.svi.org/BADGER.html>

Health There are a number of exercise machines that use virtual landscapes to encourage progress; while climbing stairs or pedaling a bike, one can ride through a town - or take a tour of a favorite city. It is certainly possible to create engaging biofeedback systems.

Much telepresence and simulation work is being done in the medical fields. Its traditional use of high technology to do visual imagery has encouraged the development of high-speed collaborative systems. But beyond all of the various work being done for surgeons and medical staff, there are many applications for the patient as well. Users can connect to a hospital's advice service and manipulate a 3D body to show them exactly where the pain is; they can use sliding scales to indicate the kind of pain that they are feeling. People can make use of medical expert systems or take courses in basic health, learn new exercises, and access their medical charts (which would be typed in and not scribbled on by doctors, hopefully).

Coupling virtual environment technology with sensing devices allows exploratory work to be done without having to touch patients and facilitates planning for many intensive, location-oriented procedures such as tumor removal, dental work, and plastic surgery. Surgical simulation can be used in medical education as well, and imaging of this type can be used to show patients what is happening (for instance, a mother may be shown three-dimensional ultrasonic scans).

Because real situations can be simulated, it's possible to help cure problems related to reality - for instance, current work is being done in using VR to help cure those with particular "space" phobias - fear of heights, claustrophobia, and so on. It is ventured that the technology can be used in role-playing and social skill learning, that it may become a valuable tool for psychologists.

- A comprehensive HITL report, *Medicine and Virtual Reality: A Guide to the Literature*, is at http://www.hitl.washington.edu/projects/knowledge_base/medvr/

Science

Research papers can consist of interactive exhibits of experiments: a scientific journal may be a world in which data can be played back and experiments rerun with different parameters. Mathematics journals could, of course, be highly interactive.

There are already numerous examples of using 3D computer graphics in scientific visualization so that they need not be discussed here; however, the ability to teach science using such technology has not been so widely discussed. Students could perform physics experiments and change any manner of variables; plant growth and cell division could be simulated and compared with other forms of growth; mathematics students could watch the graphic output of algorithms as values change, and so on.

Today scientists have been moving atoms, molecules, and DNA; traditional work has been done in the fields of aerodynamics, astronomy, and mathematical modeling.

Teleoperation, Telepresence, and Telerobotics

The ability to remotely control a device as if it were an extension of one's body has proven very useful in hazardous situations that do not lend themselves well to live human presence. Manually operated robots have been made to fight fires, defuse bombs, aid police busts, clean up nuclear accidents, and explore volcanos.

Virtual environments in which manipulated objects have real-world effects can be referred to what David Gelernter calls *mirror worlds*. Environments can be created to control factory floors, maintenance systems, intensive care units, planes, and air traffic. Because the screen can display a number of widgets and displays in a smaller space than physical ones, such mirror worlds are useful in applications that make use of many controls, such as virtual cockpits. Some projects have proposed that all conventional cockpit instrumentation could be duplicated functionally using an HMD.

Many more applications of virtual environment technology exist:

- Media research: Law students can replay 3D incidents and corroborate them with the books; film students can recreate shots and compare them to multiple films.
- General simulation and training: Training for miscellaneous kinds of vehicles, space missions, combat missions, air shows.
- Consumer products: Product visualization, industrial design, and marketing and sales. One can visualize car interiors, new fashions, ergonomic positions, and all manner of prototypes.

3

Creating Cyberspace

3.1 VRML and HTML

Now that version 1.0 of the *Virtual Reality Modeling Language* (VRML, pronounced *ver-mull*) has been specified, we can begin to look at ways in which it can be extended to support the kinds of features that future cspace systems may require. Basic HTML has been popular because it is very easy to read, learn and write; however, this “feature” would have gone unnoticed if adequate HTML WYSIWYG editors had been supported by developers early on. Although it’s possible to write and edit VRML manually, we must think about building intuitive editors now, so that users need not have to page through code or worry about typos. The “drag and drop” method of world building that programs like *Virtus VR* uses is friendly and easy to use; this is a very good basis for creating a multiplatform, generic editing module for worlds.

- VRML information is at <http://vrml.wired.com/>
- The VRML 1.0 specification is under <http://vrml.wired.com/vrml.tech/>

3.2 New Definitions

The words *teleport* and *avatar* are fairly well-used today; an avatar may be referred to as the representation of one’s self in cyberspace, and the ability to teleport gives one the ability to disappear in one part of cspace and appear in another, making it similar to the ability to quickly move from page to page in webspace. The proper way to think about movement in cspace is not that worlds disappear and appear while one stands still; rather, you move around from server to server. The ability to think this way in webspace has given rise to new forms of systems and is key to WOO development - the page doesn’t disappear to be replaced by another; you move from one hypermedia space to another.

When people start making new worlds and environments, they will encounter many of the similar problems that plague webmasters today, and it is worth taking some time to think about potential problems in order to find solutions. Those who make or maintain worlds may be called “worldmasters”; other possible terms are “spacemaster” or “placemaster” (as proposed by Ubique), but they sound a bit ordinary, as if they were furniture or kitchen tools. Since VRML-defined scenes are called worlds, it seems appropriate, and it keeps the idea in mind that one may truly create a complex world and all that the word implies, not just something in a “space” at a “place”. And “worldmaster” fits nicely in the same sentence as “webmaster”.



In time there will be worldmasters well-known for their 3D designing and programming abilities - the job will help to popularize those who work in the movie industry and will require new types of training and different skill sets. Professional teams will work in groups, coordinating with marketing departments, system administrators, digital artists, and managers, and they will slowly influence the way the Internet looks and feels and how it is perceived by business and media.

Just as people have developed the concept of a "home page," so will there be a "home world" which may be the entry point into multi-world organizations (in the same way that multi-Web server companies have a main corporate server). Multi-world systems may be called "galaxies" or something similar.

Many HTML files have ended in the `.html` suffix; there seems to be no good reason why VRML files cannot end similarly in a `.vrm1` suffix (or `.vrm` for DOS machines). It associates the file with the language, for one; the `.vr1` type is general (`world` is proposed to be a major MIME type, after all), and other 3D formats could easily claim that suffix.

3.3 Tools

We can think about developing tools and features that make it easy to organize and search through VRML files, so that browsers can display extra information about objects: for instance, it would be nice to be able to associate a description with a polygon or group of objects, so as I navigate close to objects or move a mouse over them, the browser can give me more information. This information can be used to create worlds for the blind and visually impaired (the browser may choose to speak this information). Some sort of defined structural support for keyword, copyright, and author information is needed (rather than simply putting this information in comments) so that robots and other utilities can quickly index or archive VRML information.



20: A "slice" tool can be used to create 2D imagemaps automatically in order to represent worlds in generic HTML browsers.

One popular concern is being able to visualize worlds in other ways so that other applications or displays can handle them without needing any sort of 3D graphics capability. If objects within a world had descriptions attached to them (this would be an equivalent to HTML's ALT tags), it would be possible to write an HTML convertor that simply generated an alphabetical list of descriptions with links. It may be possible to take a slice of a world, say with a large planar knife tool, and generate a GIF and coordinate file so that the particular view could be used as an imagemap.

It is also possible to create tools that translate the other way - traditional media and hypertext to three-dimensional worlds - and it is in this area in which some of the more exciting possibilities of cspace can be explored. Here are some examples:

- I tell a utility to map a Web directory; it generates a cone-shaped structure of files that I can browse through.

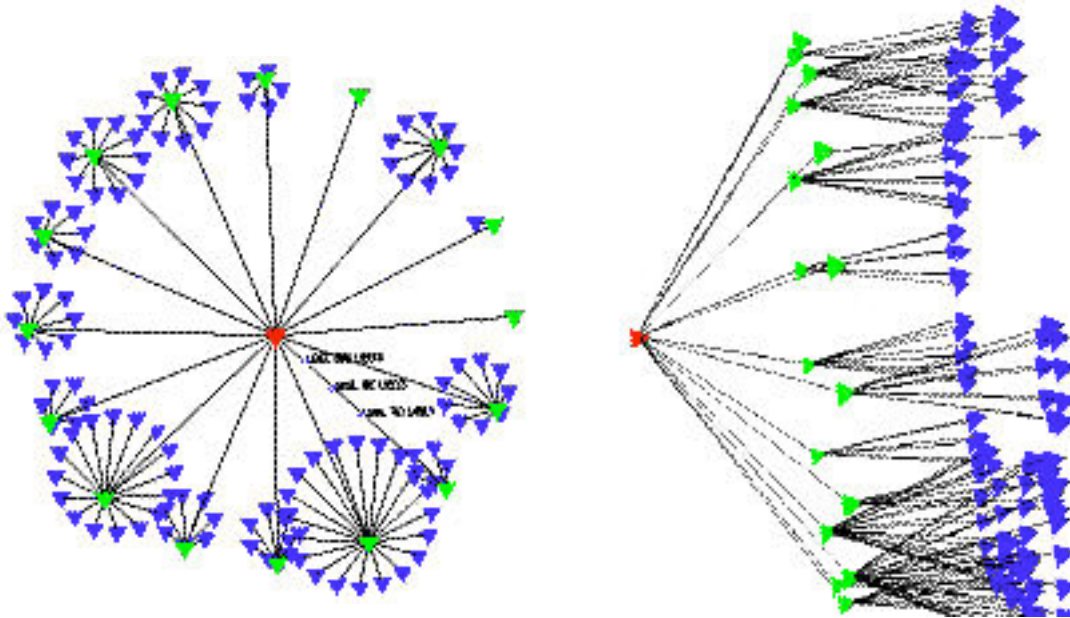
Utilities along these lines were researched at Xerox PARC; they created an umbrella-like directory structure which could be rotated and browsed, with shadows underneath which, as it turned out, were necessary visual aids to navigation. At PARC a horizontally scrolled list metaphor was also created for browsing files. Users could scroll the page by dragging on it and move and recolor files. All of this research was done some eight years ago, but today we can create similar structures with VRML and make it easy to research new ways of interacting with media. By creating general structure libraries and utilities, one can use the same code to create Web and world administration tools, check-in and check-out systems, log analyzers and database interfaces, and general tools to view and manipulate large amounts of data in different types of file systems.

- I take an HTML file which contains inline images, a few paragraphs, and linked list of ten objects - one links to a file, another to a movie. I run it through a VRML convertor. It generates a ten-sided sphere (the movie's opening frame appears on one side, the others hold links to their respective places). In the background the inline images float on planes; other planes appear that display the formatted HTML paragraphs (this may be via a texture map, generated text described in VRML, or output from an HTML module).

The idea of a spherical file browser comes from MIT's Media Lab; the structure could easily be any number of shapes or planes. A long menu may translate to a hallway, with each labelled door representing a different choice. However, it is cumbersome in a mouse-driven system to have to turn to view the information on every door, so it makes more sense to generate large circular, spherical, or concave structures (a la Stonehenge) in which a user in the center can rotate around or, better yet, have all of the options appear in the same plane, similar to SGI's "buttonfly" demonstration program.

- Multiplatform tools that make it easy for people to create worlds will be needed; one could easily imagine an authoring environment which could

handle various forms of media or a “worldmaster’s starter kit” that allowed people to easily set up and administer cspace servers.



21: An experimental 3D Web visualization done at the University of Minnesota.

Below are some pointers to current 2D and 3D Web utility research:

- An experiment in 3D webspace visualization is at <http://www.geom.umn.edu/docs/weboog/webospace/webospace.html>
- A paper on WebMap, a 2D webspace visualization tool, is described at <http://www.tm.informatik.uni-frankfurt.de/Publications/Doemel/WWW94/WWW94.html>
- WebViz, a 2D access log visualization tool, is at <http://www.elsevier.nl/cgi-bin/WWW94link/31/overview/>

3.4 MOO/WOO Thoughts

It is possible to create a core library of presentation code so that it may be plugged in into a traditional Web browser or a different VRML-based sort of framework. With such a library, it would be less difficult to create WOO systems in which 3D objects could be handled via text commands or VRML scenes could be rendered as static images (merging VRML worlds with the traditional HTTP way of doing things). Using something like CCI or an API particular to the library, one could easily create a MOO to VRML gateway, so that traditional MOO systems could be visualized in 3D.

Persistent storage is necessary; as more Web servers are developed that can handle media using databases, it will help to think about how one could store, retrieve, and manipulate worlds that consist of objects stored in them. One could ask a WOO server to create a cup; the system grabs its generic representation of a cup and makes it appear at your location, keeping all environment information in its database. You may ask it to grow or shrink, to mutate into a table, to change colors, etc. In this way one can create ways of teaching programming (think of LOGO, but with 3D output), design, and mathematics (a collaborative 3D calculator). With the power of a database behind object storage, one could ask the server to “display all the objects that look like this”, “highlight all the houses within five miles”, and “show me all animals that have six legs”.

Because of the highly interactive nature of cspace, objects in worlds will be created, destroyed, and changed very quickly - the update and reading of worlds may be faster if a database architecture or RAM is used to hold the dynamic world state at any instant, with snapshots of the current world being written back out as backup at longer intervals. Various methods of multicasting must be explored.

3.5 Coordinate Systems

It is a good thing to be able to display coordinates using different systems: particular applications may require it, and it may help to create new ways of thinking about information, not to mention any number of special effects and art forms that could be generated using non-Euclidean spaces.

Some care should be taken in developing a framework so that objects can be oriented to other objects and to their parents and that scenes can be oriented with respect to other scenes and other worlds. Otherwise, one gets mismatches of scale and orientation while traveling from world to world, and navigating worlds then becomes a bizarre trip through a virtual funhouse.

A number of common coordinate systems are in use that work well when applied to particular applications. As a default for virtual environments the Cartesian coordinate system has worked quite well. The only issue that arises is one of labelling the three major axes.

We can think about what order of labelling makes sense by taking a step-by-step approach: what is the axis that humans naturally focus on? In many instances we tend to focus on the horizontal first, and then the vertical when dealing with two-dimensional objects. This is due to the fact that our eyes are spaced horizontally, perhaps an evolutionary result of the fact that the horizon is flat, and in order to hunt and evade efficiently our optics needed to adjust to the environment in this manner. We do naturally, then, focus on the horizon, the X axis, first. The vertical, Y, becomes secondary.



The only axis left is the Z axis, the third dimension: which way does it point? Is the XY plane flat, or is it upright?

When people make models, architectural models or small constructions, we often visualize the XY plane as lying flat and the third dimension rising upwards: we focus on the flat first and the vertical second. If the XY plane is horizontal, we look at buildings on their side, something we do every day; if the plane is vertical, we then see the top view by default, placing the observer in space, which can be rather disorienting.

If the XY plane is flat, one can make use of the right-hand rule: if the right hand is curled, the X axis points out along the arm, the Y axis points at a right angle to that in the direction of the curl, and the Z axis points upwards. When all one has is information about two axes, the right-hand rule can be used to reorient oneself in the coordinate space.

In systems where a user must navigate and locate themselves, models should be constructed with a horizontal XY plane.

3.6 Editing




22: Default coordinates and the right-hand rule. The XY plane should be flat when designing virtual environments.

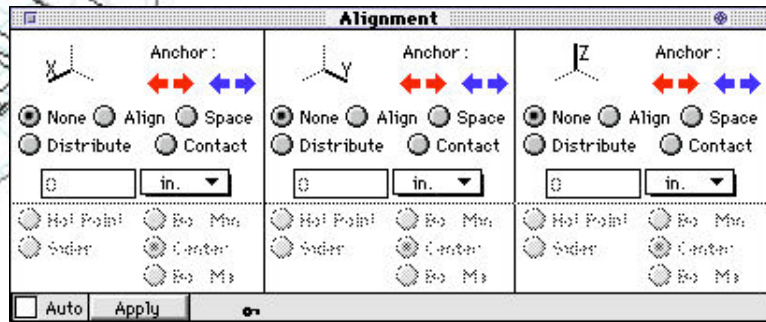
A VRML editor and modeler should be able to let users create using a number of standard units - centimeters, millimeters, meters, inches, feet, yards. This greatly helps people to map real objects into 3D representations. Quality VRML editors will take lessons in orientation, editing, and user interface from the more successful 3D editors, such as Ray Dream Designer's editor and SGI Inventor's editing functions. Among the features of a good basic editor:

- The ability to "snap-to" a grid and align objects with others
- The ability to repeat transformations and duplicate, clone, cut, and paste objects
- The ability to work in 3D or planar views
- The ability to group objects into groups and layers, which may have show/hide/color attributes
- The ability to work in wireframe (showing points) or rendered mode
- The ability to embed levels of detail within objects
- The ability to display the polygon and object count of scenes
- The ability to align, move, and edit using numerical values
- The ability to handle and edit lights (and cameras) as objects
- The ability to add objects via a drag and drop method from libraries
- The ability to simplify objects

- Interface consistency and simplicity



23: Ray Dream Designer's editing view and alignment options. Designer has won acclaim for its easy-to-use interface and clear display.

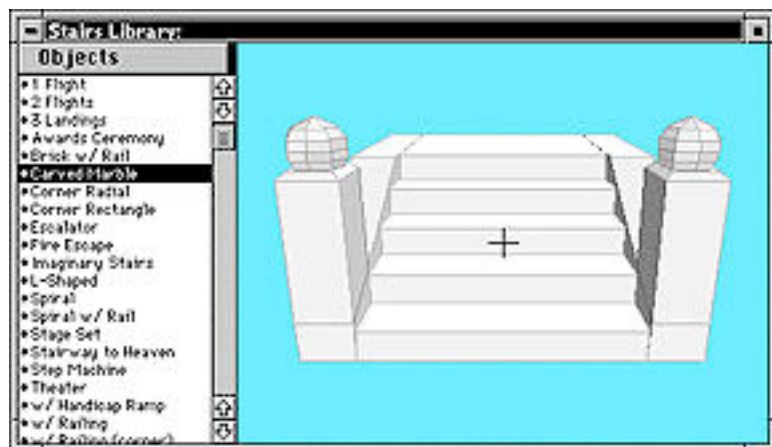


3.7 Objects

Clip Objects

Just as clip art, datasets, and FTP image sites exist today, sites for 3D objects will appear once any amount of world creation begins. Because of this, again, it is important to be able to embed comments (such as author and copyright information) via some mechanism into VRML files, and it should be easy for users to view this information via browsers.

24: Virtus Walkthrough comes with a number of object libraries which can be easily browsed. Tools to browse object repositories in this way can be made to make world creation easier.



Just as Virtus has a generic "Brutus" character and *Snow Crash* described a generic "Brandy", so too will people use generic, commonly used characters and objects. When the ability to give objects programming or associate them with scripts in some fashion is possible, newsgroups that are a cross between `alt.binaries.pictures` and `alt.sources` may appear - perhaps an `alt.sources.vrml`?

Any number of default world templates could be made: aquariums, rooms, houses, auditoriums, forests, cities, landscapes, and so on. And programs and filters could be created so that users may change different elements of the environment: put loud wallpaper on buildings, create a starry sky, plant trees and bushes everywhere using a "forestify" filter.

- Information on Virtus can be found at <http://www.virtus.com/>

Privacy and Permissions

Every object and scene should have permissions attributes, using a model similar to how UNIX or AppleShare implements permissions. If I am browsing a world, objects that I do not have permission to read will not be sent to me (this makes things like secret doors, clubhouses, and the like easy to implement), objects I can't write to will refuse to be modified, and objects I can't execute will not be interactive (perhaps their links are disabled, or their programming won't be executed).

Permissions make it easy, then, to prevent worlds from being modified by others (you can actually prevent people from littering), and being fully interactive (for kiosk exhibits and the like). If I am able to send an alternate object if a user can't read a particular object, then I can do things like make structures whose doors are opened or closed, making it very easy to see what I can and cannot do. Publicly writable objects may include bulletin boards and like. There should be a mechanism via which I may specify a password before copying an object or entering a scene or world. One can take the lessons learned from secure protocol research and apply them towards the development of a framework that can allow the fast encryption and authentication of cspace data.

Permissions for worlds in general should be a little more specific: Can users create new objects in them? Are there access restrictions by location or IP address? Can users edit and change individual objects?

The idea of using machines as boundaries is natural - as more people use their personal machines (laptops, perhaps) with their own worlds to connect to cspace, the virtual world will begin to reflect this. In a public space, for instance, a facade representing a home or corporate structure may be made, but only when a stranger enters does he or she enter the particular machine where the world exists.

The Scribble Attribute

If I give an object or polygon a scribble attribute and make it writable, then I may be able to write on the surface using any color I select in my browser. This makes it easy to create guest books,

shared whiteboards, and graffiti areas. Scribble elements that are writable via permissions may only be virtually written on - that is, no one but the owner can delete them or change their shape, etc. The reason why polygons should have scribble attributes is so users can make flat surfaces or structures on which some but not all sides can be written on. By adding texture maps, one can simulate the look and feel of a brick wall, chalkboard, or paper.

The problem with implementing a scribble attribute is that one must think of a way to store scribbled information and transmit it in a way that facilitates real-time interaction. But if done well, this attribute can help create new kinds of objects, like virtual Post-It notes, and foster new kinds of communities (cooperative art spaces, art and animation classes, storyboarding sessions, virtual meetings).

Texture Maps

Through the use of texture mapping, it is possible to create objects that look more like their real-life counterparts. It is desirable to be able to specify JPEG as well as GIF files for texture mapping, since they can help make transmission times short while keeping image quality high. GIF texture mapping is also desirable because it is possible (using the GIF 89a format) to specify a transparent color, so all sorts of effects (tattoos, decals, text, irregular images) can be made. It is helpful to be able to specify alpha channels or extra texture map information such as bump maps.

If viewers support comments within texture files, it should be possible to be able to select a texture and read the image's comment in a new window or status display. This would make it very easy to create galleries with information on the pictures without needing an HTML-based browser, and would help to clarify texture copyright issues if copyright information is included in the comment.

It is also desirable to be able to use movies (MPEG or QuickTime files, for examples) as texture maps, so that the illusion of movement can be conveyed (flowing rivers, waterfalls, moving sky). It is best to allow the author to specify a default play rate but give the viewer some control over playback.

In world creation there should be a way of defining the background, as a color or some sort of repeating texture map. That way it is easy for authors to make a semblance of an atmosphere without actually having to contain the whole world in a giant primitive whose interior was mapped. The ability to create an infinite ground (with an infinite horizon) is also desirable, as that prevents a user from conceptually "falling off" the end of a world, when in fact there need be no end at all. The floors of many 3D scenes simply consist of a large flat plane; when one reaches the end, there is simply nothing there, and that may detract from the experience. It should be possible for authors to avoid this.

It is desirable to be able to texture map live video onto polygons, so that things like live movies and videoconferencing may take place in cspace. With live video, one could easily make virtual aquariums, TVs, and other objects. However, adequate video (and audio) capabilities may be best left to other programs or modules that are designed to do such things until a robust protocol that can handle such interactivity is developed.

- Information on papers that describe architectures that allow real-time data to be handled by applications can be found at <http://www.tns.lcs.mit.edu/publications/publications.html>
- Information on the MBONE, the Internet's multimedia backbone that transfers real-time audio and video, can be found at <http://www.eit.com/techinfo/mbone/>

3.8 Constructs

The concept of the virtual document has proved to be very handy; the by-product of many gateways, CGI programs, state-passing Web programs, and other utilities is virtual documents. They may exist for a moment in any state and are always generated automatically by server-side processes. Without the capability to generate such documents, the Web would be inflexible, and a lot less capable and interactive than it is today.

We can extend the philosophy behind virtual documents into three dimensions, in which one can create virtual objects, also perhaps traditionally called *constructs*, after a Gibsonian term. Dynamically assembled constructs result from the output of some process or groups of processes; APIs can be created to interact with these constructs to create completely new forms of utilities. One may be able to better grasp the idea by thinking of constructs as being like program-controlled puppets.

The model of the World-Wide Web has shown that it may be best to distribute most of the processing power behind constructs on the server side, or in such a way that the server can distribute its processing needs. Using an API users could create constructs in many languages for any computer platform. This has been one of the strengths of the Web model. However, because cspace is so interactive by nature, servers handling only a few complex constructs (such as an articulated walking skeleton) may quickly become overloaded. It is necessary to build some simple, flexible scripting abilities into VRML so that clients can handle a certain amount of processing without adding so much that the language becomes too complex.

First, to display a construct in VRML, one might use `wwwInline` and specify a program at a site. The program, via some mechanism, could receive any amount of information: the basic sort of information provided with CGI, viewer position, local time, longitude and latitude, selection location, velocity data, detailed world

information, and so on. One may want to send only certain variables (to keep transmission time down), rather than every possible variable, so one could do something like this:

```
WWWInline {
    name "http://www.foo.com/map"
    constructVar {
        longitude
        latitude
        gmtime
        eventlist
    }
}
```

so that the map program residing at `www.foo.com` will only receive longitude, latitude, GMT time information, and particular event signals.

To provide for interactive constructs, it must be possible to be able to update objects to give them some semblance of reality. There should be a number of event hooks provided so that constructs can be created:

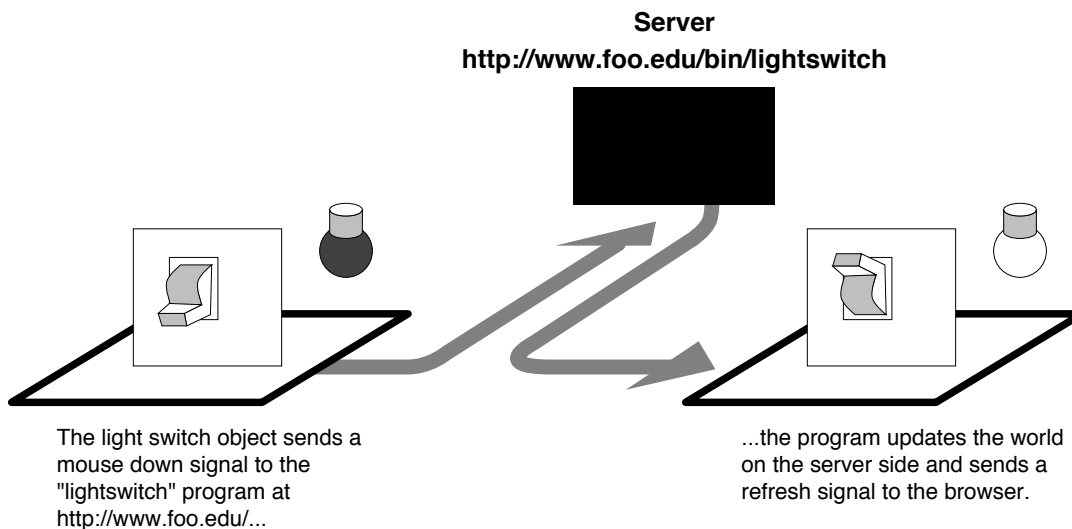
- At a particular time or at regular intervals
- When a polygon is touched (via `WWWAnchor`)
- When a polygon is held (such as when a mouse button is held down on it)
- When a polygon is dragged
- When a polygon is released
- When an object collides into another
- When a user is a certain distance from an object
- When a user reaches a particular location
- When a user enters a world
- When a user leaves a world

The ability to specify events particular to different keys and mouse buttons may also be useful.

The program will send back VRML data with the `world/vrml` MIME type; it must be possible to specify that the output of a program replace the previous output of that same program, in order to give constructs the illusion of movement. Because event loops and construct processing can slow things down considerably it may help to be able to tell a VRML browser to not send particular event signals; however, this may cause some confusion in highly interactive collaborative environments.

Programs could send back `text/html` content or URLs on any of these events. When a VRML browser receives such information, it can run a Web browser to display the data and establish a CCI connection when appropriate. So when a user enters a world, its HTML home page may appear. When a user moves into

another room, the page can change to provide descriptions and other data. If the VRML browser can receive data via CCI itself, it can be coupled with a WOO browser to create powerful collaborative systems. It is worth looking into programming and event handling methods used by such languages as Meme and HyperTalk.



25: How a simple construct works. The construct, a light switch object, has polygons linking to a program stored on the same machine as the world file. When the object detects that a user has held the mouse down on it, it sends the signal to the lightswitch program, which takes the appropriate action. In this case, it updates the world with a new light switch and bulb object and sends a refresh signal to the browser so changes can be viewed.

- A paper on Meme can be found at <http://remarque.berkeley.edu/~marc/home.html>



26: This late-1994 virtual environment, developed at the Mitsubishi Electric Research Lab (MERL), shows Mike, a voice-responsive robot construct.

Avatars

We can think of avatars as being constructs - three-dimensional virtual documents that follow a user's viewpoint. In many situations, it is a good thing to be able to get a sense of who is occupying a particular world. Therefore servers must track the locations of people and send avatar objects representing them to browsers when requested. To make environments more like reality, the browser must request avatar object updates regularly. This may be done by sending some sort of "avatar refresh only" signal to a server. To save processing time, one could create groups of signals that only update particular constructs and allow a user or browser to regulate these, so as I move towards someone's avatar to talk to them, their refresh rate and/or resolution increases to allow for more detail, while other objects to my sides and faraway objects may become more static and blocky.

One's avatar could be represented as the output of a program itself or as a static VRML object. One must give this reference to the browser so it can send the avatar to servers in order for other people to look at it. It must be possible to create avatar objects which can only be viewed by one's self and which are not sent to servers. This gives one the ability to create personal, private objects, which are discussed later. If you do not specify an avatar to use, generic ones should be provided by browsers (or they may use a previously cached avatar).

Because a world with many, many occupants can take a long time to render, it is a good thing to be able to selectively control who one can and cannot see (and completely ignore), by such criteria as name, time, IP address, and proximity to one's self. It is



27: Avatars within the MERL environment are described by a minimal amount of shapes but can be effective in communicating gesture. In this environment, operators use various input devices, from sensing jackets to physical joysticks.

best, however, for a browser to always display a list of the occupants in the same world (which I may sort according to the mentioned criteria) and be able to perform actions with any number of them (such as chat) even though I cannot see them. So at a virtual concert, I may choose to view only the players and people within ten feet of me; there is no need to render the entire audience. Servers may choose to pace themselves by putting this sort of restriction on members when their load increases, in order to quickly dispatch updates.

The ability to “filter” one’s view of participants is akin to the IRC model of channels - I could choose to read and respond to every existing channel on the server, but that would create much too much information for me to keep up with. Instead, I can decide to listen in only on particular channels that interest me. My browser would filter other users and send update requests for only the people I wish to see to the server so it would not have to do any filtering itself. In general, servers and browsers will have to communicate very closely to determine things such as:

- When objects (static objects, constructs) enter and leave worlds
- When objects are created and deleted (by themselves and by others)
- When objects must be updated (via VRML scripting, server-side programming, or user-initiated events)

If the browser gives me visual or textual hints, such as, “Bob can see you and hear you,” this can be useful, and if I am able to tell the browser to watch for certain people within a particular distance - “Alice has entered the room” - trying to find people can be made much easier.

- Information about the MERL is at <http://www.merl.com/>
- The DIVE home page is at <http://www.sics.se/dce/dive/dive.html>
- Work on instilling avatars with human-like movements is being done at New York University’s Media Research Lab, http://found.cs.nyu.edu/MRL/proc_anim/animation.html

The word *avatar* was first used to refer to people within virtual environments in Lucasfilms’ Habitat, an online MOO-like community that used two-dimensional artwork for representation. Some years have been spent in the development of its sequel, WorldsAway, which is offered by Fujitsu and CompuServe.

- WorldsAway information is at <http://www.worldsaway.ossi.com/Worldsaway/away.1.html>



28: Two avatars within DIVE (Distributed Interactive Virtual Environment)


WorldsChat is something like a chat service that makes use of virtual environments and avatars. Its avatars are displayed using textures, saving some rendering time and allowing a fair amount of detail as well as expressions.

- WorldsChat information is at <http://www.worlds.net/>

New Metaphors

Cyberspace: First Steps describes a portrait of a worker wandering through a virtual office with a backpack, in which he stores all sorts of personal references and objects - think of it as a 3D tool box and hotlist. The backpack goes with him wherever he travels. When one is able to create personal objects that nobody is capable of seeing (servers would have no knowledge of them), it is possible to create virtual wallets that hold business cards, tool boxes of constructs and objects, and many other things. When one enters a world, their avatar information may be "published" to other servers. When private objects are dragged and dropped into a world, they become published - public - and are at least readable (not necessarily writable) by everyone in that world. In the same manner, if an object in a world is readable, I can drag and drop it into my tool box, in which a duplicate copy can be made.

29: What a simple tool box may look like. One may be able to store personal information in it and drag objects from it into worlds. It may be handled as part of the browser or it may be a completely separate module or application.



If in creating objects one is able to specify that they are public and not able to be duplicated, then one can simulate the handling of real objects. I could walk through a virtual cave and collect rocks, or leave business cards and constructs from my wallet on a virtual table (only if I am allowed to create new objects in this world, of course). One has the ability to make virtual transactions. Sharing "dumb" objects or objects completely described with VRML is easy, but if I give a construct to a friend, I would need to give him the construct programming (or he may only receive the 3D output of the construct). If the construct uses server-side programming in a language he doesn't support, then my friend will be unable to use it in any meaningful way. In my browser I would have to specify an area to store things which I pick up or are given to me - this may be a directory or a pointer to my virtual tool box (which I can drag and drop things into and out of).

Using the aforementioned event handling capability, one can pick up, hold on to, and drop objects; objects that are constructs themselves can tell when they are being dropped and shatter into pieces when they hit something else. It becomes possible to create representations of agents that do some amount of work for me. It becomes possible to create programs that wander around worlds and interact with and follow other users, if they are smart enough. This entire class of constructs can be incredibly useful but can also be potentially very annoying. Much as there are guidelines for writing Web robots today, there must be guidelines for developing intelligent and respectful constructs.

Client Processing

There should be some simple methods available of client-side processing in order to take the load off of servers. The type of basic transformations and animation that Inventor provides are a good start.

Any kind of language that VRML uses should be easy to read and understand (it may incorporate functions, nesting, be somewhat object oriented, etc.) if many people will be coding manually. However, some thought should be given towards developing interfaces (such as GUI-based VRML editors and programmers) so that the average user need never learn VRML's internals. As Dave Raggett has said, polygon specifications are the assembly language of VRML.

Constructs can be used to create new forms of interfaces or to emulate old ones, to create new kinds of programs, such as self-running exhibits and simulations, unfolding structures, dynamic sculpture, self-extracting and self-constructing objects. With constructs, I can make virtual suns that rise and set within my world. They are what will bring life and creativity into cyberspace; they will make it a system worth learning and collaborating in.

3.9

Links

The ability to link to other media should be one of the core capabilities of any VRML browser. It may be desirable to provide hints that an object is selectable, either by rendering the object in some different fashion (making it turn red with the mouse passes over it) or displaying that information in a message area.

In order to create a 2D imagemap in a 3D space, one could overlay a number of completely transparent linked polygons on a surface. In this case, it would be a good thing to highlight these regions in some fashion as stated above.

3.10


Sensing Proximity and Location

The ability to detect collisions is one of the major features cspace needs in order to make it seem real. However, such detection can be very computationally intensive when many objects come into play. There should be a collision detection attribute for objects, but this may be ignored if there are too many objects in a scene.

The ability to send proximity information to constructs is also important; using this feature one could create worlds that can be navigated by the visually disabled: as I approach an object, it may tell me that I'm getting too close; a dog may bark and growl louder and louder as I approach it. Since some sort of ability may be built into the renderer (and it may be necessary when using particular protocols), the performance hit taken due to this feature is probably one that browser developers will have to live with.

3.11 Moving and Viewing

Floors and Gravity



30: Polygons with floor attributes keep a user's viewpoint a certain distance above them (top diagram), while gravity acts as a constant force, dragging one's viewpoint downwards until a certain minimum height has been reached.


Although it may be possible for users to move throughout cspace at any perceived speed or elevation, there should be hints given by objects that they contain particular physical attributes, so that a world may appear more real and solid.

One could specify a floor attribute for polygons. When my viewpoint passes over a polygon with a floor attribute, it assumes a constant height relative to the elevation of that polygon and straightens in the direction of the polygon's normal. If I pass over two floor polygons, my elevation becomes constant relative to the highest polygon. I may be able to indicate in my browser what my default height is and whether I wish to follow floor hints; doing so can make it easy for people to navigate up and down stairs, follow graduated trails, and follow paths to walk upside down or sideways. Although movement may be jerky in some instances, floor hints can make some forms of navigation very easy and can help people to maintain a general level in a social setting.

The ability to generally specify gravity within a scene may be important; with gravity on, users will lose elevation until their viewpoint reaches a certain elevation above the highest polygon below the viewpoint. This can allow users to drop from the sky and land on the ground, it is an easy way to simulate moving down stairs and ramps, and it can facilitate jumping in a fairly simple manner. When it is possible to jump, drop, and land, many sorts of action-based worlds can be created: games (a 3D Super Mario Brothers, basketball) and adventures (parachuting, mountain hiking), for instance.

Navigation and Orientation Widgets

Full movement within three dimensions requires six degrees of freedom: movement along the X, Y, and Z axes, and roll, pitch, and yaw.



31: From left to right: roll, pitch, and yaw.

There are many sorts of methods that could be created to facilitate interaction within 3D worlds; the best guideline to stick with is “keep it simple”. In virtual environments many problems occur due to users not understanding the rules of movement, so it is important to provide a simple, easy-to-learn navigation solution.

Solutions range from being very tactile (HMDs, gloves, body suits) to being very portable (touch screens on PDAs). One of the more successful solutions, since it does not require any large amount of sensory hardware and can be implemented easily, is the method of navigation using a virtual joystick. Imagine a top view of the joystick superimposed on the view of the world. If a user moves a cursor (with a keyboard, mouse, or other device) above the simulated horizon, they move forward in the world. Below the horizon, one moves backward. To the left, one rotates left, and to the right, one rotates right. Often crosshairs and tick marks in the center of one’s view are supplied in this model so users have a reference point: the farther away the cursor is from the crosshair, the faster one goes. One type of virtual joystick places a crosshair or box where a user clicks; this then serves as a central point of reference from which the mouse drags. The crosshair/box is then removed when the mouse button is released.




32: Two possible “virtual joystick” navigation elements. The left diagram depicts the widgets used in the DIVE environment. The right depicts “heads-up” mode in VRweb.

Because movement in this model is usually accomplished by clicking and dragging the mouse, it becomes problematic when one wants to select a link in a world or move an object but not

move (this is especially a problem if a user has a one-button mouse - an awkward key combination might be used). So this virtual joystick should be displayed in a panel or separate window where actions on it will not affect the world, or a different navigation mode or control enabling and disabling anchors should be offered. It may be best to put a navigation widget in a moveable palette, so that left-handed and right-handed users alike could adjust accordingly and so that full-screen views could actually take up the entire screen with the navigation palette floating on top of the view. Keys for navigation (perhaps the keypad/arrow keys) and cursor control should be supported.

A nice widget that allows movement in three axes without requiring the keyboard is used in the DIVE project. Users can click and drag on any of three shapes, each which represents a different axis. VRweb, a world viewer, uses a similar apparatus in its “heads-up” mode and includes an extra tool at the bottom: a virtual “harpoon” or “seek” tool. A user normally selects the harpoon and then selects an object or feature that they would like to view. They are then pulled towards the object or face as if they were being reeled in by a harpoon. Different viewers allow users to fine-tune the harpoon’s actions: movement may be logarithmic, so one slows down as the harpooned object is approached. One may also be able to move perpendicular to the harpooned object - this is useful to do things such as look at textures on a wall. Another option may be to push away from, rather than move toward, objects. Keep in mind that harpoons should not instantly move to the desired view; there should be some movement in between so as not to disorient users.

33: Navigation palettes. The one on the left is a virtual joystick palette with direction and altitude dials; the right is similar to one used in Home Space Builder.



WorldView uses a virtual joystick element within a palette that can be merged with the main application’s window. This type of merging functionality is highly desired, as it offers one maximum configurability. Home Space Builder uses a palette of navigation icons to move around. Although the icons are not as standardized as other viewers’ and take some time to get used to, navigation is easy because the buttons are grouped so close together. Viewer developers should avoid grouping navigation widgets too far apart; it ends up annoying users if many movements have to be made for simple navigation.

Some viewers have different viewing modes, in which models can be examined in different ways. For instance, WebSpace has a normal walk view and an examine view, in which one can rotate a virtual sphere that mirrors the viewed object. One of the novel

things about WebSpace is its use of a virtual joystick as a 3D element - widgets are incorporated into a small dashboard that can be shown or hidden.


Useful modes for viewers to have are:

- **Walk** - this is a normal walking action and should allow one the six degrees of freedom. Mice with three buttons can be used to their full advantage in this mode; some viewers use the left mouse button to move forward, backward and turn; the middle to slide, rise, and sink; and the right to look and roll.
- **Examine** - this mode is meant to facilitate the examination and observation of a model; it is not meant to be used while within the model (such as in a house). Users may be able to use a virtual trackball that can be moved, thus moving the model. Some viewers use a virtual cube for this purpose; when a facet is double-clicked; the model snaps to that particular orientation. The cube is particularly handy for architectural or engineering applications. In this mode the emphasis should be on direct manipulation on the model, not necessarily exploration through it, although zooming in and out may be possible. Viewers should always allow for direct manipulation as well as simple walking.
- **Fly** - this allows one to fly within environments, useful for things such as flight simulation. VRweb's fly mode makes use of three mouse buttons: the left starts, aims, or stops the observer; the middle and right buttons accelerate and decelerate the observer.
- **Seek** - this mode changes the default selecting cursor into the harpoon tool, so that every click allows one to get closer to objects, without continually having to select the harpoon. This is a very good way to move across wide spaces quickly without losing control of one's direction. WebView has a seek mode that can be toggled using the "s" key.
- **Plane** - this constrained mode allows one to slide, rise and sink, and move forward and backward. Usually this is coupled with a harpoon that allows one to zoom in perpendicular to an object's plane in order to observe it. This mode may be particularly good for drafting and applications requiring flat views.

There should be some display or indicator that lets the user know what navigation mode they are in, what movements are possible, and what plane(s) movement is constrained to, if any. A small diagram or heads-up display text on the screen can help new users take advantage of different means of navigation more easily. Always the amount of on-screen information should be config-


urable, so experienced users are not distracted by visual clutter and so that inexperienced users can get as much information as possible.

34: Different ways of indicating types of movement and planar constraints.



It is highly useful to be able to change into a top-level, simplified map or radar view in order to get an idea of what and who is in the environment, and to get a bearing on one's location and surroundings. Some world modelling programs such as Home Space Builder and Virtus Walkthrough incorporate such views; the 3D and map views are linked so that one can see and manipulate the observer in the map view and see the effects in the 3D view, and vice versa. In the map view the observer is represented by a dot; the line connected to the dot is the observer's viewing direction.

35: WebSpace modes and the map view.



In some 3D games, the map is a simplified, shrunken 3D or 2D model that contains the environment in a separate window. Lucasarts' Dark Forces game shows the map in a heads-up display; although it saves space, this type of map display can be annoying if it obscures the world beyond it.

Viewers should support the generation or real-time display of maps; it is a great navigational aid. Map widgets should include a zooming function as well as a "fit to window" option so a complete view is readily available.

The ability to see one's relative location (x, y, and z) to the world's origin and some concept of north, south, east, and west should be supported; this makes it easy for people to jot down views and locations. One idea is to show clock numbers or orientations horizontally in the navigation widget so that one may say they are facing 3 o'clock or north-northwest. Some indication of speed and distance travelled may also be useful. It may be possible to create personal constructs (such as an odometer or dashboard) to do such tracking.

One might be able to use movement multiplier widgets, perhaps dials, that adjust acceleration, so that one can travel very far without having to maintain a constant maximum speed and so that very detailed movement can be possible. Movement sensitivity can then be adjusted to fit different user's preferences. One may be able to increase sensitivity by simply resizing a virtual joystick palette to make the area larger (and likewise decrease sensitivity by shrinking it).

The idea of using planes and grids to orient users is one of the oldest ideas in world navigation; I may travel on an infinite grid and "snap to" certain intersections and locations; three dimensional matrices may surround objects so I can more precisely track my movement and location (this is akin to using grids and rulers in CAD programs and the like). Grids are useful because they can give a sense of how big or small things are in relation to each other, and they can give a sense of location, allowing users to orient to them. A moving grid gives a better sense of movement than a surface covered with one color or even a simple texture, and the description of the grid geometry may well end up taking up less space in a 3D file than a bitmap texture.

In many virtual environments that use glove-based input, users must perform certain gestures in order to grasp objects and walk. Often users must point forward in order to fly. For new users this can be particularly disorienting, since in the real world pointing with one's finger means different things. If I point backwards, will I fly backwards? It may be a better solution to let glove-wearing users command a 3D input device to navigate, and use hand gestures only to manipulate objects, rather than try to force new meanings onto common gestures that may not be appropriate in the first place.

- VRweb, a viewer from IICA, NCSA, and the Gopher team can be obtained at **<ftp://iicm.tu-graz.ac.at/pub/Hyper-G/VRweb/>**
- WorldView can be found under **<http://www.webmaster.com/vrml/>**
- The home page for WebSpace is at **<http://www.sgi.com/Products/WebFORCE/WebSpace/>**
- SDSC's WebView can be found at **<http://www.sdsc.edu/EnablingTech/Visualization/vrml/webview.html>**

- Home Space Builder can be found under <http://www.us.paragraph.com/>
- Information on Lucasarts can be found at <http://www.lucasarts.com/>

Color and Shadow	<p>The use of color to orient users can be very important; most of us have gained an unconscious ability to orient a sky blue field above a forest green one - it just seems to feel better. One can create structures with different colored tops to show the preferred direction of "up", buildings can fade into the distance, giving one a sense of how far they are. One can choose to make the virtual sky change colors as the real sky does so that it serves as an indication of what time and weather it is in the real world locally (and it is much more compelling than a clock in a window).</p> <p>Shadows tell a lot of information about the world and how objects are positioned in it. They give one a sense of the dimension and location of objects; they can help orient one's position. Hide-and-seek games become more challenging when one must take care to hide their shadow. Shadowing systems that consist of many elements (like a file browser) give users important visual cues: Ray Dream Designer's world view is shown in 3D and is shadowed in three axes; it makes object creation and manipulation much easier than traditional systems with edit modes that offer only a separate two-dimensional window view for each axis.</p>
Teleportation	<p>When ones wishes to teleport to different areas, one's avatar may disappear from one place and appear in another. Because having objects and people suddenly appear and disappear just anywhere can be disconcerting, annoying, and inconsistent with a world's laws, one could create teleportation areas or reception environments such as a boat dock (where people come to meet visitors) and so on. The act of teleportation should be gradual (the teleporting user may see the world gradually fading away) so that other users are not surprised and so that lags in processing time can be made up for - transitions can then appear smooth.</p>
Camera Options	<p>The ability to control things such as perspective, zoom, and focal length may be desirable so that users can render and take snapshots using a number of effects and viewpoints (this may help bring about the use of virtual postcards - look where I've travelled...).</p>
The Home View	<p>There must be the ability to specify one or more default view positions for a world - these might be called home views or starting points - upon entering, users by default may go to the default view first specified in the world or may be given a choice. Think of default views as name tags in html: instead of specifying byte offsets to go to a particular paragraph on a page, you may say something like <code>http://www.foo.com/page.html#para5</code>. To go to a particular view in a world, one might say <code>http://www.foo.com/earth.wrl#view5</code> or <code>http://www.hilton.com/hotel.wrl#room1156</code>. Of course, one should be able to specify a location and direction for a view (perhaps among</p>

other camera information) when referring to a world via a URL; this is the parallel to specifying byte offsets when going to an HTML page (although addresses with this resolution can't be used on the Web yet). This makes it easy to link a doorway in one world to a room in another world - just specify the view. There may be a list of views in a view structure that could be defined in VRML, and there may be a view list attribute for every object if so desired, in the same way that an HTML document may contain many name tags.

When people teleport into a world (or when users fire up their cspace browser), they may go to any home view specified for that world. These home views become akin to addresses - I may choose to give anyone my corporate address, so they can visit the building, yet only give a few people the address for my house in my world.

It would make sense if name tags could be affixed to any node or set of nodes within the file, not just camera view nodes. Then if I wished to include a particular function or behavior from another file, I could reference the particular node or nodes directly for inclusion.

Camera Paths

It may be a good idea to allow camera paths within scenes, so that users may follow them to take tours and experience demonstrations. Here's a possible way to add path information in VRML:


```
cameraPath {
    name "United Nations Building Tour"
    pathInfo {
        ...
    }
}
```

My browser could give me a choice of paths to follow within a world; path information may consist of point, vector, or curve information and data related to velocity, timing information, lists of URLs, etc. I may be able to select a "smooth path" option so that my motion along paths defined by points and lines is not jerky or irregular.

Much in the same way I would be able to enter a world and go to a default view, I may be able to enter a world and follow a default path. For instance, the URL `http://www.foo.com/floor1.wrl#building.tour` would take me to a world and start leading me on a tour. I may have widgets such as sliders, play, pause, and stop buttons to play tours; in an editor I may be able to record paths using similar widgets.

If I can display the pathways with their names, I may see how to get to different places very quickly; by selecting any displayed path I may be able to “snap-to” the path and follow it.

36: Paths that can be viewed and followed make an excellent navigation aid.



Controlling Detail

Although one may specify different objects with different levels of detail in Inventor, it is a good idea to allow browsers to control detail according to one’s specification. There are a number of ways objects could be represented:

- Bounding boxes, no hidden lines
- Bounding boxes, hidden lines
- Points
- Simplified wireframe, no hidden lines
- Simplified wireframe, hidden lines
- Full wireframe, no hidden lines
- Full wireframe, hidden lines
- Flat rendering
- Simple shading

...and then any number of options could be applied: different types of shading and smoothing, texture map detail, transparent image detail, shadows, and so on. It is up to the browser to take into account such factors as window size, world size, viewer speed, refresh rate, network speed, processing capability, and proximity to objects.

The speed of navigation often varies with the complexity of the scene. As one moves from a detailed room to a more simpler one in many systems, movement sensitivity suddenly increases and users may find themselves suddenly rocketing towards an unknown direction. Browser writers should make an effort to compensate for this variation to some degree to avoid sudden and uncontrolled movements.

As one moves through an environment, the object detail should be reduced according to set preferences: I may wish to see high-quality renderings while moving at the expense of speed, or I may wish to move in wireframe mode. When one stops moving, image quality should be progressively defined.

Generic Objects When a browser loads an object or scene that cannot be read, something to indicate an error should be displayed. This may be a generic object, such as a cube with a question mark on each side. This use would be similar to the way that many browsers offer a generic icon or message when images can't be found.

Much in the same way that browsers and servers have generic icons for data types and ways to represent other services (such as Gopher or FTP) in HTML, VRML browsers will have generic objects and ways to represent such services in VRML. Although it is completely possible to represent such things in a traditional 2D manner using VRML, one should explore new ways to make understanding and use of such services easier. One can make 3D equivalents of international symbols and icons common to most 2D interfaces today; one simple way to map these icons into three-dimensional space is to use them as texture maps onto primitives or generic objects. To call attention to international symbols in a virtual environment, it may help to put the symbol on a bright background of the appropriate color (white, red, or yellow, for instance) and then instantiate this symbol on multiple sides of a primitive of neutral color (such as a medium gray).

Here are some possible ideas:

- To display a generic object: A gray primitive (cube, sphere, cylinder, cone, etc.)
- To display unknown data: A generic object with question marks on it

These could be used as simple defaults by browsers. The ability to specify replacements for generic objects, such as music-playing constructs in place of sound files, should be offered to users.

Because cyberspace is an international medium, some research will have to be done to ensure that its popular symbology can be understood by its various participants. Education can also play a large part in teaching new users any common visual language of the medium that may arise.

New Structures Of course, there will be many new ideas to be explored in order to create interfaces to systems that have no real-world counterpart. Should WAIS searches return a stack of cards? Why not use file cabinets or fields of spheres to represent directory spaces? I may explore databases by walking down hallways that fork off until I get to the information I want or visualize collaborative spaces as arenas (like a sumo ring) in which people can contribute and manipulate information and objects as well as converse with others. In common areas, participants can show off data, exhibits, and other constructs while allowing other users to view them. In a room with friends I may select a avatar's head to "see what they are thinking" or get their home page, resume, or business card.

Liquid Architecture When it is possible for constructs to create new structures, one can imagine dynamic environments that change according to any number of criteria. When applied to virtual habitats, this concept

may be referred to as “liquid architecture”, a term that probably originated with Marcos Novak of the Advanced Design Research Program of the University of Texas at Austin.

I may create a door construct that slides open temporarily when it detects an avatar within a certain proximity - the equivalent of an automatic sliding door in the real world. This is a simple example of dynamic structure.


Today it is possible to analyze Web server log files so that one can tell what are the most popular files, directories, and paths through Web sites. Intelligent programs can be created to rearrange lists and links so that the most popular information is moved to the top of the list or popular pages become easier to navigate to, or so that information that is useful to the areas people are visiting from is made more prominent on a user-by-user basis. Dynamic programs that change how they interact with others depending on their environment have what are called “soft” interfaces. When we extend this idea into three dimensions, one can easily imagine buildings that rearrange themselves to fit their occupants. Structures may change shape and color depending on the kind of contents they hold and the type of transactions that are occurring within them. A building that appears with many lights when many people occupy it can gradually get darker as people leave it. FTP sites could be represented as objects that look very hot when busy and cool when available; their size may be related to how much data is stored in them.

One can think of self-tuning information structures - a building construct that detects that most people bump into certain walls may move them to different areas. Structures can rearrange themselves to help users find and navigate through information more quickly. Say I’m walking through an expert system; I tell the system I’ve made a choice by choosing a particular fork in a road and walking down it. As I read questions that I am asked, new forks in the road appear, and so on, until I’ve reached my destination - a room with the most probable solution in it. I then can keep a map of my choices; the next time I use the system, my physical sense reminds me what paths I’ve taken before. Virtual roads could age with use; footprints may be left on the floor and fade out with time so I can see if I’m in a popular area.

Novak has been creating highly liquefied environments that dynamically respond to algorithmic changes. The spaces which are created can be written and experienced as three-dimensional

music and point the way to new media forms such as what he calls navigable music, disembodied dance, and habitable cinema.

37: One of Marcos Novak's works, a four-dimensional chamber.




Needless to say, the ability to create, play with, and learn from such structures can have applications in the real world: in particular, the fields of architecture, information science, psychology, and artificial intelligence may all benefit.

- 3D architectural works exist at http://archpropplan.auckland.ac.nz/misc/virtual_tour.html
- An interview with Marcos Novak, with pictures of his liquid environments, is at <http://www.altx.com/interviews/marcos.novak.html>

3.12 World Issues

If a world, scene, or object is completely readable it may be possible to grab them and collect them in other worlds. Just in the same way that it is possible to copy a Web site remotely by using a program to do this, it may be possible for a program (or construct) to copy, distribute, and move entire elements in cspace. This can make this easy for administrators and people who need to move worlds to different machines and can make the assembly of worlds by assembling different scenes and other worlds easier.

Portals



38: Movement between worlds can be made more realistic by preserving users' camera and velocity information and using default views for common entry areas.

When different worlds and scenes do not exist on the same server but must be connected (for instance, different rooms within a house), there aren't any easy solutions that allow one to simulate smooth, fluid movement between them. For instance, I may see a doorway with an inlined scene of another room behind it, but although people and things may be in it and moving I can't see this information (unless my machine has sufficient power to completely mirror that world). A region in a doorway may be a link to the next room, so when I hit it, I am transported to a default view of the next room. My orientation information may be preserved, so I can pass through this simulated doorway at an angle and not have my path changed.

When others view my avatar passing from world to world like this, they will see me moving towards the doorway and at some point disappearing (dissolving or better yet blinking immediately out of existence) completely. Others won't be able to see me through the doorway unless their browser mirrors the world beyond.

In order to avoid this inconsistent form of movement, one should create structures that make travel from world to world believable. Elevators are a very good solution in this case; I may pass into an elevator and hit a button (a link that triggers two elevator door constructs to move and close), and I am teleported into the next world with the appropriate floor. When I enter the world or move forward in the elevator, this triggers the doors to open, and I am able to explore. When viewed by others or with others, this action appears realistic and natural.

The creation of entry and exit points, more commonly referred to as portals, can be facilitated by making portals generic elements in the language, as is done in the Sense8 system. A portal object may have hooks to behavior, inputs and outputs, and how it appears when users look through it or how it appears on both sides, if the portal stands alone by itself.

It is not necessary to make a portal like a doorway, although this is the most common method of describing them. In one of the cyberspatial worlds that Rudy Rucker writes about, portals are spheres, so that people can fly into them from any direction.

Tiling

If I am able to specify a tileable 3D container in the language, then I may be able to create large environments with a minimum amount of source. For instance, say I make a 3D rectangular cube with a square as its base. I can make a tileable hill contour as well as a cloud cover. When the environment is generated, the desert and clouds can be tiled to infinity; as I move forward in the environment, tiles ahead of me in the distance can be generated and displayed dynamically.

I can also make shapes that tile not just in two dimensions, but three dimensions as well, so it is easy to make matrices, mazes, and worlds with multiple rooms.

3.13 Other Issues

Cyberspace Protocol

Although HTTP may be adequate for transporting simple read-only worlds with little interactivity, a new protocol must be developed to handle the new kinds of interactivity that is desired. Servers must be able to act as reflectors, so data may follow optimal routes among them and so that information can be multicasted. There should be a binary form of the protocol so that transmission can be as fast as possible.

- *Internetwork Infrastructure Requirements for Virtual Environments*, a basic overview of infrastructure needs and possibilities, was presented at an NII 2000 Forum in May 1995. The document can be found under **ftp://taurus.cs.nps.navy.mil/pub/NPSNET_MOSAIC/npsnet_publications.html**

Dynamic Building

If many users are building a virtual scene at once, how should the system know when to “publish” elements so that the changes can be viewed by everyone? File locking mechanisms may be used to make sure that users don’t edit any single object at once.

VRML and CCI

The ability to control Web and VRML browsers with CCI is a step in the right direction: if one can make VRML references with the ability to specify camera location and speed (using default views and paths, perhaps) then simple tours and 3D slide shows could be viewed. If an interface language were well developed enough, then one could create stand-alone world editors that use VRML browsers to do the rendering, communicating with them via CCI, or traditional MOO systems that used CCI to transmit world information to cspace browsers.

- The CCI specification is at **<http://www.ncsa.uiuc.edu/SDG/Software/XMosaic/CCI/cci-spec.html>**

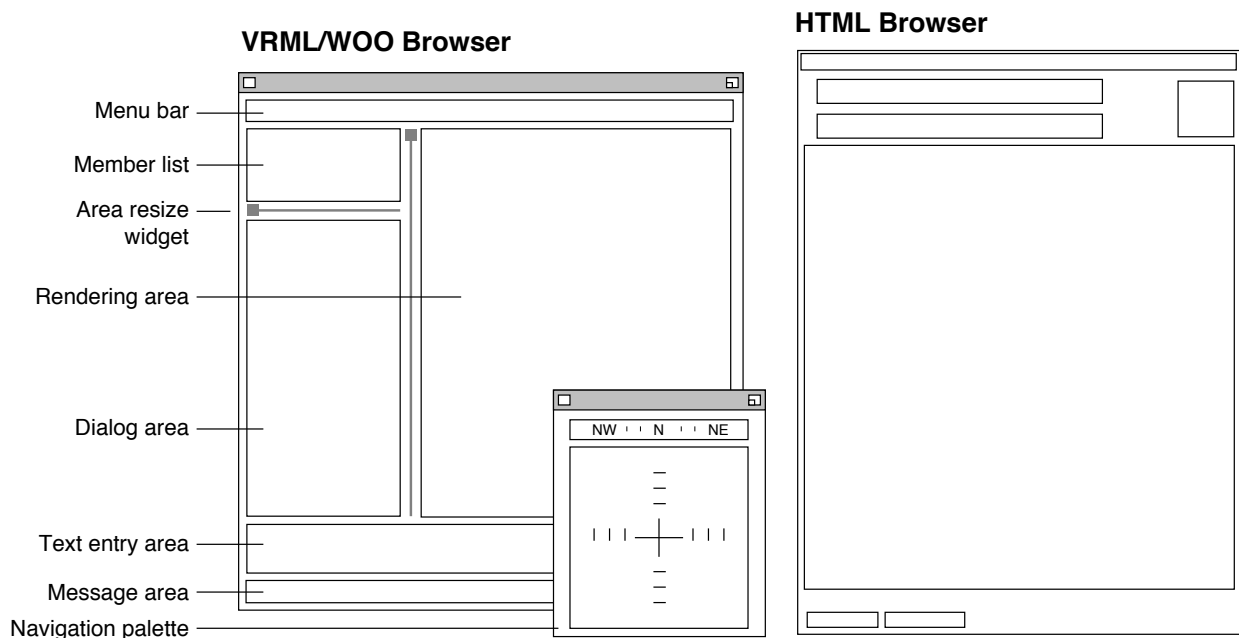
3.14 A Prototype VRML/HTML System

A prototype interactive VRML/HTML system should consist of a number of basic parts. They may each be applications in themselves, modules, or palettes, and could be combined in a number of different ways:

- VRML display area - This area displays 3D rendered objects as required by a process - it may help simulate movement throughout virtual spaces or depict static scenes as requested.
- HTML display area - This area displays rendered HTML documents as requested by a user or another process.
- Navigation widget - Palettes may include VRML navigation and orientation widgets and elements common to HTML navigation (back, forward, abort transfer).

- Menu widget - This is an area in which program options and user preferences can be indicated.
- Text entry widget - This element would allow users to type in text so it could be sent to other members within a common area.
- Dialog area - This is an area in which users can read what other people and themselves have typed and possibly what actions they have performed.
- Participant area - This is an area that would hold a list of members within a particular space and/or would show their respective icons and other status information.
- Message area - Using this area, the browser can display alerts, updates, and other information (location, address) that it may need to convey.

Perhaps the best way to visualize such a system would be to imagine a VRML browser coupled with a WOO. A fully featured system would at least allow users to send and receive any manner of real time media such as video and audio, and would incorporate WYSIWYG HTML and VRML editing capabilities.



39: The layout for a VRML browser with basic WOO capabilities. It may interact with a normal Web browser or each application may be palettes within a larger framework.

Browser Capabilities

There are a number of things that browsers should support that range from basic to advanced capabilities:

General Navigation

- Open local files and URLs

- Back and forward, history, and home options
- Hotlist with sample default links
- The ability to set a current scene as home

3D Navigation

- Walk, examine, fly, seek, and plane modes
- Generic navigation and harpoon tools
- Virtual joysticks, trackballs, or the functional equivalents
- Palettes, possibly configurable ones
- The ability to specify absolute motion or motion relative to one's viewpoint

Window Options

- Status area
- A progress bar or animation

Model Features

- A menu of multiple viewpoints within the file
- Ability for user and author to set background
- Ability to show node labels in the model view
- Ability to show a file's title and URL
- Ability to show other document information: inlined worlds, file comments, SceneInfo nodes, number of polygons, language version.
- Ability to show parser warnings

Rendering Features

- Different rendering options: wireframe, hidden line, Gouraud, and so on
- Automatic generation of normals
- Automatic level-of-detail
- Configurable degradation during movement
- Ability to show frame rate
- Collision detection
- Constant velocity control
- Control of speed of movement
- Camera focal length control and zoom
- Dithering
- Ability to look directly at scene with no specified viewpoint
- Ability to change how anchors are rendered

- Ability to change brightness of the world

External Viewer Features

- Support for .mailcap and other standards
- Saving files and other objects
- Ability to save worlds in different formats
- Ability to print
- Ability to view the source of a world or a link

Miscellaneous

- Accelerator keys
- A help menu
- Support for multiple types of input devices
- The ability to automatically uncompress files

Visions Here's a scenario: I have a text entry space, the contents of which I can send to any one whether they're in or not, any one within my sight, or any object. If I hit a mail widget, I can mail the text; I can save it, or use it as keywords for doing WAIS or Web-Crawler searching.

I go forward; my 3D environment moves with me. I can drag things around, keep them in a personal wallet, tag along or hang out with friends watching plays. And as people move on the virtual stage, I can follow the script in a HTML viewing container somewhere in my space, or see it using some generic Web browser. If the play were run by a program, I could fast forward it, make changes to the script, talk to the players and have them do different things. If someone else wishes to rewrite the script with me, fine; we can gather at my personal site somewhere and make the changes to a personal copy.

If I want to give this URL to a friend, I can type some words to him, drop a link on a selected phrase, and send him the hyper-text, or I can put an object on the screen, a big red arrow (I have a few different types in my tool box), make it rotate around, follow him, and blink so he notices it, and when he does, it tells him to click on it; it's a note from me. If he grabs the arrow, he can go to that place I wanted, or maybe I can send him a copy of the annotated script with my changes.

So I go over to a bunch of friends with the corrected script; there's a PostScript image in it which I shove through a VRML filter and float above everyone else's heads. I see some folks off in the distance (my binocular sensitivity is set pretty high), so I give them a ping to see if they're interested; they zoom in for a second to take a closer look but decide to hang out at IUMA watching QuickTime clips. I listen to alt.scripts for a little while and keep a few messages in my wallet; I've offered to talk to some posters for a bit at their site to pour over some actors they've constructed.

When I get there I find their place a bit messy so I get rid of some of the detail, like their wallpaper, and get a bird's eye blueprint of the place. A forwarding object tells me they're actually at someone else's site, so I punch it and when I get there I find they've already built the stage. They send me some actor programming if I feel like working on it. I'm not sure how it works so they point me to online documentation with a page of MPEGs showing what the results of certain functions are and some simple actors to get started. I figure I could talk with them to get more help, but I don't feel like doing a face-to-face today; it's a bad hair day. I decide to use large text over a calm JPEG of a lake to save bandwidth.

A gaggle of students on a field trip from Monte Vista High (their school logo trails after the group) passes by in the distance. Their chat is public, so I tune in and find they're researching Shakespeare. My biff follows me around shouting that I have mail. I touch its nose to tell it to shut up. My first message is a gripe from a user who says my "Hello, world" program crashed when he ran it. There's a pointer to his world, so I go to his public section (his private area has big black and yellow construction borders all over the place) and find my code with the buggy pieces highlighted. I put a message in the source telling him to run it with a certain flag, wrap it all up in a self-extracting object, and leave it in a corner that looks like it has been fairly well-traveled (there are a lot of footprint outlines on the floor, fading by the second). I send mail to the guy so he knows where to find the code again. If he has a biff going, it will find him and let him know, but I decide not to wait around for him or his utilities to show up.

I notice a little cube with a nice texture map on it, so I ask it to open itself up and I notice it has no programming and no links. I try to grab the texture map to put in my wallet, but the cube refuses (I don't have permission), so I take a snapshot of one of the sides instead.

There's a small shelf of FTP sites here too, the bigger they are, the more they hold. I copy some of the objects to my wallet for saving. I don't care how much they hold, though; so I sort them by last date modified and the sizes don't change much.

My next message is a pointer to a magazine and a list of folks interested in working on the next version; I go to the first GIF on the page, browse the "objects to the editor" section, and link to the mailing list. When I get there I find people are adding all sorts of things to it and it looks like a junk pile. I sort out the PDF documents only and send them home for reading tonight. I close my wallet (I don't care much about leaving it around my public place, since it's got a pretty long hash key on it) and shut off my Power-Book.

3.15 Cyberspace Guidelines

A cyberspatial etiquette will no doubt evolve, just as there are guidelines, do's and don't's for webspace and MOOs. And because there is great potential for chaos, it may not be a bad idea to create principles for cyberspace - a sort of ethics, much like the hacker ethic. There will be so many people participating from so many places, ways to organize societies and bring about communication will be very much needed.

Below are some rough guidelines that follow a basic philosophy: when people create new interfaces and worlds for people to experience, things in the environment should be consistent in their use and actions. Otherwise, people may get confused and disoriented. You expect doors to open and close in real life. When one sits down, they don't fly through the ceiling. Although creativity should be encouraged (you can build anything you like in cspace, after all), public areas by default should mirror reality as much as possible; exceptions should be clearly marked (they may be signs saying "zero gravity area here").

Again, these are only guidelines and not rules - but if cyberspace architects follow them, we can create common spaces that are more liveable and usable.

- Create, but don't confuse.
- When in doubt, keep things real.
- By default, your avatar should be opaque to all.
- It is not polite to pass through someone.
- If you do pass through someone, your appearance should become transparent or changed in some way to denote this action.
- The appearance of objects should not be changed when they are passed through unless they have been programmed to do so.
- It is best to be yourself; otherwise, assume a neutral shape and gender.
- It is impolite to teleport into and out of places unannounced.
- For short distances it is best to walk.
- You can only exist in one place at once.
- Users should be made known in some way to all others in the vicinity.
- Users may decide who to see and display according to any criteria.
- Identical objects cannot be in more than one place at once. Aliases, agents, and symbolic objects must be visually distinct.
- It is best to keep structures within a consistent range of height and dimension.

- For places meant to be visualized by others, it is best to stick with Cartesian space.
- For structures meant to be traveled by others, it is best to stay away from overlapping, extradimensional constructions.
- Objects should appear the same to some extent to their viewers.
- Users cannot sense into domains in which they are not allowed. Only owners of objects can change their outward appearance, unless such objects are in the public domain.
- Users should be close to each other in order to converse, unless a virtual communications medium is used.

What is the search for the next great compelling application but a search for the human identity?

Douglas Coupland
Microserfs

4

Living In Cyberspace

4.1 Next Generation Environments

The task of creating a system through which our minds can manipulate information efficiently is one of the greatest challenges of our time. It is not necessarily about designing the right interface for the task, rather than about developing the right environment. "Interface" implies that an invisible barrier exists between the computer and the user and emphasizes the barrier. "Environment" emphasizes the sum of the experience and the explorability of the system.

If you look around your environment, full of devices, chances are that something is not configured quite right. Perhaps the clock on the VCR is blinking, or you have trouble finding certain buttons on the remote control, or you always end up locking the car keys in the car. None of these things are your fault; each of these problems can be easily corrected by a little thinking on the designer's part. There is no good excuse for bad design in the things that we use every day; if we cannot devise an appropriate computer environment for people to work in, we must try again and again until we get it right. It is an evolutionary process.

It is a given that systems will be complex in the future and that they will be fast and getting faster. Designers should create with an eye towards allowing people to use the computer both as a medium and as a tool. The environment of tomorrow will be our remote control, our VCR deck, and our movie theater. It is an ultimate environment.

The new environments should take advantage of the things that humans are good at, rather than trying to let the computer completely take over these functions. Our brains excel over machines in doing things as recognizing patterns, evaluating information, and providing a sense of context for what we experience. We are so good at all of this because we do it routinely every day; humans have been doing such stuff for a very long time. New systems should assist users in these tasks so that their use becomes a natural part of operation. When we create systems that try to completely mimic this behavior, like pen-based handwriting systems that assume to know everything about your penmanship beforehand, they often fail. We need devices that we can work with, machines that can learn about us in order to help us work with and learn through them.

Certainly cyberspace will give us more options than any medium ever has before. Because of this, it may end up stretching our innate pattern-matching and evaluation abilities. If so, an important question to ask is, will it help us think better? I think it can, if people are properly educated in learning how to learn. Without this important skill, learning in cyberspace would be like giving the Encyclopaedia Britannica to a person that can't read.

Right now almost all computer users are accustomed to the WIMP interface: Windows, Icons, Menus, Pointing device. The idea is not necessarily to get rid of this environment at once, but to take its best points and merge them with new ideas.

In time the difference between "interface", "environment", and "world" will blur. We will not be using an interface, but an information workspace that will make use of aspects of architecture, environmental design, and psychology, among many other fields.

The goal is to create an environment in which things relevant to users can be found and manipulated out of extremely large pools of data, in which information can be visualized in an efficient manner, and in which users can create, develop, and share many kinds of information and experiences easily.

Jakob Nielsen has noted a number of characteristics of the future computing environment:

- **There is no syntax** - things will be accomplished by manipulating objects directly, rather than by having to memorize a particular syntax or language to get things done.
- **There is no taking turns** - users will be able to act on processes and feedback in an asynchronous manner; dialog boxes and background events will not interrupt processes.
- **High bandwidth is used** - the environment makes use of high-resolution, true color displays; multiple media streams and dense information structures can be displayed and manipulated quickly.
- **The environment is all over** - displays are integrated into the physical environment; computing is ubiquitous.
- **It knows about you** - the environment learns one's preferences and is able to guess what tasks to perform and how to present itself.
- **The environment is controlled automatically** - the computer handles many tasks invisibly, rather than have the user directly act on many objects.
- **One programs by demonstration or graphically** - programming by example and visual programming languages are made use of.

- **It is not application-centric** - tools, modules, and media make up objects in the environment, not applications and files.

Although computers will in time make use of lots of bandwidth and high resolution displays, it is important to also concentrate on building environments that are appropriate to today's technology. Trying to represent systems that are too complex for the hardware can lead to failures, that is, the system ends up being too slow or crippled in order to run. This kind of implementation then discourages people from thinking about how similar systems can help them. It is important to concentrate on doing fast, low-end things with what we have rather than do slow, high-end things.

The new environments will not be application-centric - one does not have large programs that do specific things, but tool sets that can be combined to do anything. Instead of document-like files, emphasis is put on creating and editing any sort of media, media that can be merged or taken from anything else in the environment. We should follow Ted Nelson's urgings not to simulate paper - we might as well be simulating etching on clay tablets.

Likewise, Nielsen urges developers to stop using the book metaphor - it limits hypertext's search and navigation potentials. With the computer, any kind of media structure is possible. Why do windows have to be impenetrable and solid? Why does media have to be stratified and separated? Why should browsers be different than any other application that handles words and layout? Why can't I have hypertext and hypermedia in everything I create, whether it be a spreadsheet or calendar? Why can't my calculator program do things in my CAD program?

Instead of programs, we then must think of tools, and how they can cooperate and interact with each other. The best scenario is that they interoperate with each other, popping up and disappearing in a seamless fashion depending on what one wishes to do. There is no "switching" between "applications", some of which work with "files"; the emphasis should be on doing things to the media you're working on. One day, as I work on something, I may discover I need a certain tool. I can tell my environment to search over the network for this tool, and it will look around until it finds one. I may be able to subscribe to the tool as a service, so I can use it for a certain period or for a certain amount of time, or I can purchase it outright for myself. In this manner people can build their own environments and tools.

Mentioned below are some of the features that future environments are likely to make use of, in one form or another.

Parts In creating new environments, it is helpful to be able to break everything down into basic parts. Most computer environments contain both programs and data: programs act on and with data as well as produce it. Data are your things, your media, objects in the environment. Here are some of the basic things that people like to do with data:

- Create
- Delete
- Find
- Edit
- Manipulate
- Experience

All of these things can happen passively or interactively. Other people can do these things with us and for us, or the computer can do these things to save time and make life easier.

For instance: I may wish to create something. I might use templates and defaults to help me. Or I may ask the computer to make something for me. Or I might do this thing completely from scratch. Or I could get this from someone else, ask them (perhaps pay them) to do make it for me, or have them help me in doing it.

If I wish to find something, a certain piece of media that I or someone else created, I can ask my computer to look for it, or things similar to it. I might sketch an object and be able to retrieve files off the net that look similar to it. I might search manually, or ask a searching service, or other people that may have encountered similar media.

If I wish to edit something, my computer may be able to make changes that it knows I will agree with (such as correct for bad spelling). It may find and display all the various editing tools I'll need for the job. I may ask someone to assist me in editing.

Should I want to examine or organize my material, I may have many options for viewing and sorting. While I'm doing other things, the system can perform many background tasks: cross-referencing, indexing, publishing.

I may be able to experience information as it is sent to me or after I have received all the necessary data. My computer can set up the appropriate viewing modes and filters according to my preferences. I can share the experience with someone else or other groups and send live streams to them as well as canned media. My system can analyze and view media in the background and notify me when it finds something of interest, or compile packages of possible interest for me to view from time to time.

Note that information retrieval, perhaps a far too general term, plays an important part in the entire picture. What is also important is the ability to create and handle data. If information retrieval can be considered a science, why not information creation and information manipulation?

Connectivity

In time, everything will be connected. The question of where becomes unimportant. Resources are where you find them.

In a society in which people can send and receive, create and view, it is our computers that will serve as gateways to the world. They will be open to the networks and on all the time; they may be continually connected or they might sleep, waking up only when something needs to be received or sent out. The system can constantly be searching, indexing, maintaining, asking other people and systems of things, and receiving answers.

It will be unimportant where one's central processing power is. At some point, CPU cycles may be distributed - other people and processes may be using your cycles when you're not doing anything. Or you may want every cycle for yourself, and disable this kind of sharing.

Much work will be done not just by collaborating with people, but also with the things that people have made - agents, processes, and other media. Certainly people will have a greater opportunity for collaboration, but they need never exercise it if they do not wish to.

This type of connectivity has a philosophy associated with it: one does not concentrate on where; one concentrates on everywhere. It will be trivial to provide everything to everybody; it is more useful to be able to provide particular things to certain people that need them. The ability to tell people what I have and ask for what I want become necessary skills.

Pointers Almost every system that has ever been made to help people find things has used pointers - things that represent data, that point to it. Cards in a library catalog reference the original media; pointers in a C program reference memory locations. The advantage of using pointers is that groups of them can be manipulated and viewed more easily than what they represent. They are most useful when the original media can be easily accessed with the pointer. In a library, there is no guarantee that I can access the media after finding its pointer, its reference. When both pointers and media are online, the original media can be accessed very quickly.

Pointers serve as shortcuts not just to media but also to actions and ideas. Many systems make use of pointers under various names: symbolic links, aliases, shortcuts, shadows, and hotkeys (accelerators). Any time I can represent a certain media, action, or structure with a pointer, the potential exists for me to manipulate and view more things easily.

A simple example of using pointers in today's systems might be using a pointer to a remote service such as FTP. The pointer is represented by an icon and a name; it may be labelled "Aunt Betsy's FTP" with a picture of her computer as the icon. When I perform the action of opening her icon, I may then see a view of her file space, in which I may upload and download things. Or I may be able to drag media onto the icon, sending my things to her system automatically. Or, as I hold my mouse button down on

the icon, I may be able to navigate the file structure and retrieve things using a hierarchical menu, saving a good deal of screen space.

Instead of having to know the protocol, or the network address, all I would have to remember is the icon and/or the name, and where the pointer is located. If I can simply tell the computer, "Aunt Betsy's," I don't even need to know the location. We will have so many network services coming in and going out of our systems, it is essential to be able to symbolize as much information as possible using pointers that we can understand. My online address book should only have names and information that helps me to remember them; the computer can take care of things such as where to go and how to get there. The phone number can be a thing of the past: instead, we simply exchange pointers.

Labelled icons in today's environments can serve as pointers to external storage devices. Because these spaces are functionally equivalent to internal storage, they should be able to be treated and viewed in the same manner. Systems that label storage devices with arcane names (such as A:, B:, and C:) and expect the user to read them are hopelessly backwards and should be rethought.

Customization

Everyone has a different view of the world; the idea behind customizing environments is to give users environments that are comfortable but that still allow them to operate with the rest of the world. Customization allows people to discard or change elements they don't like; it also allows people to use ways of interacting with people and things that they find appropriate.

Elements that the user is expected to deal with often should be customizable in an easy manner. One may be able to change how objects look, act, and respond. We may be able to see different views of information, sorted by any number of criteria. Our environments will be extensions of us, systems that we find familiar. There should be a number of common methods for sharing resources and information with other people, however, so that communication can flow easily.

Computers are devices that act as extensions to our brains: we adapt new systems so they become like us; network protocols and media formats are the languages connecting ourselves with everyone else. We must always have common languages that we can use, or the appropriate translators.

Because computers and their I/O devices can be everywhere, the concept of having a personal computer becomes less important than having a personal environment. I may sit down at any display device and continue my work; the important thing is to be able to transfer information about environments among devices. This requires a common method of obtaining preferences; if I want any service to know about me, it should be able to access my profile in order to better help. Future systems should provide a standard way of receiving and sending preferences over the net and within themselves.

Defaults When users start out with new environments, they are like blank slates, ready to be filled and customized. Although users should be able to customize things how they like, they should also be provided with a set of consistent and professional defaults to serve as starting points. These are environmental templates.

Things such as defaults, guides and hints help the user become less naive; they guide users toward suggested solutions and help propose new ones. They should allow users to be more efficient with the tools they have.

Users should be given a choice regarding defaults: there may be sets of defaults appropriate for new users or experienced users. Some will not need hints; others will need all the guidance they can get. It has been written that the appropriate starting point of view should be the one that gives the most information, in order to present an overview of things.

The system may have default rules of behavior and rules regarding layout. One common example is the “snap-to” 2D and 3D grid: it allows media to appear consistent and orderly by aligning their positions. Default colors in the environment may be tuned to facilitate working, or to enhance particular elements.

Consistency Environments with a few consistent rules of behavior and layout are easier to understand: there are less elements and rules to learn, and less surprises. Things do as you intend them to. A consistent environment allows users to ignore everyday things and concentrate only on the important differences.

To begin with, environments that are made to primarily visualize and manipulate information should be visually clean, collapsible, and static. Users should not have to deal with a wide array of default fonts and elements. The system should have a recognizable look that enhances its uniqueness.

Behavioral consistency can be achieved by applying a core set of rules to every action; these are constraints. These constraints may be explicit, observable things, like graphic boundaries that can't be crossed, or implicit, inferred rules such as “you can view a document by clicking twice on its icon”.

Icons are good at depicting objects: when you encounter an icon of a page with a bend in one of the upper corners, chances are that it's a file. Directories often have a consistent look to them.

Semiotics, the study of signs, classifies such things into three categories. An icon is an image of the thing exactly. An icon representing a train will be an image of a train. An index is a suggestion of the thing; an index representing a train might be a picture of railroad tracks. A symbol is a conventional representation of an object, some particular image that has an agreed meaning. If convention dictates that an image of a circle represents a train, then the circle is a symbol.

Within systems, sets of signs should consistently be either icons, indices, or symbols - mixing the types can lead to confusion and misunderstanding.

Feedback Feedback at all levels must be provided in the environment. Completely retroactive feedback methods, such as dialog boxes that force the viewer to acknowledge something, are not things that should be added to the environment as the last step in development. They should be integrated into the overall design of programs. Because users will most likely be running many processes at once in tomorrow's environment, the type of feedback that is needed will be continuous and intrinsic to each element. Users should be able to glance at an item and quickly get the information they need.

One example of this is the scroll bar - rather than simply moving the screen within a window up and down, it also gives an indication of where one is in the document as a whole; it provides a quick birds'-eye view of things. Another good example is the everyday flexible tube of toothpaste - at a glance, you can tell about how much toothpaste you've used up by the size and shape of the tube, and thereby know when you'll need to buy another one. Newer, stiff, opaque plastic tubes have no such indication; no feedback has been incorporated into the system.

In an environment where many tasks will be occurring, some standard method or area for providing feedback should be available. Today in many windowing systems this feedback is particular to applications, so people must switch to a certain application to see what progress it's making on something. If one is working on something else, often this feedback from other background processes is obscured, grayed out, or in an inconvenient position for checking up on. One possible way to solve this dilemma is to set aside a particular area in the environment - a tool bar, a region, a status palette - for showing the status of applications as tasks. If I send a program out on the net to find something for me, begin backing up my hard drive, and convert a series of graphics files into another format, I want to know what is happening in each process - where it is, so I may be able to cancel the action, observe it, or get a good idea of when the task will finish.

Another reason to provide continuous feedback is to allow people to watch something while they are waiting for a process to finish; people are becoming more impatient with their online devices. This is because users are becoming used to seeing faster technology, and thus have the expectation that things will always be fast; this is also because the speed of our devices is slowly catching up to our speed of thought. As the bandwidth of cyberspace increases, people will expect it to take on more characteristics of mental activity; eventually the gap between the two speeds will close. The smaller the difference is, the more we try to mentally complete the distance between action and thought, leading to impatience.

The Importance of Finding Things

Only a few years ago, outspoken members of the hypertext community felt that the idea of being “lost in cyberspace” was a myth: the fact that one could “get lost” in a hypermedia system was of little importance, unproven, or missing the point entirely, since users that viewed such systems using navigational models were conceptualizing everything the wrong way anyway.

It has been shown that users of many types, new and experienced, use navigation terms in referring to experiencing different media in hypermedia systems. One must never blame users if their mental map doesn’t correspond to what the designer intended; *it is the designer’s job* to make the kinds of systems that do not fight with people’s expectations.

Perhaps one reason that people form this mental model of navigation when dealing with hypermedia systems is because of one’s short term memory - as a user browses through the Web, some sort of mental image of the last visited pages is kept in mind, and this is likely to evolve into some sort of linear trail in the mind’s eye that one can move along or branch off of. Users may wish as some point to experience a prior piece of media again, and thus they move backwards along this line: they “revisit” a document, “go back” to a page “where they came from”. The questions users are most liable to ask themselves while browsing is “Where did I come from?” and “Where do I go?”. Both of these questions can be better stated as “How do I find this?”, and to answer this question elegantly is one of the Holy Grails of user interface. And the Holy Grail is something that people have been trying to find for centuries.

The case can be made that the node-and-tree structure of much of today’s hypermedia forces people to think in navigational terms, but even if every piece of media on the Internet were unrelated, it’s likely that people would still be trying to “go back to” things and “get somewhere”. Time, for all intents and purposes, is linear, we live sequential lives, and tomorrow is tomorrow as surely as today will be yesterday. One of the great myths of hypermedia is that humans are non-linear creatures. The truth is, we are both sequential and non-sequential. We may often think and associate non-linearly, but we live one day after the next. Humans can cause great chaos, but they have tried to impose order wherever they exist. The bottom line, then, is that any system we use to create and find things must allow for both of these modes: it must integrate the two as surely as Yin and Yang.

Hypertext was created as an alternative to linear text, not a replacement. The prefix *hyper* connotes extensibility, something extra. People will always want default modes and suggested order, just as surely as they will wish to rearrange a work to fit their own views. A good example is online books on the Web: many are split up into multiple pages, with each page having “previous” and “next” links. By selecting these links, the user can follow the default, suggested order of the work, but they can also read it in different ways as well.

Internally we may think in webs, but externally we tend to think in hierarchies. This may be a symptom of the real world: our minds free-associate easily, but in the physical world we have to deal with organizing things, putting things in places and arranging them where we can find them. This is why humans make piles of paper and stock bookshelves. Externally, we organize by location, subject, color, and so on precisely because we cannot access everything instantly with a mere thought or recall it in memory right away.

Often in looking for things people don't necessarily recall the exact thing, but rather something related to it. They then follow these relations until they have found what they are looking for. Authors of hypertext can help readers discover resources by offering related material: "If you liked this, then try this." With related suggestions, users can find what they are looking for more efficiently than by random browsing alone. If it is true that all things are related by only six degrees of separation, then why is it that Web indexing engines and related agents don't make use of this rule? Many don't index deeply enough.

Many kinds of research in different fields have corroborated that there are basically three kinds of users:

- Those who are wandering around with no particular goal
- Those who are wandering around looking for interesting information
- Those who are looking for something in particular

These types are what Mark Heyer refers to as grazers, browsers, and hunters, and usually people switch among all three types from time to time. The net should support them all: those grazing may wish to find passive media, broadcast media or online galleries; browsers may wish to look at subject-oriented areas and try interactive types of media; hunters will want to use search facilities, maps, outlines, online help, and indices to find what they are looking for.

When people have found something, they often ask themselves a few questions: "What do I do now that I've found this?", "Is this what I want?", and "Should I stop looking?" Environments can make answering these questions easy. They can allow the user to do many things with the resources that they find: save a local copy using various methods, send it to somebody else, keep a pointer to it, or edit it. They can help the user determine if they have what they're looking for by offering pointers to similar resources or material, offering them more choices to continue searching.

It is important to be able to find what one is looking for on many levels: in the immediate area, among all documents in the environment, among information on the network. The goal is to make


a scalable, consistent solution that makes the Internet and all its resources seem to disappear, by making everything seem so easily obtainable.

J.C.R. Licklider estimated that people spend 85% of their time finding information. Anything that can be done to get people to their goal faster should be examined; and every information environment should incorporate some method of doing this.

Programming

With every generation the average computer user becomes less naive. Introductory computer classes are taught in most K-12 schools today; the public can easily find such classes being taught in many places. Languages that were once only taught to Ph.D. candidates are now being used by ten year-olds; thirty years ago it was ludicrous to think that people might want to learn how to touch type to use computers. It is always the older generation of technology users that must play catch up to the ones accustomed to new devices.

Environments should give users the ability to program and customize elements, change the system, create new tools. At some point people will expect this capability. In the kinds of environments that are likely to appear in tomorrow's cyberspace, what will programs look like, how will they be created?



40: Elements in a three-dimensional environment may be able to be opened so that their programming can be accessed and manipulated visually.

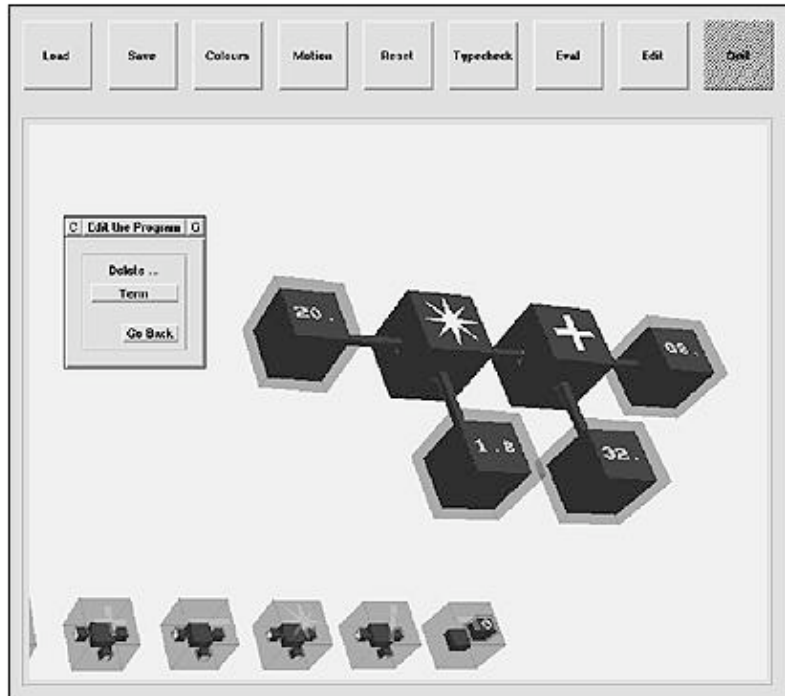
Programming may in large part be done by example and will likely be highly visual. Some of the core concepts in programs are things such as constraints, events, functions, input and output, pointers, and logic. All of these things can be displayed and assembled in a graphic manner. In a three dimensional environment, for example, a program may consist of a shape with input and output plugs in its side. One may be able to go within the shape and view other constructs representing functions, each connected by pipes representing the flow of logic and data. Filters and constraints on data may be represented by transparent, spherical membranes that resist or keep in data unless it meets specific criteria.

Why would we want to use a three-dimensional programming language? Marc Najork, the author of Cube, the first such language, quotes the following reasons:

- Graphic information can be understood faster. Programmers often make flowcharts and diagrams to visualize programs; graphic images and shapes can be remembered and grasped more readily than textual functions.

- There are more dimensions of expression. Images can encode more information than text; program areas may have variations in color, shape, size, depth, and transparency.
- Pictures can provide concrete metaphors.

41: A screen shot of a Cube program. Note the use of transparency to denote containers. Programs elements are cubes connected by pipes which represent the flow of data. This program converts Celsius values to Fahrenheit and can do the reverse calculation by reversing the data flow.



One drawback of traditional two-dimensional visual languages is that they usually do not efficiently use screen space. 3D programming can help alleviate this problem because it can make use of the extra dimension; it can be viewed in multiple ways to see obscured or hidden parts. Languages like Cube can certainly convey more meaning than 2D languages; one subjective advantage is that they may be more visually aesthetic, allowing one to create memorable works.

A great deal of work needs to be done in this area to fully explore the potential of this field. It is possible that better programming environments can be made by combining aspects of text-based and visual programming, for instance, I may be able to visualize parts of a text program within floating palettes that are color-coded and interconnected; I can visualize the entire program as a world in which I can zoom in and out and explore functions.

Today's programming environments should make use of hypermedia, allowing programmers to select a function call and jump to the function, drag and drop media and graphic elements into functions, and navigate code in different ways. It would be nice to be able to collapse a program and its documentation into an out-

line, so it can be viewed as an abstract. By controlling a widget, I may be able to slowly increase the detail in the outline until eventually I have the entire expanded program.

Three-dimensional programming will be best done with immersive displays and devices that take advantage of all six degrees of freedom. Such work today is slow and awkward when using traditional equipment such a mouse and monitor.

It Knows You

The environment should be able to learn about its users, learning their preferences, their movements, their habits, things they are likely to do. By doing this, the system can perform many tasks that users would normally do, saving time and satisfying expectations.

For instance, I should be able to toss and fling things to different areas using my mouse: the environment should know where I like to move my mouse; I may have set hot spots that I can automatically jump or accelerate to without having to move my mouse manually across a large screen.

Our current user interfaces have the great potential to be cluttered. When I look for things, the system should have some memory of the things I have looked for, how I like to organize my material. All of these things can be stored in a preferences database that can be transferred to different systems should I wish.

One example of a helpful agent-like program is Brenda Laurel's Eager - as one works, Eager looks for patterns of action and tries to complete them. If a pattern is detected, an icon appears that the user can select, completing the sequence.


One could imagine if other devices had similar capabilities: Compact disks could contain writable tracks for preferences; when I put a disk into my player, my graphic equalizer would be set, the sequence of play would be customized, and the volume level would automatically adjust to how I like it. Because user actions are not always defined easily, the use of fuzzy logic in devices can go a long way in helping decide things to do for the user. Machines should be able to know their users: "If she did this, it is likely that she will do this", "He has indicated that he prefers this."

In developing systems that do this, the user should always be able to take the initiative to do something first; environmental support should be optional so that it can be overridden at any time.


Fantasy and Reality

In writing, authors can choose to create linear texts, in which case they have the most control over the action and sequence; or they can choose to create completely non-sequential hypertexts, in which it is intended that the reader have the freedom to read

any part in any order. Hypermedia should be created so that its



42a: A scale of extremes for hypertext.



42b: A scale of extremes for virtual environments.

initial state is always a default; in a traditional novel, I may be able to read it in a suggested sequence, but the author should allow me to rearrange it and read it in different ways. In something like a reference manual, where I can read any part out of order, the author should provide some suggested linearity or sequence; this is often done in the form of pages and page numbers.

In designing environments, a spectrum lies between fantasy and reality: authors can give the reader many alternative possibilities, allowing them to do things that would be impossible in the real world, but they do so at the expense of confusing users. Authors can also make highly realistic environments, so that they behave and look real, but they do so at the expense of constraining the user. Defaults should be created so that users can slide between the two extremes: they may be able to change the physics or create impossible objects. One should be able to control the degree of magic in the environment.


Of course, each extreme can be well-suited to particular applications, and it is not necessary for authors to always allow users to change the environment. Art pieces may be highly segmented and fantastic, while historical works and government reference material may be highly linear and realistic.

4.2

Vision

Color, light, shape, and depth are all things that are pre-consciously processed by the brain; that is, one requires no special knowledge to decipher their meaning - understanding is innate and immediate. Conscious processing, on the other hand, requires a sort of learned skill as well as conscious mental effort to understand. By making intelligent use of preconscious elements to convey information, one can increase the speed of understanding, the time it takes to recognize data. This section provides a brief overview of some of the preconscious elements that can be used in designing environments.

Color Elements



43: Ten-step and twelve-step color wheels, useful in picking color schemes. The ten-step wheel is focused on red, green, and blue as its main colors and is suited to on-screen design. The twelve-step wheel focuses on red, blue, and yellow and is more suited to printed or painted works. Note that color spaces defined in this manner are subjective and are not mathematically consistent or precise.

Color in the physical world can be seen as consisting of three major parts: hue, saturation (or chroma), and value. Hue is what most people refer to as color; it is what one sees in rainbows: different hues. Saturation refers to the intensity of a color; it is the difference between a light blue sky and a dark starry night; between new and faded jeans. As colors lose saturation, they gradually fade out. Value refers to how light or dark a color is; as colors get brighter they become more white, and as they get darker they become more black.

In addition to hue, saturation, and value, visibility (or acuity) and contrast are also essential to understanding color. The more saturated a hue is, the lighter its value is, or the greater its contrast with surrounding colors, the more visible it is. Brightness is often used to refer to value; in the correct sense brightness is a subjective, psychological measurement.

The eye can discriminate among 128 different hues and many thousands of colors. Our eyes are most sensitive to yellow, blue, and green colors that are equal in brightness and saturation. There are a number of peculiar facts relating to human vision, for instance, it is known that if green and red light is detected, only one passes into the eye. For this reason designers should avoid potentially disturbing combinations of these colors.

Color can be used in a variety of ways to indicate or highlight information: it can orient the user along particular lines of sight, and in many applications one may wish to emphasize a vertical or horizontal massing of color, making short things appear tall and thin things appear wide. Color can be used to group similar kinds of information and behaviors together, it can separate out layers of data or help achieve a sense of community. It can show recurring themes and point out patterns. Over time color may decay or change to indicate a change in status. Using color alone, transparent effects can be made to generate a sense of depth, and color boundaries can create a sense of space.

Color and form are inseparable and should be designed as one. Color enhances directs the mood of environments, it indicates changes. Users should be able to color elements in their environment to fit them and display information they feel is relevant.

In applying color to worlds, designers should not forget the color-blind. At some point, environments should be visually checked in black in white, or a black-and-white alternative should be provided.

Below are a number of color characteristics, subjective impressions of colors that have remained popular over time.

- Red - festivity, dignity, gravity, elegance, blood, danger, financial loss, revolution
- Cyan - sky, coolness, neutrality, contemplation
- Green - nature, affirmation, security, fertility, poison, corruption, environmental consciousness

- Magenta - detachment, royalty, melancholy
- Purple - royalty, sensuality
- Blue - sky, understanding, quiet, reliability, corporate strength
- Yellow - being happy, warmth, joy, sunshine, caution, gold, wealth, luxury
- Orange - being informal, being cheerful
- Brown - being informal
- White - purity
- Black - land, horizon, witchcraft, sophistication, authority, death, fear, mourning

In general, warm, bright colors like yellow and orange are associated with leisure and recreation; cool colors such as cyan and blue connote efficiency, work, and background information. Bright hues may help mental activity.

Black makes a good background, but if used as such a context is needed so the viewer does not feel "lost in space". Yellow when judiciously used can imply nobility; black and white is one of the most noticeable color combinations, due to their high contrast. Thus we see yellow and black school buses, caution signs, and taxis: they want to grab your attention. Yellow and black then infer being businesslike.

Red when used in large space can imply a sense of dignity; when used in fine detail it can be quite elegant; thus Chinese temples and red calligraphy denote importance. Blue is a serene color and is capable of making spaces appear larger than they really are. Red and blue in combinations are often used to denote hot and cold pipes.

It has been shown that green is the easiest of colors to live with constantly. This may be because we are accustomed to nature as a pleasant thing; but green can also imply rot or sourness. This is why sugar packets are not green but instead are white or pink, which imply purity or sweetness, respectively. Artificial sweeteners are often packaged in pink to imply that they are just as sweet as anything else.

Much imagery is associated with different colors and their meanings are often culturally particular. When stop lights were first installed in China, people would go when the light was red and stop when green. Why? Red is the Imperial color of China; there it means joy and festivity. Below are some associations that should be kept in mind when designing for an international audience:

- Japan - White is associated with death
- India - For Hindus, white is connected with mourning; in general, red implies prosperity

- Latin America - Purple is associated with death
- Moslem countries - Green is a holy color and should be used cautiously
- Ireland - Green and orange are colors of warring factions in Northern Ireland and should be used cautiously

Value and Contrast A number of basic color contrasts exist; some of them are:

- Cold-warm contrast - Contrast between cold and warm colors can help depict aging with time and increasing danger or activity.
- Complementary contrast - Colors that are opposite each other in the color wheel are complements. Pairs of these colors often conflict; in this case it is best to use the brighter, more active color sparingly against the other.
- Simultaneous contrast occurs with two colors that are separated by two or three positions on the color wheel. This kind of contrast should be avoided, as it emphasizes the border at which the colors meet.
- Analogous colors are close together on the color wheel and tend to produce harmonious combinations.


The sensors at the back of our eyes are generally grouped into rod and cone cells, named for the shape of their structure. Rods detect brightness, and each eye has 128 million of them. We have 8 million cones in each eye; these detect color and detail. Because each eye has far more rods than cones, this is why we detect changes in brightness easier than changes in color.

A brighter environment can help us discriminate among different hues better. Users should be able to change the contrast and brightness on their display devices in order to see things more clearly. Contrast is essential in design; it defines shapes and allows authors to emphasize particular elements in a scene.

In the more traditional architectural schools of drawing in the latter 19th century, particular guidelines for sketching shade and shadow were established. Sunlit areas in shapes received the most light - these were not fully white, but an off-white, usually a warm color. Shade values were applied to sides receiving reflected light, and shadow applied to the areas from which light was subtracted.

Typically equal steps of tint were chosen between sunlight, shade, and shadow areas, and the actual color of the lit object fell between the sunlight and shade values.

44: A traditional method of determining shade and shadow values. The actual color lies halfway between sunlight and shade; in this example it is a 20% tint.



This process works well for two-dimensional grayscale images; it may be useful to apply such knowledge to the construction of information environments.

In any image, people will tend to look at the brightest areas first. Do not overuse brightness for this purpose, however, or the user will not be able to pick out details. The problem of not enough or too much contrast is a common one in graphic design. As a test, one should squint at the work at different distances in order to determine that elements are properly balanced.

Developers should use logarithmic rather than linear scaling when using contrast to show values. Otherwise, differences may be difficult to pick out.

If the environment is meant to be viewed using a head-mounted display, a 3 to 1 contrast between background and foreground is recommended if they are using an LCD display. The contrast ratio on CRTs is typically much better, so contrast should be muted somewhat; fully white objects may appear glaring; documents with completely white background should instead be slightly gray or a muted, neutral color that contrasts well with foreground elements.

Because HMD displays usually encompass one's vision and are close to one's eyes, users will perceive color much more strongly - the number of colors, especially highly saturated colors, should be kept to a minimum.

Format Elements within layout should be placed so that one's eyes are guided naturally to the relevant information. Blinking is very good for attracting attention but can be quite annoying if it cannot be stopped, if it is overused.

When reading we process, in order:

- Visual patterns
- Grammatical constructions

- Sentences as a whole
- Word shape

We tend to look at the more general level of detail first. To create realistic three-dimensional environments, the attributes to pay attention to are, in order:

- Overall form (contrast) and color
- Behavior and movement
- Simple shapes
- Texture
- Complex shapes
- Shadow

Convincing texture is more important than overall shape, as textured bitmaps can emulate shadows, bumps, and various shapes without being polygonally complex. Proper behavior and movement is more suggestive than shape and is more important in making the environment appear real.

When encoding information, one should keep in mind the old guideline: seven plus or minus two. This is the number of elements that people are most likely to remember unaided. Here are some other guidelines for encoding data in different ways.

- Position coding is overall the best method, and it most effective in environments with low information density.
- Color coding is the second best method. A maximum of 11 levels is recommended. Note that coding only by value and texture is not reliable.
- Coding by angle can be as good as position coding but its overuse can be confusing. A maximum of 24 levels is recommended.
- For coding by shape, a maximum of 15 levels is recommended for geometric shapes, and a maximum of 30 levels is recommended for images.
- Size coding has greater impact than length coding. A maximum of 5 levels is recommended.
- Coding by length is good for comparing ratios of things. 4 to 6 levels of coding are recommended.

Space and Depth

There are a number of ways to indicate depth within an environment. Things that allow us to perceive depth in static images include interposition, or laying shapes on top of one another; shading, brightness (blurriness), size, perspective, and texture gradient. Texture gradient tells us how far away an object is by the detail of its texture. For instance, one may see individual grains and specks of sand near you on a beach, but the farther away one looks, the less resolved the detail becomes.

One cue that allows us to perceive depth in moving images is static parallax - things closer to you move faster than objects that are further away. One can see this being used in many video games and animation today. To achieve the effect, separate layers are moved at different speeds. As a character walks, mountains in the far background move the slowest, nearby buildings move faster, and the character moves the fastest.

When we focus on objects that are within 20 feet of us, both of our eyes converge on the object, each located at a different angle. We sense this difference and it is a physiological cue that helps us to determine whether an object is far away or not. When viewing objects, our brain processes and compares the images from each of our eyes; it is the difference between what the left and right eyes see that allows us to sense how far away things are.

Blurriness and changes in color are things that can be incorporated into environments in which realistic images are more important than displaying as much detail as possible. Normally objects become faded, desaturated, as they move farther away, and light colors darken as dark colors lighten. Things become blurry in the distance.

This is so because at the back of our eye our cone cells, the receptors responsible for color and detail, are concentrated mostly in the center. As we focus, the image coming into our eyes becomes focused there. But the amount of cone cells decreases outside of this location; this is why our peripheral vision is blurry.

Without things such as atmospheric effects, haze, and blurriness, environments that wish to look real tend to look contrived instead. Developers should make use of any number of depth cues so that users can have a good sense of distance when navigating worlds.

4.3 Input/Output Devices

There is some apprehension that the use of computers with its various display devices will replace paper. At least for now, the idea is not to replace it but to make more efficient use of what we have. Companies regularly stockpile warehouses of unused mailings, brochures and forms; the U.S. is incredibly wasteful of tree-based materials. When paper is eventually replaced by something, it is likely that we will have progressed to a point at which we will no longer miss it.

We like paper because we can cluster and store large amounts of it, it incorporates a high amount of resolution, we can easily stack it in a compact way and pick out individual pieces, we can visualize large amounts of discrete documents, and we can easily place, compare, and manipulate many papers at once.

If a display and environment were to displace paper, it would certainly have to satisfy the aforementioned needs:

- The display would have a very high resolution

- The display space would be large or encompassing
- We could easily position the display and/or its information in comfortable ways and so that we can see what is happening in the real world as well
- We can easily and quickly take notes, compare, find, edit, and read through large amounts of information
- The information or display can be given or shown to other people
- We can position information in the environment in such a way that a quick glance at it can tell us what information is the most pertinent
- The device is small and lightweight, being highly portable.
- The device incorporates some sort of static, persistent storage and backup mechanism

Computer Monitors

Today's display devices are far from displacing paper, but they are continually improving. We show our bias for paper in the programs we design; Web browsers are "pages"; we try to use page-like backgrounds; the display area of programs is often taller than it is wide. This is designing for paper.

We should instead design for the display, or design for display independence. Many CD-ROM-based media is created for the screen; their areas appear wider than taller, and many good examples exist of such applications making efficient and elegant use of the entire screen.

The idea is that media should be created in a size and format that takes advantage of its capabilities, that allows communication to flow easily, that allows authors to express their ideas. This can and should be done. If we are able to create scalable media, movies and images that can be changed in size without loss of resolution, filters that can convert media into different formats as it is being experienced, then the need for fixed sizes and formats goes away. This does not mean that the future will or should be full of mutating, shape-shifting collages. Intelligent defaults, suggestive hints, and author preferences should always be present to provide a baseline, a mental anchor for experience.

Today we design for the flat rectangular screen. Developers should realize that the things they create will be experiences; they may wrap around one's field of vision, or drape on furniture, or envelop buildings. Bruce Sterling has envisioned computers as handkerchiefs, with circuitry woven into the thread. The images that computers generate can be put on all manner of devices and things. There need be no barrier between computers and users. Screen implies veil; there is something we cannot see through. Lift the veil.

Reading is 25% slower when done from CRT-based computer screens. This can be attributed to poor choices of type and layout, bad resolution and image quality, and glare. It is also very difficult to position a large monitor in a comfortable reading spot. The luminosity of CRTs makes our eyes more sensitive to flicker. When used in head-mounted displays, this problem can be compounded, which is one of the reasons that LCD panels are the display of choice in HMDs - there is less glare, and more subtle contrasts. Using most LCDs today, however, it can be difficult to see text and fine detail in environments.

We must have better input devices as well as displays. There are devices that we manipulate, such as mice, and then there are input devices that monitor us automatically. Why shouldn't the computer have sensors to see us and listen to us to help us do what we want to do?

Pens Early on in the history of computer screens, light pens were a popular device; users could draw directly on the screen. Today we have pens primarily on art tablets and hand-held devices, but why aren't our computers monitors touch-sensitive? One can guess that is proved to be too tiring for users to write on a vertically-placed screen. Horizontal devices that can be placed flat or in comfortable writing positions often use pens for input, often because there is no space for a usable keyboard. But all of these devices should offer multiple kinds of input, so I can use a mouse/pen/keyboard combination with my computer. When using both hands for input, people are able to do more, more easily. The more inputs we can manipulate, the more senses we can convey. Why not two mice at once, new devices to take advantage of the fact that we have two hands?

Organizations that frequently give presentations often have whiteboards or butcher paper available to show ideas. Their size-to-price ratio is one advantage over current computer displays, but they point out the importance of being able to scribble down ideas quickly while using a minimum amount of effort. Environments should allow this.

Keyboards Chord keyboards, in which letters are produced by playing chords as on a piano, were developed decades ago and proved themselves to be a superior alternative to the conventional keyboard, much in the same way that the Dvorak key layout was superior to today's QWERTY configuration. Their lack of acceptance points out the fact that input devices should be immediately easy to use by beginners. It also points out that fact that sheer inertia sometimes wins out over positive progress, so it is important that news of new, efficient devices created today be widely disseminated.

Using a see-through HMD or within a 3D environment, it is possible to create virtual overlays for keyboards and other devices with buttons or labels. My physical keyboard may be blank, but using an overlay any of the keys may be customized with different col-

ors and symbols. The blank keyboard provides a way to provide tactile response, allowing users to quickly position and home their hands on the device.

Mice How can today's mouse devices be improved? Certainly they should be able to support at least three buttons (a la Douglas Englebart's first mouse); such mice have proved to be very convenient in navigating in 3D environments. Without the extra buttons, users often find themselves using awkward key combinations for functions that should be fluid and closely grouped. Perhaps a mouse can have interchangeable parts, so one can substitute three buttons for one, or sliding controls to facilitate 3D movement. Users may be able to turn it over and use it as a trackball.

The Spaceball is a sort of trackball that can be raised and lowered; generic forms of this device are becoming popular among home video game players, who use it in first-person 3D games. Hopefully a solution like the Spaceball will become a standard for 3D input much like the traditional mouse has.

There is a similar need to create 3D version of the artists' tablet, a device that allows people to sketch in three dimensions. Devices used in special effects work translate models into 3D representations by recording points along a grid that has been drawn on the model; the process is painstakingly slow. Work is being done on the creation of scanners that can analyze an object, create a wireframe model, and apply the appropriate texture. Devices that track motion and shape in three dimensions using lasers may one day be cheap and portable enough to commercialize.

VR/O Wearable input devices for three-dimensional environment systems should incorporate some sort of force feedback. If I grab for a virtual object with DataGloves, it is difficult to know when I have it in my grasp or even when I'm touching it. If I touch something, the object may blink, or I might hear a certain sound, but it is not enough. It doesn't really need to feel solid, rather, all I need is some hint of solidity. Even a little sensation can go a long way if the system is responsive enough. Today gloves are infused with inflatable air pockets or shape-memory alloys in order to provide feedback. Much work needs to be done to make portable, low-energy tactile devices like these.

Traditional VR with heavy HMDs and gloves is too physically demanding. We may point and gesture in the air, stand and walk around, and do any manner of things in immersive environments. The problem is that it is tiring to do this all day; this is why we sit at our computers and let our arms and hands rest on armrests, desks, keyboards, and mouse pads. Devices should not promote fatigue after extended use. Methods such as retinal and body tracking and sensor suits allow the user to move fairly freely in the virtual environment without having to move a great deal in the physical world.

Even though it is desirable to provide some sort of acceleration of movement in environments, disparity between user and virtual movements can cause Visually Induced Motion Sickness (VIMS), also called simulator sickness. The sickness, often accompanied by a feeling of nausea can arise when:

- One's virtual viewpoint is moving without the appropriate body cues
- One experiences frequent starts and stops and sudden acceleration
- The movement in the display is not matched with body movement

One can avoid inducing this kind of sickness by keeping parameters regarding size and movement close to those in the real world. If the user is expected to move around a great deal, acceleration should be gradual and some sort of autolevelling feature can be enabled.

Display mismatch can also be alleviated by increasing the frame refresh rate and keeping the response time fast. Users can be quite sensitive to lag time; those wearing HMDs can detect lags of 200 milliseconds. In terms of the frame refresh rate, 40 to 60 fps is recommended, and the higher the resolution, the better. Sickness tends to wear off after using the system for a while: users should be given a warming-up period to get used to things.

Vision

Our entire field of view encompasses 270 degrees, but to gain a proper sense of immersion, we only require a field of view equal to or greater than 90 degrees.

Reducing the field of view to 60 degrees horizontally only can help prevent sickness. Many HMD systems have a smaller vertical than horizontal field of view because people do not tend to use the vertical field that often - they don't look up and down that much. It is for this reason that 3D holography discards vertical parallax in its systems - most people don't notice it missing at all. As one views such holograms moving up and down, they look like static images, but they look appropriately three-dimensional when moving from side to side. Discarding vertical motion in this case saves a considerable amount of computing power.

At the Human Interface Technology Lab, great progress is being made on Virtual Retinal Display (VRD) devices. Using a coherent light source such as a laser, VRDs scan images from a video source directly into the eye. Currently full-color, 1,080 pixel by 1,024 line, 72 Hz images can be achieved. The device uses little power and provides a field of view of over 120 degrees.

To keep out ambient light and to prevent the device from slipping (as it needs to be precisely aimed into the eye), it is likely that any commercial device that makes use of VRDs will be form-fitting in some way, much like swimmers' goggles are.


Another invention that will influence the development of display devices is the Digital Micromirror Device, developed at Texas Instruments. The device is basically made up of extremely tiny digitally controlled mirrors etched in silicon; it is basically a high-quality color TV display on a microchip. A standard DMD offers better resolution than a 640 by 480 pixel computer screen on a surface that is only 5/8 of an inch wide. DMDs will be widely used in projection televisions; they should be on the shelves within a couple of years. It is likely that the television of the near future will mechanically resemble something like the CRT-and-Fresnel lens combinations found in the movie *Brazil*.

- Information on the Virtual Retinal Display can be found at <http://www.hitl.washington.edu/projects/vrd/>
- DMD information is at <http://www.ti.com/corp/docs/tech/dlphome.htm>

The Portable Environment

If we combine miniaturized VRDs with goggles, force feedback gloves, and a small computer, we have the beginnings of a portable VR set. It is likely that a keyboard will be a part of it; a blank keyboard can be added with tactile ridges.

Such a set may be a small, thin, briefcase-like object that folds out and locks flush into position so it can be rested on a lap and held in a fairly static position. The device powers up when opened; one side holds the CPU, display hardware, and other circuitry and the other, more hollow side holds the goggles and gloves. Both input devices can be connected into the front of the keyboard. Three head and hand tracking sensors are mounted flush around the keys.



45: This portable set could be folded up and carried like a briefcase or stuffed into a backpack. Input devices are stored in one half of the device and are connected into jacks in the side.

The keyboard, whose keys may be somewhat rubbery to prevent slippage, is a kind of device that Rudy Rucker has called “feely-blank”: Feely-blank objects have no exterior labelling, rather, the labelling and shape is created and shown in the virtual environment one sees.

Key ridges guide the hands into typing position and into position to use other buttons and controls. A number of blank shapes are available to customize for different functions. A trackpad or track area may be incorporated into the device; a track area would be a particular 3D space in which the user moves the hands solely in order to navigate. It may be possible to have ridge templates, different kinds of keyboard ridges, ridges appropriate for sketching, architecture, or any number of applications.

The device may be able to talk to similar devices in the area, setting up a LAN. In this way groups can be quickly formed for work or play.

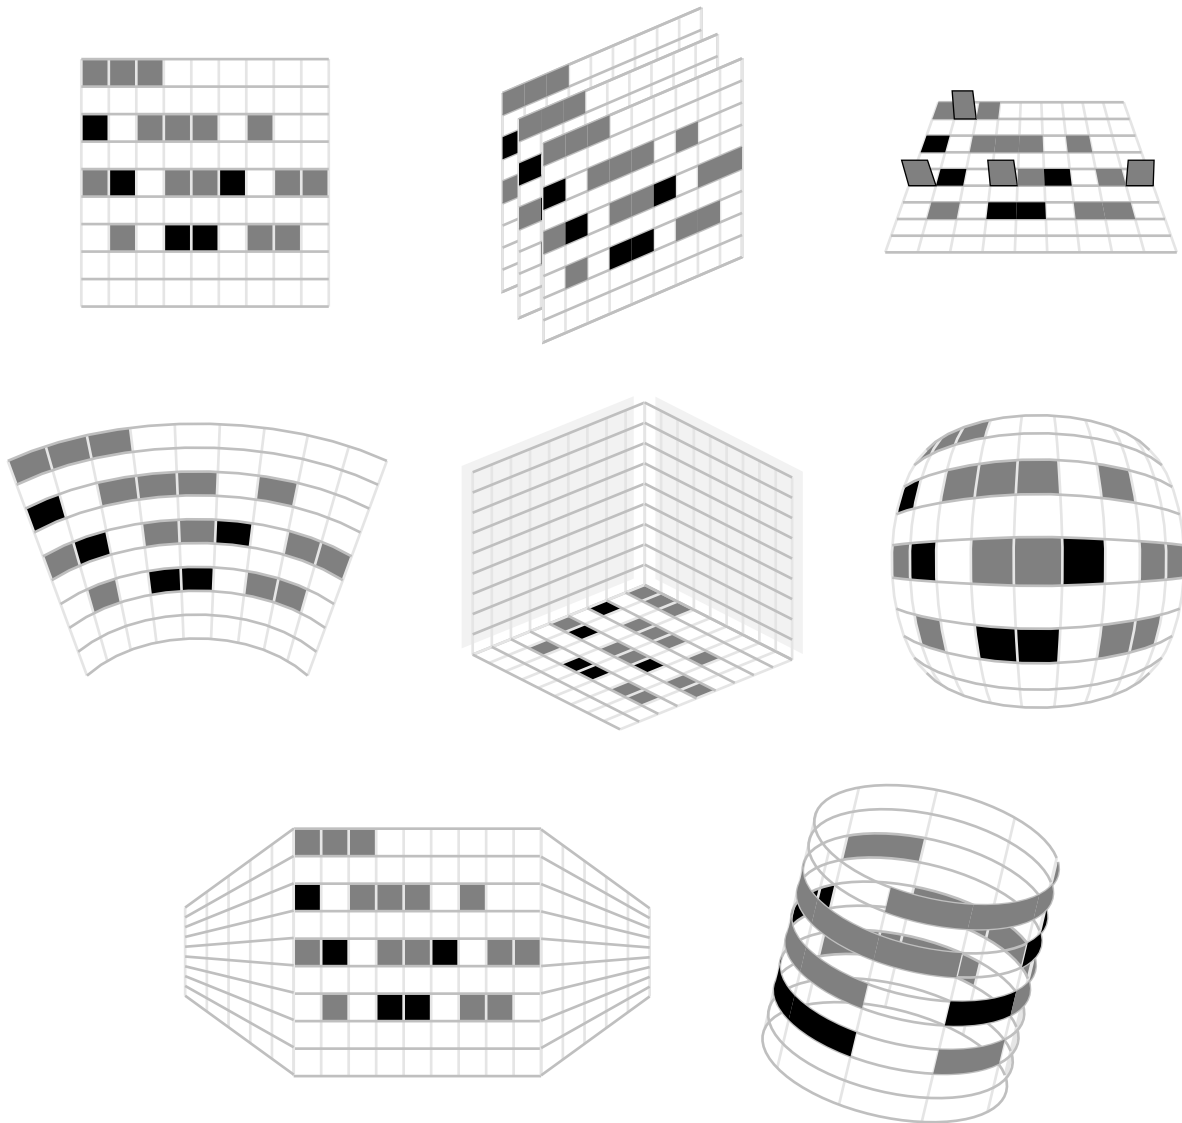
When the input devices we use are so close to the body (gloves, glasses, suits, etc.), there is some concern over health risks, particularly in public areas where the devices may be shared by many people, like handsets in public telephones. Because these things will be so portable, it is likely that any number of things will happen: devices will be personalized, used by only one person, and will be worn as fashion to some extent; and they may end up as interchangeable devices like headphones, kept around the house and in drawers when not used.

4.4 Possibilities

If the button is not shaped like the thought, the thought will end up shaped like the button.

Ted Nelson
Computer Lib

Layout There are a number of strategies one can follow in creating a new kind of information environment; in this case I will start with the basic layout of the system.



46: Common grid-based layouts in environments. From top to bottom, left to right: Vertical plane, matrix, horizontal plane, cockpit, cube, fisheye, perspective wall, cylinder.

Many environments are oriented towards a particular grid type as a basis for their layout; described are some of the more common ones.

Vertical plane - Programs such as the Macintosh Finder and other windows-centric environments make use of a vertical layout on which elements are placed. In order to see items in detail, the items themselves may “come closer” to the viewer; as in the Media Lab’s “Put That There” data space. Some systems do not allow items to move but rather allow the viewer to come closer to the entire environment, zooming in on the particular item of interest.

Matrix - Multiple planes are stacked; they may be semi-transparent, allowing users to see elements that are obscured by planes in front of them. The planes may be grouped and manipulated individually or as a whole and is somewhat similar to a 3D windowing environment.

Horizontal plane - This is one of the oldest layouts used in virtual environments. Elements are stored on a plane that can be flown over; the plane may recede into the distance, giving the impression of infinite space. David Gelernter has described a system that keeps track of a number of variables; they are visualized as square elements on a horizontal plane. Values or action items are indicated by the angle at which the squares are raised on the plane, easily allowing important information to be noted and compared. Multiple parallel planes may be used as levels to group objects and functions. Virtual worlds are typically built on a horizontal plane.

Cockpit - This layout wraps partially around a user, much in the same way that instrumentation wraps around the pilot in older airplane cockpits. All the users need do is turn their heads to view different areas or media. A completely immersive cockpit may take the form of a sphere that engulfs the viewer. Sun’s Starfire is a desktop display concept that appears as a cockpit.

Cube - This is commonly built along three intersecting planes - the layout may be a cube or corner. Elements may be stacked or grouped at the bottom or hung in space within the corner’s boundaries. Matrices representing databases may be created using this layout; widgets along the three axes can help pinpoint data.

Fisheye - A plane bulges or expands at the point of interest, showing the elements at that location in more detail. Xerox PARC has done some experimentation with this type of layout.

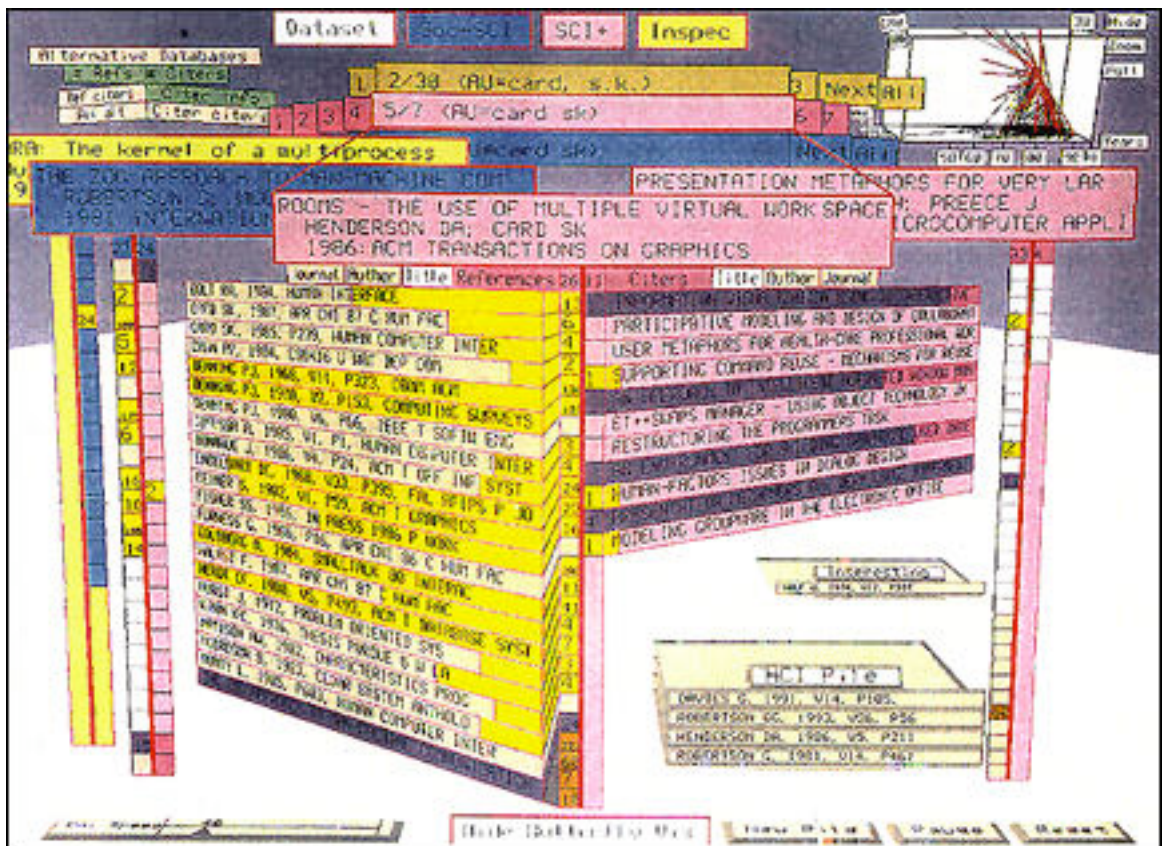
Perspective wall - The perspective wall layout was popularized at Xerox PARC. Elements may be sorted on different vertical levels on the wall; they can be scrolled on the plane horizontally like a paper scroll.

Cylindrical - Elements may be placed along the outside of a cylinder; users can manipulate it to find what they are looking for. A similar layout has been used in SGI’s interactive television prototypes.

All of these layouts may comprise smaller elements within a larger context. Each layout lends itself to particular ways of grouping, finding, and manipulating elements; the choice of layout should be determined by its use, content, and users. Here are two other layouts that have been popular:

Cone - A cone-based structure is suited to hierarchical applications; typically a directory is placed at the apex of the cone and its contents are distributed along the circumference of the cone's circle. To avoid obscuring elements, lines or shapes making up the cone shape are usually thin or transparent. This structure was another layout pioneered at Xerox.

Sphere - A sphere may have elements placed around its outside; it may be faceted so that separate elements appear on their own facet. The sphere may be rotated to examine elements and may make up nodes in a larger web-like or hierarchical structure.



47: A screen shot of Butterfly, an application developed at Xerox PARC for visualizing bibliographic citations.

Butterfly, an application with three-dimensional elements, was designed to display a great deal of textual information. Note how the spacial and color grouping highlights different data. The bright yellow items draw the most attention. Although cluttered, there are some interesting elements within the workspace. Piles allow

citations to be clustered together and possibly sorted; this kind of stacking of elements is an idea that is emulated in many kinds of systems. The problem with stacking many elements into piles, however, is that they may in time obscure other elements on the screen if they are tall enough. If it is possible to move stacks elsewhere in the space or represent whole stacks with a smaller pointer element, this may help. One may be able to “push down” or “pop up” stacks from the ground plane, or compress them in a particular way in order to conserve screen space.

Fonts When one sees the kind of work done at Xerox PARC, one can see how the quality of typeface can make a difference in legibility in 3D environments. There are a number of problems involved: How does one label elements? At what point do labels disappear, when the element or the font gets too small? In complex layout, it is possible for the computer to place labels automatically for maximum legibility, as is done on some maps, but this may be done so at the expense of cluttering up the screen.

Elements and their labels should not compete with each other. Rather, they should cooperate in an elegant fashion. The original Xerox Star, which made the windows-based environment popular, showed labels within icons, sometimes cutting off text and splitting words awkwardly. Systems can change font size, style, and type automatically according to the position and importance of elements, and this may be done in a useful way, but developers should be very careful to make labels appear consistent and not confusing. No more than one or two main typefaces should be used as the main system fonts.

Kids and old folks need big type, and environments must be able to accommodate this. Whether one wishes to choose a serif font (more common in the U.S.) or a sans serif font (more common in Europe) is mostly a subjective choice, although sans serif typefaces may work particularly well with technical, tabular matter.

In three dimensions things are a little different. Xerox uses a serif font in its environments perhaps to an advantage, because at an angle viewers can tell where the tops and bottoms of the letters are. Sans serif fonts at a horizontal angle, especially thin ones, may appear more like isolated lines. Sans serif fonts in 3D should most likely be thick; serif fonts should be readable enough to be recognized easily. Thickness along the third dimension should be used in typefaces only sparingly, unless the thickness is necessary to convey information in the environment.

At some point typefaces that are optimized for three-dimensional visualization should be created, much like Apple’s Espi Sans, which was designed to look highly legible on 2D bitmapped displays. Espi Sans is a comfortable font that suits the medium it was made for; it is one of the best fonts of its kind ever designed.

Text in all capitals is harder to process than lowercase text, because it is harder to pick out word shapes when the heights of all the letters are highly uniform. Systems should encourage the use of lowercase text.

Traditional computer environments have separated text from graphics, but why should they be displayed or treated any differently? They are simply images, and in the same way that movies and constructs can be dynamic, text should be responsive and dynamic as well. Today text and graphics are separated on the screen and handled in different ways on different layers. New media and future environments should not have to separate the two.

General Elements

There are a number of particular elements that can be created in environment that help the user visualize in useful ways or help them perform certain actions. Today's systems have made things like windows, scroll bars, menu bars, and palettes in virtual environments common knowledge. How can we improve upon these elements? How can things be easier to use and understand?

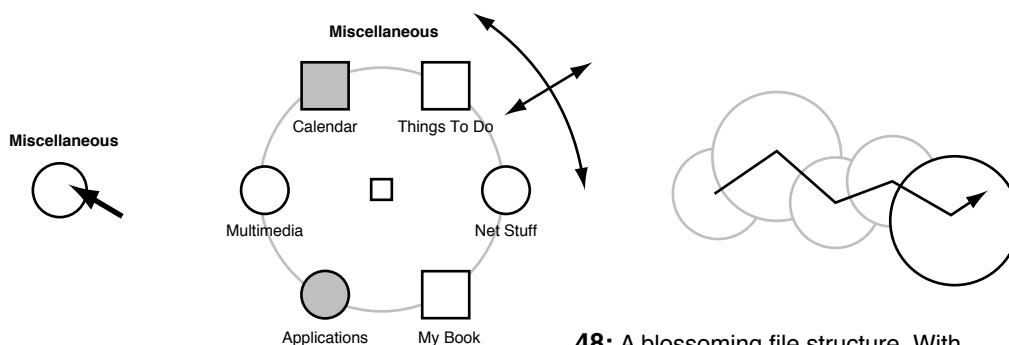
When our computers are so interconnected that services and resources are constantly coming in and out, there should be a way to check the status of this information. How many resources do I have left? What am I sending out? How long will it take for a certain task to finish?

One way of representing changing data is by representing values as pipes; the thickness of pipes can indicate how much data is flowing in and out of my system, whether a process has finished, or not. Programs have their own inputs and outputs; I may be able to monitor or even manipulate them using pipe objects. I may be able to chain pipes and tools together to create customized filters. If a pipe is farther away from us, it may represent an uninteresting task. Closer, thicker pipes or pipes whose color is highly saturated may represent urgent or important data.

If they are easy to read at a glance, dials to represent values can be better than using numbers by themselves. When we look at dials, we remember position, shape, and angle, things that are preconsciously understood by our brains. If numbers are involved, we must spend slightly more processing time in understanding the number and then associating it with a status. With dials, we associate form with condition. Dials should be thick and contrasted well with the background; often poor choices of shape and color are used, making dials hard to understand.

The interesting thing about computer environments of the last two decades or so is that we have been recreating physical widgets that have been around since the 1950's: the dial, the switch, the button. With computers, we can make entirely new devices and controls that may be more appropriate to the virtual medium. Many of the widgets are based on physical ones because when

we see them on the screen, we impart certain physics and behavior to them, thus helping us understand how they operate. But developers must be willing to explore new ways of doing things.



48: A blossoming file structure. With single, fluid movements, files can be singled out and located.


One example of a different structure is a structure that opens like a flower, showing its contents in equal steps along the circumference of a circle. The circle can be resized and rotated and elements can be sorted so that more important elements can be put in noticeable locations. This type of widget allows us to better remember things by their angle, as is done with dials. Using these elements, one can traverse structures in a hierarchy in a single, continuous motion to find a particular item. If accessed enough times, the gesture then becomes associated with the location of the file. This can help people locate elements more easily.

The drawbacks to such blossoming structures become evident when many elements are within a structure: the widget may become visually cluttered and then too difficult to use.

In visualizing interconnected webs of elements and hierarchies of elements, tree-like structures are often used, with lines between parent elements and their children or elements pointing from others. The problem with these kinds of structures is that after a certain amount of elements are in the system, the structure collapses into a spaghetti-like, unreadable mess, an image that I would liken to something I would draw with an Etch-a-Sketch on acid.

In order to make such structures more coherent, filtering can be used, selecting only particular data; one can use weighting, which may make particular elements more prominent according to an algorithm; or one might use clustering, in which items meeting some criteria are grouped together. It may help to see only the area of interest, rather than the entire structure at once; this would provide a sort of fisheye-lens view of the structure, empha-

sizing what is needed and downplaying what is not. Outline views and hierarchical menus are other structures that people have created to help navigate and perform certain functions these spaces.



49: Variations on similar structures can be created to serve different requirements. All of these structures are variations of the hierarchical menu.

Processing is faster when we can see the whole picture. Although it is important to see a general overview, users should be given the ability to see only certain elements and structures they feel are important.

Instead of trying to create structures that can visualize every element in the system, we should explore the creation of systems that do not require the visualization of so many elements: these are systems that have no concept of the “file” or “application” as we know them today.

Speed aside, one reason why some people like the old command-line based environments so much is because they do not have to visualize any elements in the system if they do not want to. To recall a file, they could type in a single command to use it; they need never create or view a single directory. Tomorrow’s system to some extent should allow both extremes. Command-line operation fails when the user forgets names of files and directory names; then they need an overview of the files on the system to remember. Graphic environments provide the ability to find files and remember where they are even if their names are forgotten, but they can become cluttered and confusing with too many elements.

One compromise is to provide a palette or some sort of input widget in the environment that is always available that allows the user to find things, whether these things are locally available or on the Internet, whether they are files, services, or topics. This search widget can spawn agents or filters to look for resources; it is the user’s hunting dog. The user then does not have to navigate through hierarchies or visualize hundreds of elements; instead they can simply type in a few related terms to get what they are looking for.

An environment that has no concept of files may instead keep all elements in a database. Users can construct media by mixing and manipulating pointers to elements in the database and saving these combinations with related terms and media to help

them find it later. Media and topics may be found by looking at particular subject areas instead. The creation and manipulation of media, rather than the storage and hierarchy of files, is emphasized.

An environment that has no concept of applications consists instead of tools that are able to cooperate and communicate with each other to do what the user wishes. This system does not make the user ask, "Which program should I use?" Rather, it makes the user ask, "What do I want to do?" Each tool in the environment would do one or two things only, but it would do them very well. Depending on the task, tools could configure themselves and group together. As users create media, they may select certain modes of working; the environment and tools around them change automatically depending on what the user wants to do. One may connect tools to do certain tasks and make pointers to the new constructions so they can be easily initiated.

One's cursor is a tool; it can be a pen, a gripper, or any number of things. Developers should take advantage of changing the cursor to indicate status or other information: in many applications it is the element on which the viewer's eyes are focused quite often.

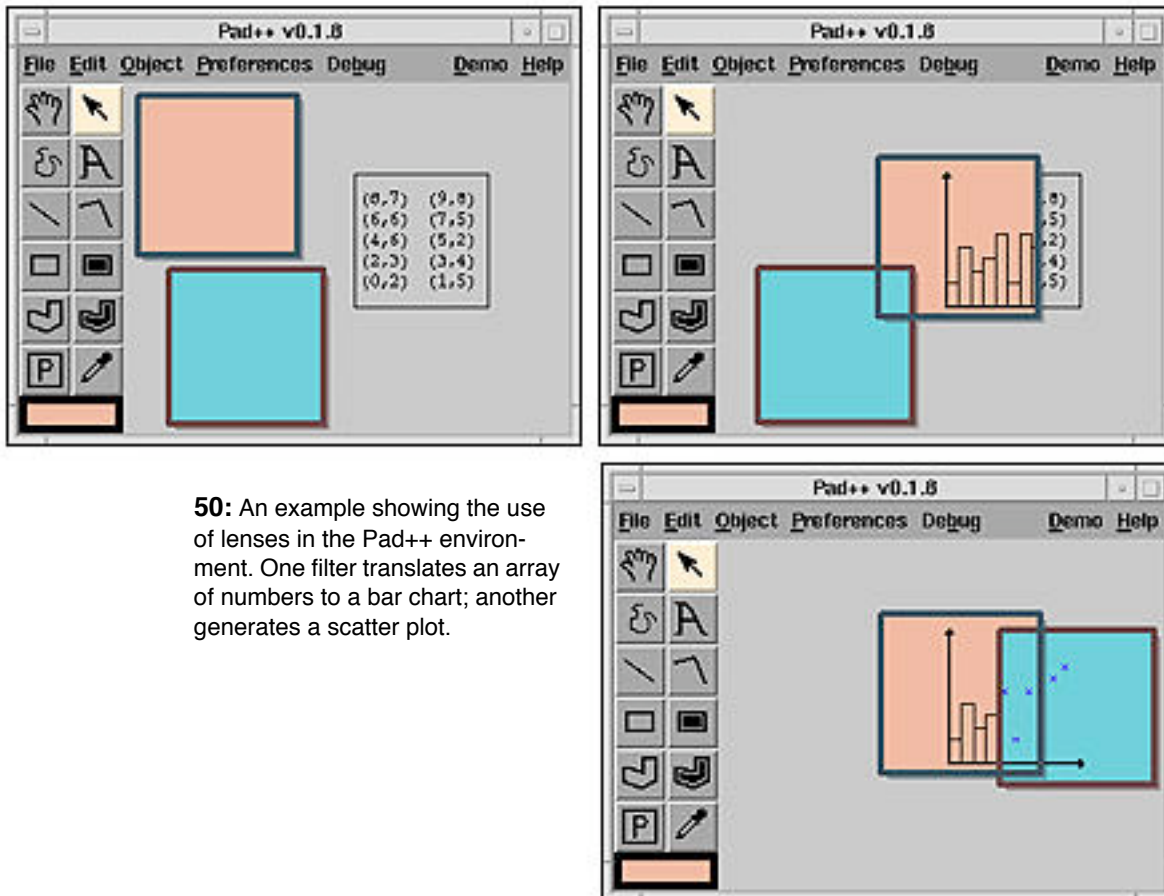
Many traditional windowing environments are able to store media in a buffer so it can be included in other works. This is often been implemented using a "cut and paste" method, which has become quite common. Today the new method is not cut and paste but drag and drop: media are directly manipulated, there is no invisible black hole that they sink into temporarily. In tomorrow's environment different methods may be created for doing all manner of things; it will not just make use of drag and drop but also surround and insert, merge and split, connect and isolate, enlarge and shrink.

When our computers are continually connected, always able to send and receive services, it will be useful to implement a pointer to services that are frequently used. One can think of storage devices and/or services as places; thus we may have something like a portal element in our environment.

Portals are like small holes, doorways, anything that implies space. If I mount a hard drive, a remote volume, or a CD-ROM, a new portal appears in my environment. I may be able to toss elements into the portal, sending things to whatever it represents, or I may enter it and explore and manipulate what's in there. Since portals allow two-way communications, I could represent output devices such as printers, but the element needs to show what it represents in some so as not to confuse the user.

With portals I may expand a virtual space, adding rooms to the environment. If I wish to send an agent to the net, I can send it through a portal. If I want to disable a service coming in and/or going out, I may shut the portal, or cross it out to remove it.

Developed at New York University's Media Research Lab, Pad++ is an environment that makes use of zooming - all elements are represented on a plane; users then zoom in and out of the plane to find things. Pad++ makes use of dynamic filters called lenses that can be moved over data to view it in different ways. When one has to use many filters on things, using lenses may become cumbersome, but it is an interesting idea.



50: An example showing the use of lenses in the Pad++ environment. One filter translates an array of numbers to a bar chart; another generates a scatter plot.


- Pad++ information can be found at <http://found.cs.nyu.edu/MRL/pad/pad.html>

Within new environments, developers should make use of direct manipulation as much as possible, rather than resort to pull-down menus or elements that require great movement in order to make a selection or perform an action. Instead of just clicking, one should be able to check, scribble, flick, and jot things in order to do things that today we use the keyboard or menus for. We directly do what we want to the object as much as possible.

As an example, one can take a traditional alignment feature found in many design programs: in many cases, the objects to be acted upon is selected and then an alignment menu option is chosen. Sometimes one must choose among multiple alignment

buttons in a toolbar or widget. How can one use less space, make a more precise tool, and put the action where the object is? Here is one idea.

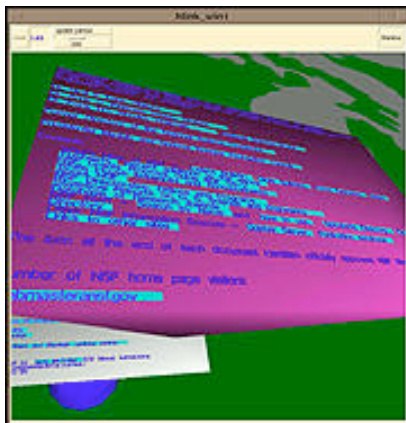
First, one may select an alignment tool or hit a modifier key to enable alignment. Then one can select a group of objects one by one, in which case a bounding box appears around them, shrinking to the maximum dimensions of the group. Alternatively a user may create a bounding box that is larger than the group of elements. This box has two axes on it that may be slid horizontally or vertically; when a user clicks once or twice on an axis or side of the box, all of the selected objects align along that axis.



Using this type of tool, users need never move far away from the objects being acted upon. In order to streamline the environment and make it visually more efficient, one should try to combine functions to some extent if it makes sense.

Navigation

Jakob Nielsen has written that users tend to visualize movement among pages as left to right, while linking to and from hypermedia is visualized as in and out movement. Can it be possible to take advantage of this view, to allow users to move in and out of documents as well as left and right, literally?



52: An HTML document as parsed in 3D in the Tecate Visualization System.

- The Tecate Visualization System allows navigation in this manner - users can open up new plateaus of documents, rising above or sinking below them. More information is at <http://www.sdsc.edu/SDSC/Research/Visualization/Tecate/tecate.html>
- Tecate is part of the Sequoia 2000 project, which hopes to advance earth sciences with the use of computer technology. Information is at http://www.sdsc.edu/0/Parts_Collabs/S2K/s2k_home.html

A great deal of research needs to be done in using trails as a navigation aid. As I wander throughout a dataspace, I may leave a history of media that I have visited. I may be able to save these histories from older sessions as outlines of sorts, and I might have the ability to splice them and merge them together, incorporate them into other media, or play them back like a slide show. How can trails be used to compare data?

Lara Catledge and James Pitkow recently completed a study in which Web browser usage was examined: using a modified version of NCSA's Mosaic for X, almost every user event was recorded: buttons selections, using hyperlinks, using other widgets in the browser's environment. About 43,000 events were recorded and then analyzed. The most commonly performed actions of all recorded events were:

- Selecting hyperlinks (this represented 52% of all recorded events)
- Selecting the "back" button (41%)
- Selecting an element in the hotlist (28%)

All other actions represented less than 2% of all recorded events. One conclusion that can be inferred is that *many of the elements in browsers today are unnecessary*, that is, they could be omitted or relegated to less prominent places and not be missed by the average user. Even back in 1987 in a study of videotex systems Nielsen found that people use the backtrack function frequently and that it was one of the most important navigational aids.

The Catledge study also pointed out the fact that people do indeed use hotlists, and that they can be quite useful if implemented properly. If users find themselves getting "lost in cyberspace", then that can likely be at least partly attributable to faulty browser design; navigation is not always easy to do. It is clear that developers need to do their usability homework in regards to developing hypermedia systems; without at least consulting past research, it is as if they are designing cars with the steering wheel attached to the floor.


Common actions in environments should be able to be taken by using the keyboard - Mosaic for X allowed users to hit the "b" key to go back and page up and down easily with single keys; a number of other functions were able to be done using keystrokes. Why should I have to move the mouse to a back button that may be located in an obscure area on the screen every time I wish to move back when I can hit a single key? Why should I have to deal with a messy and obscure hotlist interface? People use these functions the most; doing such things should be easy and not difficult.

It is clear that a better widget needs to be devised so that users can move backwards and see at a quick glance where they've been. As Alan Kay has said, what we need is the equivalent of page numbers on the Web.

People often use the back button in order to back out of a Web site: users may visit a home page first, then move deeper into the hypermedia hierarchy as written by the author. Often people move backwards to get back to a home page or main node of the site, backing out of trunks of the hypermedia tree they have been exploring. How can this be made easy?

One idea is to make use of a widget that succinctly represents the last few pages or media a user has seen; this may be done using a text-based list, but I would prefer the use of thumbnails, since they can take up less space (as opposed to a list with long titles) and can be visually more noticeable. This widget might display the last seven things seen by the user; allowing one to keep this short visual history memorable, although one could scroll to see more of the list.

53: This visual navigation aid allows users to quickly see where they've been.



As the user moves through the Web, new thumbnails are generated and added to the list. Outlines or color coding would group elements that are in the same Web site: if I visit a home page first, then link to another page under that, I will see two elements in my visual hotlist grouped by an outline or some sort. This outline allows me to quickly determine where the home page is, should I want to go directly back to it, and it also indicates how many links deep I am in the current site I am looking at. Catledge found that users accessed an average of 10 pages per server; if this list were 10 elements long it may very well take care of most navigation problems.

I may be able to drag thumbnails out of the widget into my environment in order to copy the media, or I may be able to group and rearrange thumbnails within the list. I can drag thumbnails to a bigger, textual hotlist, save it, compare it with other such lists, etc. Thumbnail navigation has been around at least since the advent of HyperCard, which made good use of the technique. Why not use it for the Web?

In addition to a list of thumbnails, other widgets may create small navigation trees that grow as the user browses; any node can be selected, taking the user back to a page or other media.

Through an analysis of repeated paths that users took while browsing, the Catledge paper also supported the existence of three kinds of users: those who browse for nothing in particular, those who browse for interesting information, and those who browse for something in particular. Another interesting report from the study was that people often use their home pages as jumping-off points to go to other areas. Users need to be able to create their own indices to things; this is a call for a better implementation of hotlists in all Internet-based software today.

For environments in which the user is expected to explore, the ability to see an overall map is very important. Such a map may be a simplified top view, a web, or an outline in a separate window by itself or it might be part of a heads-up display that can overlay other information. One can take advantage of the benefits

of fisheye views in maps and outlines; some systems have reported great success in emphasizing the data that the viewer is viewing and deemphasizing surrounding media. If I am exploring a particular room in a world, a top-view map can be generated that shows a blueprint of the world with the room I'm in blown up so I can see things in detail; surrounding rooms are scaled down yet still large enough so I can visualize them and select them.

In environments heads-up displays may be able to be used to much advantage; status panels may be slightly transparent so I can see underneath them; I might be able to change the panel from being opaque to being solid. Three-dimensional environment can make good use of transparency in information-dense environments.


Editing

When all media can be described, transferred, and experienced in a similar fashion, there is no reason that they cannot be edited in a similar fashion. The phrase "word processing" emphasizes the fact that one is manipulating words. If I work on images in my word processor, what am I processing then? If I make a Web site with images, sound, and video, am I processing words because I have to write in HTML? Or am I designing layout? If I am designing layout, how come my Web browser is so different from my desktop publishing program? If I am publishing data on the Web, how come my desktop publishing program can't do this?

The idea of having entirely separate applications to create and edit media is obsolete. It will not matter what the final form is of the things one makes; new environments will instead emphasize the creation, manipulation, and transfer of ideas.

In creating a consistent creation and editing environment, one possible technology to use is transclusion, invented by Ted Nelson some twenty years ago. In what Nelson calls the transclusion paradigm, any portion of a work may be reused in others, pointing to the original, and different contexts may be compared. New works can be made of transcluded media, and any part can be moved, edited, deleted, or subdivided. Ultimately, it does not matter what the media is - transclusion encourages comparing media, finding its original source, and merging media together.

54: Transclusion allows people to see the same media referenced in other works and compare contexts. This diagram shows two text-based documents; certain sentences have been transcluded from one document to the other. The double line represents transclusion: the material on both sides is the same. The single line represents a link: the material on both sides is different.




The main concept is to be able to find and edit information more easily. The World-Wide Web began as Tim Berners-Lee's personal information organizer. If environments cannot support doing these things in a straightforward, consistent fashion, then they are taking a step backwards. Tools should allow people to create as well as experience media; it is even better if there is no difference between the two acts.

Editing environments can also benefit by incorporating version control; such systems may appear as graphic trees, with different editions and revisions of media represented by nodes of the tree. Different portions may be color-coded or labelled to indicate who has certain editing permissions in a multiuser environment; other information may be shown as well, such as whether the media is public or private, if it's encrypted or not, or what file system it's on (although ideally, the idea of file systems may be completely transparent to the user).

As I edit text, I may mark up different portions with a sort of highlighter: Depending on the context, this highlighter may change plain text to a certain style as I pass the tool over it. By dragging in media from other places, I might create a hyperlink to them or include them in a document. I can strike out words with a quick stroke of the mouse, or scribble and type annotations in an available area. One set of annotations may have a certain color; I can strip out or add different colored layers of annotations, so I can see comments from other people or make notes on different topics. In Adobe Photoshop I can create and merge different layers of graphics; why shouldn't I be able to do this for any media?

While editing something, if I need a new tool to perform a certain function, my environment may be able to find one over the net, or I can shop around for one to do the job. Everything is modular: the tools I use, the media I create, the environment I do these things in. One can create tools and elements that merge with the media: much in the same way applications in Microsoft Office allow users to configure toolbars and rip off palettes, media in new environments can be merged and tools can be placed and configured as the user sees fit.



55: Merging applications and tools in new environments allows for maximum flexibility. Instead of the method of *cut and paste*, the new way of editing makes use of *merge and divide*.

If I wish to add a new piece of media to a document made with a layout-based program, I just merge the two. My environment can supply hints as to where it thinks the best placement may be for the new media, and may automatically apply certain styles based on my preferences or styles from other media that I've seen elsewhere online. I can embed my tools within media: if I send my creation to somebody else, I can include custom tools for manipulating or experiencing the media that they may not have. My tools can then be integrated into their environment if they wish.

- Sun's HotJava, an extensible Web browser, makes use of Java, a transportable programming language. HotJava can gain functionality by requesting Java tools over the net. More information is at <http://java.sun.com/>.

Participation

When our environments can be online all the time, there is always the possibility of doing things with other people or programs. The important thing is to be able to include people and programs in shared activities as well as prevent others from coming in: I don't want people intruding into my space all the time.

Collaborative systems at NTT allowed users to see each other behind or beside their windowing systems: they found that it is faster to point to something with one's finger than use a mouse. If the other member(s) can see your gestures, this is a much richer communications channel than a simple whiteboard. Almost thirty years ago, Douglas Englebart's systems overlaid data on top of video. Why can't something like this be implemented now in the common environment? I may be able to push elements in the environment aside to see somebody else, or their video may be shown within another piece of media or an element such as a window. What they see and what I see would be equivalent: my image would have to be mirrored for them to see things they way I do, and vice-versa.

If I wish to take video calls as I'm working, I may be able to set certain options: I can have my system perform a certain action when somebody asks to see me, such as play a sound or show an image. I may want subtle, silent indications or very noticeable ones. I should be able to let people in automatically, the equivalent of letting people wander into your house or office, or I may turn on an answering service that filters calls or give users options to send email or communicate in other ways. If my camera is always on, and I let people wander in, there should some indication of when somebody is looking and who they are. I might be able to use digital shutters on particular elements in my environment, so only particular people can see them. If I am monitoring a camera that's somewhere else, I may be able to ask my environment to let me know when it detects motion. This can be a basis for digitally-based security systems.

Much personal communication consists of "drop-in", casual contacts - this helps greatly to foster the sense of community. These systems must facilitate "just dropping by", but they must also

allow access privileges, so that the entire world does not simply “drop by”. I can make a group of friends I allow to drop by any-time, a group of acquaintances I allow to see me at any time, and others I allow to only leave messages during business hours. In communications, peripheral awareness is very important: I may be able to hear and see people doing things in the distance.

The typical three kinds of users have been found in the collaborative environment as well: there are wanderers looking for nothing in particular, people looking for something interesting to happen, and people looking for something in particular. In most cases, people looking for something are usually looking for other people: “Have you seen Bill?”. Should videoconferencing become popular, this points out a need for connections to somebody at an information desk or a connection to other media such as a Web site that can provide this kind of information.

If I am going to participate in a prescheduled videoconference, my environment should cooperate with the other participant’s systems: clocks counting down the time until the meeting may be displayed, so members can prepare themselves. The systems can negotiate media formats and methods of sharing and notify people of potential conflicts in services, available bandwidth, and so on.

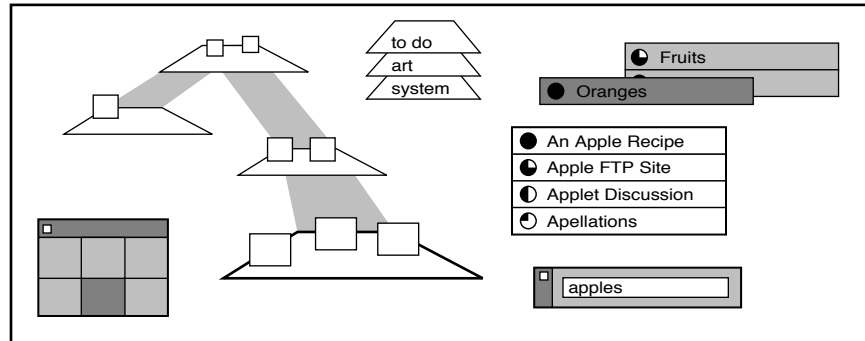
Portable cameras should have a wide, adjustable field of view and be small and unobtrusive; opening and closing the shutter may activate certain services, allowing people to drop by. If there is no shutter, a light or indicator in my environment should indicate operation. The more area the camera can capture, the less I have to sit and stay in one place. If I can walk around and make gestures with my hands, perhaps letting the camera automatically track me, better communications can be achieved.

Very high-quality video communications can be bandwidth intensive: Bryan Lyles from Xerox PARC estimated in 1993 that to support full two-way traffic, a transfer rate of at least 100Mb per second is required - this is based on the estimate that a single high-quality video stream would require 6 to 30Mb per second. For a system with 100 members that is continually online, about 10Gb per second would be needed. Given that the ATM rate in U.S. is 45Mb per second today, going digital with studio-quality services is going to take fast equipment. Until such rates can be sustained, many environments will have to make do with highly compressed video of smaller size and resolution. In most cases, the quality and size of digital video within a few years will be more than adequate for the average person for all manner of applications. Today slow but adequate audio and video communications can be achieved over the Internet’s MBONE at a combined rate of 192Kb per second; 112 to 192Kbs is what I would call the low end of today.


Putting It All Together

One can imagine any number of environments that visualize information in different ways; one example uses multiple plateaus to visualize file spaces and other media. Such plateaus may be stacked and manipulated in three dimensions. A palette may provide an overhead and/or side view for easier access.

56: One example of a future environment.



In the above example, a search palette allows users to find media locally or across the network. Results are returned as pointers to the object: they can be manipulated, stacked, and integrated into other media. These pointers may change color with time or may indicate other information.




57: How is information space to be visualized? Should it be a space?

One's environment may be represented as a world of expandable rooms; objects representing portals can transfer media in and out of the environment. If the concept of rooms is used, users should be able to leave the room, rise above, and see any number of views of the world: the claustrophobic, "box-in-a-box" effect should be avoided. The main idea behind rooms is that they can act as landmarks and provide common places for holding particular media.

If objects can be stacked, how can users determine their contents easily without examining the data within? Should objects within worlds have shadows on all three axes to make visualizing their position easier? Should the concept of windows be extended to

three dimensions, allowing for expandable boxes of media, media panes that can roll up, or is the concept of such structures useless or unnecessary in three dimensional environments?



58: Today's window-based environment, as represented in three dimensions.

Research needs to be done to determine whether any of these particular kinds of environments are useful. Today we are again in the age of carving in stone. We must ask ourselves what we really want, what tools are appropriate, whether these systems should be replaced by an entirely new paradigm altogether. The computer is possibly the most flexible device humans have been able to create thus far. What do we want to do with it?

4.5 Designing Worlds

Worlds in this sense are basically virtual environments developed with an emphasis on exploration and collaboration. While one's private environment (such as a window-based system) may be more suited to organizing, creating, and visualizing media as well as acting as an operating system, worlds are more public. Worlds can be created for people to meet, play, and hang out. Private environments are extensions of the brain and embody one's thought processes; worlds embody the architecture of the earth and can be extensions of many people.

In order to create worlds that lend themselves to exploration and finding things, it makes sense to borrow the best ideas from the real world. We can think of worlds, as well as Web sites, as airports. People from all over the world are flying in from other places. Many have particular destinations and need to get somewhere in a hurry; others are wandering around, biding their time and looking for interesting things to do. Not everyone speaks the same language. How does one design for this environment?

When you enter a new airport, how do you know how to get to the luggage claim? Are there clear, bright signs, symbols and arrows pointing the way at regular intervals? If you need to go somewhere in particular, such as the restroom, are there maps and overviews easily available to help you? Is there an information

desk? Large airports with high ceilings and artwork take on a cosmopolitan feel. Small, poorly lit airports seem cramped and neglected.


What spaces in the real world lend themselves to education? For starters, we can look at museums, children's museums. Signs and images are spread all around in lit, easy to read areas. The shape of well-designed exhibits is such that it encourages wandering along a certain path. Gadgets, objects, and simulations encourage people to interact. Brightly colored displays foster mental activity.

Newspapers like USA Today are grouped in color-coded, well-defined sections. Images and articles are laid out in an easy to read fashion. Summaries are available so that people can get a quick overview of daily events.

Good world design can take advantage of hypermedia as well as real-world techniques to make environments compelling. This does not necessarily mean that worlds should look like airports or be reduced to what might be considered shallow entertainment; rather, they should have the potential to be simple and primitive as well as complex and sophisticated.

The design of virtual environments in some ways is easier than real-world architecture: there need be no restrooms, kitchens, plumbing, electrical wires, or insulation. The need for enclosed, grouped spaces will likely come from the need for privacy and the need to group similar information together to make visualization easier. A sort of secluded public area may be available in order to change one's avatar or converse privately with others in a populated world. A building may represent a file space through which I can navigate.

Style Below are a number of the more popular basic styles of virtual environments. Note that such environments do not need to have a ground, or even boundaries as we know them; our spaces may be caves, impossible objects, or membranes.



59: Common styles of virtual environments. From top to bottom, left to right: Immersive, floating spaces, matrix, real landscape, fantastic landscape.

Some of the more basic styles of environments include:

Immersive - Elements within the world may surround the user in a circle, semicircle, sphere, or cockpit-like structure. The advantage of having objects placed like this is that a minimum amount of movement has to be done on the viewer's part to see information: one simply has to turn one's head. Using such environments is easier if head movements are constrained or "snap-to" areas are provided so that seeing different planes of information does not become a navigational chore.

Floating spaces - Worlds without a ground may float in space; colonies of such environments may be created. Worlds in which the initial viewpoint is not grounded or in which there is no main defined horizon on which structures rest may be considered as floating. Structures may be built on simulated floating chunks of earth or giant meteors; the environment tries to convince the user that it is in mid-space.

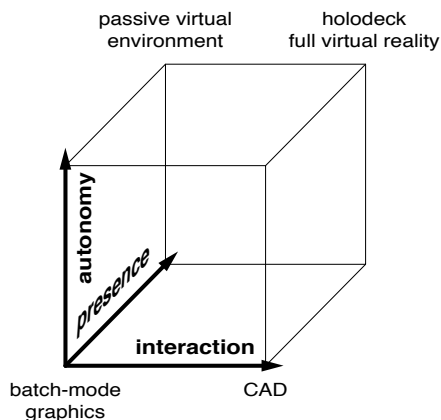
Matrix - This is one of the more common views of cspace; objects may be distributed in regular intervals and the viewer is often positioned within the resulting 3D grid. The problem with such worlds is that they may be too uniform to make navigation easy, and that they may be too complex or crowded to make sense of individual elements.

Real landscape - Worlds in this style may be similar to real-world environments: they may be faithful reproductions of cities or worlds with similar behaviors, objects, and proportions to the physical world. Being a landscape, this kind of world would be oriented along a horizontal plane.

Fantastic landscape - Worlds in this style are not likely to exist in the real world: surreal structures, strange behaviors, and other-worldly proportions may be prevalent. Buildings may be built out of giant primitives, creating shapes that would collapse if made under real gravity.

Look and Feel

In addition to the overall look of worlds, there is the issue of how they feel. Although the issue is ultimately somewhat subjective, David Zeltzer at MIT has devised a way to characterize the reality of virtual environments, using a cube that measures three qualities, each with a value ranging from zero to one:



60: The Zeltzer Cube.

- Autonomy, the objects' ability to react to events and stimuli
- Interaction, the degree of access to the variables in objects
- Presence, the fidelity of sensory input and output

The "realness" of an environment is measured as a three-dimensional coordinate: autonomy, interaction, presence. An environment at (1,1,1) is as close to reality as possible: things are highly reactive to stimuli; they are highly interactive; they appear real to all of one's senses. An environment at (1,0,1) may appear reactive and real, but cannot be interacted with. This would be a passive virtual environment, perhaps like a three-dimensional equivalent of TV.

Brenda Laurel suggests that one should focus on designing the action first; all else is then subsidiary. By designing the action, she means designing and influencing what incidents will occur and in what order. One can make paths, nice views, good things to look at, places to wander, grand perspectives, hidden objects, and peaceful spots. What is the purpose of the world? What do you want people to do? Where do you want them to go?

When considering behaviors, it is recommended that objects in the environment must move realistically even at the expense of other attributes. Use consistent rules of behavior; make the world appear real in its own right. We enjoy it when we engage in fantasy; the designer should let us believe in the world. Make representations that feel right and that act consistently within the universe.

Because using today's devices we can only visualize worlds, what we see must make up for what our other senses cannot detect. Thus designers should create textures that evoke a particular feeling, images that elicit smell as well as sounds. These things will greatly contribute to a feeling of immersion.

Building the World

When users first enter an immersive environment, among the first things they do are explore one's virtual hands and body if they exist, position themselves next to large objects, look for dramatic perspectives, and look within hollow objects (if there is no collision detection). Worlds should incorporate large, open, visible landmarks; they will help orient users and serve as natural spots for people to talk and hang out.

One can begin worldmaking by starting with sketches and rough, low quality drafts. Do different areas serve different purposes? Are there landmarks, places for people to wander, things to see? Make sure the initial viewpoint in the world specifies a reasonable perspective, point of view, focal length, and distance from the ground (if there is one). Some worlds may have users start from mid-air. If doing so makes sense in the context of the world, this is fine, but it may be disorienting when moving from world to world. It is better to start on a ground or in an entrance structure. If many avatars are within a world, one may appear in a calculated position so as to not bump into someone else as one enters.


What does the overall design of the world look like? Does it link to a number of subworlds, or is it self-contained? Are structures consistent in scale? When we view objects we mentally increase or decrease the scale of things in order to achieve a balance. If objects are wildly different, this may create some degree of visual stress. If we are viewing such unbalanced worlds on a traditional monitor, we may not notice this, but those using immersive HMDs may certainly feel disoriented.

Here are some basic guidelines to keep in mind when designing worlds:

Avoid low ceilings - If ceilings are very low, or the initial viewpoint is too far off the ground, the world will appear cramped, and the viewer may get claustrophobic to some degree. Open structures up; provide windows, sunlights, and arches.

Make large portals - Some users will not have the benefit of precise navigation. If doorways are too small and narrow, one may easily miss them. Collision detection should help alleviate this problem somewhat by allow the user to slip along surfaces they collide into to help them get where they are going.

Allow users to navigate with harpoons - Users with slow machines may wish to use harpoons as a navigation aid to move towards objects quickly. Providing large, flat surfaces for people to harpoon onto can help.



61: The Müller-Lyer arrow illusion and the Titchner circle illusion. Both horizontal lines are the same width; both circles in the center of the diagrams are the same size. Our brains try to keep objects at a constant scale.

Allow for variety - Mix buildings and structures; design different areas of the world with certain colors and styles. Build multiple paths and viewpoints to follow and experience. Make virtual menus of links to take users elsewhere in the world or somewhere else entirely; make virtual galleries of media of all types.

Provide common centers - Park-like areas with large vistas and good views, large landmarks and comfortable-looking rooms can serve as places for people to meet and collaborate. Make arenas, town centers, and amphitheatres.

Be natural - Simulate sunlight, sunsets, and stars; make grass, trees, earth, plants, and flowers. Greenery and sunlight loosens people up and makes worlds more appealing and easier to work in, despite the fact that they are computer generated. We fall for fake plants in restaurants; why should virtual environments be any different? Sunlight can help indicate the passage of time in the real world.

Complete the illusion - Make worlds self-contained. Do not make them open-ended so users can “fall off” a side of the world; enclose it in a texture map or walls or borders of some type, implement an endless horizon, or make a floating world.

Avoid dead ends - Design worlds so that users can navigate through them without having to double back a lot. Create main streets to handle heavy traffic; make all pathways point to common centers. People are the blood of worlds; make arteries they can flow freely in.

Of course, if one is designing worlds for dogfighting or hunting others, then all manner of elements can be introduced into worlds to make them challenging: lots of hallways, doorways, tunnels, trapdoors, and stairs; unexpected turns; extremely large rooms for flying; secluded areas for hiding; obstacles and mazes to maneuver.

For some reason, many worlds make little use of wide open spaces. Why not? One can have gigantic lawns, plains, and as much three-dimensional space as they like. Such spaces do not have to be monotonous; they can be used to simulate desert, oceans, and grassy plains. But wide spaces only make sense if they are not so disorienting and if it is possible to traverse them quickly. Crossing the virtual Sahara should not take as long as crossing the real thing.

Open spaces can be used for dramatic effect, particularly when large landmarks exist in them. There should be enough variety in worlds so that they are easily found. Along the same lines, developers should not link every single thing in worlds: Jakob Nielsen has said that if everything is highlighted, then nothing has prominence. Be extreme in design only when it makes sense, otherwise do everything in moderation.

Today to increase rendering speed it is best to design worlds using a minimum amount of polygons. Make use of built-in primitives. Large flat walls and thin struts can convey a lot of space and detail without taking using up a lot of polygons.

Instead of letting worlds float in a boring black background, color the background instead to fit the mood of the environment. At least apply a texture, a ground color, or something to indicate immersion within a world. Provide a plane or plateau for the viewer to anchor onto.

We must provide closure, don't just terminate. The audience, the users, are in a reality you expect them to believe in. Don't just punt them out and boom, have a new world appear. Provide some transport, fly to worlds, fade in and out, or zoom in from far away. Start in some default room, and empty space. If we are truly going to a world, then make it act so.

If users wish for privacy in a world, their wishes should be respected. Collision detection may be turned on in a certain area or for the entire world; browsers should certainly implement it. Detection may prevent users from accidentally disrupting private gatherings; there should be some way of indicating security in a world, and this may be done with things such as doorknobs and security pads on objects.

Examples In June 1995, the movie *Johnny Mnemonic* was released and was the first major motion picture to portray a vision of cyberspace on the Internet. Its sequences were done on IBM-compatible computers and were the first full-frame, 35mm effects to be completed using such hardware.



62: The Internet circa 2021 as portrayed in *Johnny Mnemonic*.


The movie's cyberspace is predominantly blue: the color evokes a mood of contemplation, of expanding space, and is very appropriate for large environments. But despite the color it is very messy: How are people expected to find their way? If you could design an ideal environment in which navigation and finding things were easy, would you make it look like this? It is a flashy, beautiful world, but it is perhaps better suited for entertainment rather than for practical purposes.

The movie introduces a compelling idea: the use of hand gestures or object manipulation to gain access to information. In one sequence, the main character must properly manipulate a segmented pyramidal shape in order to gain access to a certain file: He twists the segments in a certain order as one manipulates a Pyraminx (a triangular version of the Rubik's Cube, popular in the 1980's) until the proper sequence is performed.

In the real world, we perform complicated sequences in order to do things without thinking about them: We drive cars, cook meals, and operate our computers. Muscle memory allows guitar players to play sequences of notes faster than the conscious brain can process. Instead of remembering strings of alphanumeric codes, wouldn't we tend to remember actions better?

Those in the traditional VR field have found that by themselves hand gestures can be difficult to remember and are not always intuitive. But in the case I mention, the gestures are acting on an object; the user doesn't need to remember gestures, only what to do to the object. One can imagine gaining access to one's Internet account by properly manipulating a virtual Rubik's Cube with colored keypads on its various segments.

The manipulation of objects in a particular sequence is essentially ritual. If the manipulation is guarded and the goal secret, then using such gestural systems becomes sacred, in the sense that the ritual of entering one's PIN number in an automatic teller is sacred. But in cyberspace, movements can represent access codes, encrypted information: ritual can represent knowledge, and not necessarily physical goods. A discussion of what this implies for future society is beyond the scope of this work.




63: Two of Clay Graham's worlds: a "Neo-Retro Inner Thing" and "Stock World".

Clayton Graham at Silicon Graphics designed two worlds early in the history of VRML. One world, called "Stock" or "Stock World" was designed as a temple to Silicon Graphics stock: a user could see a continually updated graph of the company's stock in the middle of the world. In April 1995 at the Third International World-Wide Web Conference, Stock World became the first VRML world to be officially displayed by Silicon Graphics. Perhaps it is no accident that the first world shown to the public was a temple to money.

Both of Graham's worlds are exceptional in that they both convey a particular mood and lend themselves to exploration. The polygon count is small, allowing the worlds to be rendered and transferred quickly. Thin borders serve as world boundaries and promote a sense of place.

How is one to design a corporate world? One such world was originally designed in November 1994 for demonstration walk-through purposes at Enterprise Integration Technologies.




64: Above: A view of a hallway at EIT. Textures were taken from digital photographs of the original surfaces. Upper right: The entire world rendered in wireframe view. Right: The entire rendered world without the roof.



EIT's corporate offices, the entire fourth floor of a building in Menlo Park, California, were modeled. One can imagine corporations providing worlds for general use; interested people could take a tour, read product literature in the lobby, and speak with

available developers. If private areas are maintained, the world may serve as a way to get in touch with employees or provide customer support.



65: EIT's collaboration space. Top left: Outside a meeting room. A panel outside the door provides information about the status of the meeting. Top right: Interior view. Bottom left: A bulletin board for archiving and posting media. Bottom right: The virtual whiteboard application.

EIT's prototype collaboration world was designed in February 1995 to illustrate a possible use of cyberspace. This world may consist of any number of meeting rooms; in front of each room is a status panel indicating whether or not a meeting is in progress, what kind of meeting it is (if it is private or open to the public), how many members are inside, and how long it has been going on. If I have the proper permissions, hyperlinks in the panel can take me

to media describing the members as well as a current, ongoing transcript of the conversation, so before I go in I can be informed of what's happening. The entryway also has a small "glass" panel that may be transparent so I can quickly glance into rooms and see what's going on: this facilitates the idea of just "dropping by" to talk.

Inside the room is a table or some space in which objects can be shared; an entertainment area with links to games or magazines should people wish to drop out of the conversation, a bookshelf holding reference tools or needed material, a virtual bulletin board, and a virtual whiteboard.

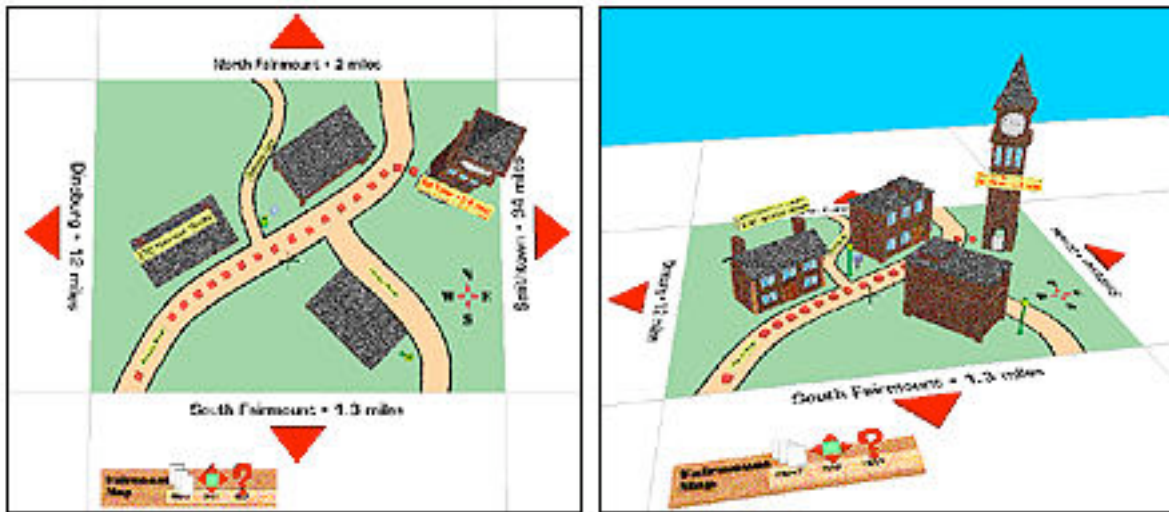
Both boards represent shared applications: the bulletin board is meant for archived media and media that is posted by others, and the whiteboard is for manipulating and sharing ideas and media in real-time.

The white board has a status panel that displays the members' names and shows what color pen they are using to draw on the board. The names may be links to personal information; I can also see what address members are coming from, the time and date, and other information. The board is exceedingly simple to use: I select a pen and draw or I select an eraser and erase. I may be able to selectively erase certain colors or all colors.

The bulletin board also has a panel through which I can sort its containers or view the media in different ways, sorted by different criteria. I can view posted documents, which may have been typed in, or I may see notes written with a pen-based device: a user with a Newton may have sent a message to this room.

Both boards are media-savvy: streams of audio and video as well as documents and images may be ripped off the whiteboard and placed on the bulletin board, and vice versa. I may be able to

drag media from my environment and place it on either board. Members can show up as avatars in the room or video streams on the whiteboard.



66: The Fairmount universe. This virtual town consists of three buildings and a clock tower. The top view of the town can be used as a map.

The Fairmount universe was developed at the same time to serve as a demonstration of educational and commercial worlds. Fairmount is a fictional town; the world may be an interface to a GIS application or may be a virtual world that has been built in a collaborative manner. As a GIS application, I may ask how to get somewhere, such as the Bell Clock Tower in Fairmount, and a path showing me how to get to the tower may appear. I can follow the path to visualize how to get there and can display or print out a list of directions to take with me in my car (or the application may be already built into my car).

The area is expandable; it is displayed as a rectangular patch with indications as to how far away neighboring towns are. I can select a red arrow and the map moves in that direction, showing the next patch of land. I may be able to change the scale of the area in order to view larger or smaller areas.

Note that if the world is seen from the top view, it becomes a two-dimensional map that can be displayed or printed elsewhere. It is important in worlds to be able to see paths, to easily get a bird's-eye view: one can easily zoom in on what one is trying to find. Labels are yellow rectangular objects that are labelled on all sides to allow for easy viewing from different viewpoints. The place I am trying to get to has a label with red text, others are in black.


An icon bar allows me to print the world or save relevant information to other formats or services. A miniature map control with four small active arrows allows me to move around the area or

easily see the top view. Selecting a help icon can provide instruction on how to use things as well as better descriptions of what the things are that I am looking at.

The main Fairmount world links to three different subworlds, two buildings and a tower: Harrison House is a historical bed-and-breakfast locale, Bell Tower is the town's clock, and Gray's Chocolates is a local candy store.



67: Above: Harrison House, a historical exhibit. Upper right: The clock in Bell Tower can be examined and run as a simulation. Right: Gray's Chocolates, where people can order goods and browse the company catalog.



Both the Harrison House world and the Bell Tower world are designed more as museum exhibits: people can select links to different media explaining the history behind the site, or they may take an automated, possibly narrated, walkthrough and tour through the worlds.

The Bell Tower world contains a simulation of the clock mechanism: people can hit a “play” or “stop” button and explore how things work; the mechanism may mirror the actual time the clock is showing in the physical world. People can ring the bells and hear their sounds, glance out the window to view the town, or look up at the tower's attic to observe the architecture.

Simulation encourage hands-on learning; parents may wish to take their children to such a world to show them how clocks work, or may wander around simply to explore or meet people.

Gray's Chocolates is an example of a virtual storefront: people can wander in and look at the things they have to offer by browsing through a hypermedia catalog, read related material, or order a custom box of candies by dropping goods into any number of packages and ordering. A virtual business card stands upright on the counter; people may drag a copy and save it in their hotlists or media wallet. A bell may ping someone at Gray's to help me if I have any questions. Instead of working at convenience stores as clerks, will teenagers in the future instead be working at virtual counters?

These worlds are meant to encourage exploration and learning. There is much more work to be done before this technology can reach its potential, and many people are beginning to contribute to a wide body of knowledge whose direction is taking shape. Where we go is up to our imaginations.


- Clay's worlds can be found at the VRML Repository under <http://www.sdsc.edu/SDSC/Partners/vrml/examples.html>
- As of this writing, the EIT worlds have not yet been publicly released.
- A 3D graphics book list, available graphics engines, and 3D programmers' resources are at <http://wclrs.bham.ac.uk/GamesDomain/gamedev/gprog.html>

4.6 Common Problems

Environments made for manipulating and visualizing media often suffer from the problem of incorporating too much clutter and visual noise. In many cases, less is certainly more. This is one of the advantages of command-line interfaces; there were often few elements to mess up the display. Just as it is possible to have gratuitous links on Web pages, it is possible to make gratuitous use of 3D elements. If the elements hinder more than they help, then they are unnecessary. If users can easily use and understand the environment, perhaps their use is not so bad. But developers still have to be careful about making use of excessive detail.


In many cases it does not amount to a black-and-white issue; the environment should be tested on people and feedback encouraged in order to make things more understandable. What features can be omitted in applications? What do people really need?

Here are some common problems with visual systems:




Click on the fish and select the umbrella to work on your document.

The problem with icons



Open the folder and get a piece of paper and a pencil to work on your document.

The problem with metaphors



Drag the triangle to the square to work on your document.

The problem with symbols

68: The problem with using signs.

One problem with icons is that developers sometimes choose the wrong icon to represent objects. The wrong choice of icons can confuse users and leave them asking, What does this mean? What am I supposed to do?

Many two-dimensional environments make heavy use of metaphor in their structures, but sometimes the elements cannot do everything the metaphor implies, so users become frustrated; improper metaphors are used, leading to confusion; or metaphors are used to their maximum extent, making the use of media needlessly complicated. If I have to do everything exactly as I do in the real world, I may understand how to do it, but it won't necessarily make my life easier.

The problem with using symbols to represent elements is that the developers and users must have a common understanding of what the symbols mean. Language is symbolic in this sense: I cannot make you understand my ideas if you don't know what I am talking about. Just as with color, developers should make sure they are not using imagery or elements that may be culture or language dependent. For instance, the "thumbs up" sign is considered offensive in the Middle East, and the "OK" sign can be insulting in other countries.

Avoid overloading and information overloading. Although it is recommended that functions be combined in an elegant manner, it is possible to overdo it and make one action or element serve far too many purposes.

Developers should also refrain from incorporating excessive assistance in the environment or making things far too obvious. Once users are used to a system, all they may need is a quick glance at some sort of reference manual from time to time, rather than help texts and dialog boxes popping up at every turn.

4.7 The Open Interface

We must ensure that information environments remain open and freely available. We don't patent the concept of a book or magazine, yet these are human-interface information systems. Patenting common elements will only serve to impede the advancement of human knowledge and science; not promote it. In time the use of such systems will appear as familiar to us as the opening of a book or the switch of a channel. And nobody has patented the look and feel of a television.

We see in consumer products the free adoption of useful interface elements, like the VCR jog shuttle, lights on rechargeable shavers, trackballs and trackpads on laptops, and icon animations in browsers. Common graphic elements propagate throughout the Web as memes, as people discover new ways of presenting information. What would happen if these things were patented, what if we patented the scroll bar, regardless of the fact that it represents a new and unique invention?

When I can completely customize my environment to act and look exactly like someone else's the point will be moot anyhow. The danger lies in locking people into one world view and expecting them to live with it. Everyone is different; perhaps the search for the next great environment really is a search for the human identity.

5

Tomorrow's Internet

5.1 The Future of the Web

In time, the current frenzy over the World-Wide Web will pass, as all things do. The excitement will wear down, and people will begin using the Internet as an everyday tool, as a normal appliance. After that, what next?

One can answer this question indirectly by describing what the Web lacks today and why other hypermedia systems failed to become as popular. Why is the Web so terrible?

This is not to say that it isn't great. It provides a simple, easy way to find, transfer and navigate through information. In large part, people can communicate what they wish to say to others. It is just flexible enough so that people can be creative in the medium, and the basics of HTML are just simple enough that people can learn how to write it. So these are good reasons why it should be popular.

Other hypermedia systems that weren't so popular include the innovative Intermedia system from Brown University. It allowed links to many kinds of media, did not distinguish among authors and readers, supported document version control, and came with editors for many kinds of media. But it could only run on A/UX, Apple's UNIX, and came with many different parts that made up the working environment that may have been difficult to manage and install. Intermedia did not produce uniform layout, choosing to represent links with dashed boxes and arrows that broke up line spacing and made documents difficult to read. Awkward link-oriented functions were performed in dialog boxes, not the most intuitive way to connect one thing to another.

Many other systems slowly faded away for many other reasons: no working code was available or it only ran on unpopular systems, no public distribution was offered, and no research or results were put online for people to see. Often systems would be tucked away in obscure proceedings, never to see the light of day. Systems that did exist were difficult to set up and configure, hard to use, expensive, or just plain looked ugly.

So it appears that new systems should:

- Be implemented and disseminated as soon as possible
- Be conceptually simple at the most basic level
- Hide unnecessary internal elements from users

- Show users exactly what they have created
- Allow users to easily create, view, and link media
- Be affordable

How does the Web fit into this picture?

Standards are not implemented in a timely manner. Extensions to HTML, URLs, and CGI are being implemented left and right in commercial browsers and servers without any say so from anyone else; in many cases these extensions conflict with what official proposed standards say. If a standard is ignored long enough, eventually conflicts appear that wedge themselves into cracks and slowly split things apart. There are already a number of different browsers that only display media tuned to their application the best; media that is supposed to be universal then becomes “optimized” or “enhanced” for particular commercial applications. What should the average person then choose to buy? For what application should the author create media? The infamous blink tag started as an undocumented joke. Do developers really need to inflict this kind of misery on their customers?

The Web is getting more complicated. Users have to deal with the guts of the machine. New users should not have to know the internals; we may not know how cars work, all that is important is that they go. Many people who have never developed a popular human-computer interface in their lives are creating the methods by which every person wanting to view or put hypermedia on the Web must follow. People should be given a break; many don't even know how to hyphenate “World-Wide Web” properly.

I should be able to type in a `www.something` address directly into my browser; the application can then infer that it is an `http://` address or guess at other protocols according to the name. Now that URLs are becoming ubiquitous, there should be no need any longer for people to have to say “colon slash slash”.

Some browsers and servers are forcing people have to learn hexadecimal numbers, lists of coordinates, and strange codes in order to do what they want. People should never have to experience hexadecimal numbers unless they are learning assembly language. Instead of having to type hexadecimal numbers or measurements, action-oriented layout creation should be implemented.

Applications of SGML, the Standard Generalized Markup Language, are often nasty little tongues for people to have to write manually. SGML, a language that in some form or another has existed since 1969, got its big break in the early 80's, when organizations such as the U.S. Department of Defense's CALS (then Computer-aided Acquisition and Logistic Support) initiative and the Association of American Publishers (AAP) adopted SGML as internal standards. The AAP used SGML as part of their Electronic Manuscript project, and their application of the language included support for things such as tables and mathematical

equations. In 1985 SGML was adopted by the Office of Official Publications of the European Community, and thus quickly made its way to CERN.

The use of the phrase “electronic manuscript” makes about as much sense today as the phrase “horseless carriage”. SGML as a standard was born during times when Text Processing was a serious field. Text - numbers and letters.

Text has been heading on a collision course with every other media known to humankind. People transfer far more things than just text over the Web; is this why is the creation of hypertext is so much harder than the creation of other media? There simply aren't any good, fast, intuitive, cross-platform authoring tools yet for the Web. Current HTML or SGML-based editors and browsers still look like they have been written by people who are completely hung up on text.

In one sense all of this is good, because these things are forcing users to become less naive, in the sense that changing a flat tire helps you know your car better. But wouldn't it be better to create solutions that help prevent getting a flat tire in the first place?

Authors never get what they expect. SGML has pursued the laudable goal of separating data in media from its presentation format by using generic structures. The reasoning is that it allows writers to focus on the creation of content rather than layout.

The problem is, without proper layout, content can be rendered quite useless. It is fine to have systems to concentrate on one thing or the other from time to time, but the Web is an unbalanced system. Many things that are taken for granted in common page layout programs are simply not possible in the Web environment. How content is presented can be just as important as the content itself, and people should be given maximum flexibility to get their ideas across just so. As the average user becomes wiser, they will demand this.

It is getting more difficult for people to create, view, and link media. Create a table of five rows by two columns in Microsoft Word. Now write it in HTML. Which is easier? Create an alias of any document on your Macintosh. You have just made a link to that document, which can be put anywhere in your desktop environment. Now create a link between an HTML file in your Web environment and a QuickTime movie. Which was easier to do? Make a Word document on a Macintosh and then open it up in WordPerfect on a PC. Does it look similar? Now make an HTML document using your favorite HTML editor and then open it up using your favorite Web browser. Does it look just as similar?

It should be so easy to do so many things. Developers need to change their way of thinking and put the user first. If they had to make a single program to manage all aspects of their lives, how would they implement it? What current systems or ideas would they emulate?

In time browsers will shrink away, its features becoming incorporated into our computing environment. Then there will be less emphasis on viewing read-only pages, or a text-centric view, and more emphasis on getting things done with whatever media one chooses to handle. Why should systems force people to design to a particular layout or size? This is what the Web is increasingly doing.

It might be easier to move away from such an emphasis on reading, clicking, and filling in forms on the Web if a powerful, easy-to-understand interactive scripting language like HyperTalk were implemented. And should HTML style sheets come into being, it will be much easier to create platform-independent layout while still maintaining a good deal of control on the author and viewer's sides using defaults.

HTML is not programming, although the Web should make use of something that is. If one defines a program as "a sequence of coded instructions for insertion into a mechanism", then one may make the case that HTML is just structured enough to be considered a program, although it's a very, very basic type of program. It is like calling an audio compact disk a giant program - the structure and sequence of the encoded instructions exists because in a large part it must be so for the disk to be used. The author has little freedom to change how these instructions are encoded.

Calling HTML programming is an insult to programmers who must deal with encoding behaviors, threads of logic, and interactivity into their creations. HTML cannot support these things to any reasonable extent and is not the proper language to do so.

Writing HTML today is the equivalent of using a highlighter. Calling HTML programming is like underlining words on a page with a Magic Marker and then calling it art. Both VRML and HTML may one day have features that allow complex programs and behaviors to be incorporated into media, but developers will have to take a good hard look at what they are trying to do and make sure it makes sense.

The World-Wide Web Consortium needs to keep an eye on the corporations that would incorporate proprietary or closed protocols in their Web software. There is no problem with having such things coexist with open standards, but the danger lies in the creation of incompatible media, or media that has a very high barrier to entry. How can we create a common medium?

The W3C should apply itself towards the design of global hypermedia systems! The Web will change so drastically that in ten years the only sure thing about it is that it will be global, and it will transport media. There are far more proposals and solutions in the world that can be tried out, and many good ideas exist all over to make the Web more flexible, robust, and intuitive than it is today.

As more people begin to use the Web every day, it will become increasingly mundane. I use it to read the daily news, cartoons, magazines, and do all manner of work; for many it has ceased to

be a wonder and has become integrated into their daily lives. But current browsers make inadequate personal organizers and awkward calendars, are terrible at presenting layout and presentations in general, and are in no way suited to become one's sole computer environment. Instead, what will eventually happen is that the best aspects and philosophy of the Web will likely become integrated into tomorrow's virtual environments and media. The W3C needs to realize this and must begin research into areas that can take us all far into the future.

The Consortium needs keepers of the flame, a *clear written philosophy* and a set of developers' guidelines that everyone can refer to so people don't keep trying to fit parts into holes for which they were never meant. There is so much to be done today that all those wanting to improve the Web need do is simply choose a direction and go.

5.2 New Business


Information wants to be free...
Information wants to be expensive.

Stewart Brand
The Media Lab


To put it more correctly, users want information to be free, and providers have traditionally wanted information to be expensive. In the new online economy, content itself will be relatively free - what many people will end up buying and selling are *services*.

The old model of publishing information - the "push" model - is not efficient in the online world. It puts a great deal of emphasis on the publisher as an inflexible institution - it edits, assembles, designs, and prints or manufactures what the author gives it. In many cases the publisher depends upon the distributor to make its products available to the public; in short the author has pushed their work through a long process so it can be eventually purchased by users. The emphasis is on the pushing of information along this chain; it can take some time for user feedback (in the

form of market demand, for instance) to propagate all the way back to the author and even then it may have gone through multiple middlemen to get to its destination.



69a: The traditional push method of distributing information.



69b: The pull method of distributing information in its simplest form.



69c: The pull method using third-party services.

Distributors can differentiate themselves from the competition by pricing competitively, offering a specialized or more broad selection of products, and by positioning themselves in a geographically strategic position, among other things. The Internet changes the picture because geography is no longer a factor, the entire process takes place in the digital world, and everyone can be afforded the same quality of communication - fast and direct. The emphasis is put on pulling information; it is users that can initiate the entire process, it is authors that can edit, assemble, design, and publish, but publish in the new sense does not mean *giving something to others*, but now means *allowing others to take something from you*.

The process then becomes boiled down to: users take something from authors. This is similar to the way the do-it-yourself 'zines market works: I send a self-addressed, stamped envelope (with extra stamps inside) and perhaps some form of payment to an author; and they send me something. When payment facilities are enabled that can allow this type of transaction online, there will most likely be a lot of it. However, much online information publishing will make use of new types of middle-chain services. These services are the hybrid ancestors of today's publishers and distributors, and they will exist to serve many purposes:

- To edit and convert works to appropriate formats
- To handle and analyze large volumes of transactions
- To advertise and store works
- To serve as a top-level repository for particular information
- To summarize and send feedback to the author

- To receive and assimilate information from the author
- To bundle the information within and among value-added packages
- To communicate with similar services to coordinate information flow and share resources as needed
- To provide tools to allow the author to automate any of these processes

These services are competitive but can also be cooperative - general service providers can farm out specialized work to smaller ones. The process becomes efficient, it is open, it is a win-win situation for all involved parties. Many larger conglomerates today try to build their own "content generators", in the form of new departments or spin-offs, despite the fact that they know nothing about the field. Just as those who think the Internet is simply the world's largest television set, they are trying to get the new business to do the work of the old.

One of the main reasons that these new services exist is because authors cannot always do everything themselves. It is easy to say that independent authors will be able to rival the products of the largest corporations, because sophisticated authoring tools are getting cheaper and more widespread. But at some point we will all have sophisticated tools; those who can afford more sophisticated ones will make use of them. The traditional rule has been that cooperatives of people can pool together more resources and produce better works than people individually. We will see whether or not this holds true in the coming decades.

When everyone has technology, the important thing is not necessarily how fast the computer will go, or what kind of software one has, but how do you use it? Are you trained to be a generalist, so you can work with these tools? This is a human resources problem, an education problem. If tomorrow's work force is not educated or trained in some manner, there is no work. If public education systems fail, companies should make use of on-the-job hypermedia training, vocational work, hands-on networked seminars for employees. Learning by doing is the main idea.

The New Library

These services are unique in that they merge author-oriented services with user-oriented services. They can maintain constant communication with authors, sending them analyses and up-to-date reports from the field. On the user's side, they will have to be able to provide easy-to-use methods of information search and retrieval, value-added information, personalized information, customized information. When everything is online, users must have a way to:

- Be informed of new information
- Find information that is relevant to them
- Find information that is trustworthy

Don't always assume that information is truth. When paper was the common medium centuries ago, books were sacred objects. Publishing a work meant validation of its contents; it implied that one took the time to arrange and carefully plan their thoughts before committing them to a fixed medium. Our common medium will not necessarily be physically fixed, but information can be fixed in versions, it can be compared, cross-referenced, and verified in ways that paper media never could. Users can check facts for themselves, creating their own annotated media or contributing to a common version. The new services must be able to provide or at least support ways of doing these things.

When we reference paper works in a bibliography in a scholarly report or thesis, we often have little choice but to accept them as truth, because we may not be able to find the original in the first place, or simply because of traditional inertia. When such references are electronic, we can see them for ourselves, but the danger lies in the fact that the author may have edited the referred version without our knowing or that someone else may have altered it without the author's consent or review. Without version control methods, permanent references to physically reliable electronic repositories, or common security mechanisms, the online world must remain in suspicion by users, providers, and governments. Without the ability to permanently fix works, the information that we experience every day can quickly become irrelevant and monotonous.


For businesses to simply have trustworthy information is only part of the solution: much of the problem then lies in trying to get users what they are looking for. Successful services will do as much as possible to prevent users from trapping themselves in informational blind alleys when trying to find something. People may wish to know where they can find the latest news and resources: there is always something happening on the Internet. In fact, there are so many things happening on the Internet at once, it is extremely difficult for users to figure out what is going on. What is needed? Abstract services, guides, online newsletters, reports on community happenings, the electronic equivalent of the town crier. Quality reporting that is personalized so as to be relevant to the consumer will certainly be paid for.

One of the often-cited problems of trying to find relevant information is that one can't truly tell if something is truly relevant without looking at the information. This problem can be alleviated somewhat by doing things such as providing accurate abstracts, releasing trial versions of products, and encouraging users to use word of mouth (or line of email).

"Be sure to check often". Providers do not need to make use of this repetitive begging if they offer something interesting to users. How can providers tell users they will be offering something interesting? Provide a future schedule of updates and new information. Let them know exactly when a new issue or version of something will appear. You can build it, but they will not necessarily come until they are told about it first.

There are many businesses that will revolve around manipulating information: How much will you pay me to bundle your information with something else? How much to send this information to somebody and verify it? How much to let me help you make the most of the information you have? There are more than enough opportunities already - the important thing is to design for the new medium, and not for the old.

Differentiation



70: Information flow in modem-based services and Internet-based services. The top diagram represents the flow of traditional modem-based services in which users connect to a central provider. The bottom diagram represents a fully distributed approach in which users are also providers, the classic view of the Internet. The flow map of the Internet will probably end up like the middle diagram, in which larger services and smaller user/providers coexist.

New services can gain lasting audiences in a number of ways. They can focus on working with specific types of information and particular kinds of products. If a service continues to provide quality (i.e., trustworthy and relevant) information, it can do well. The ability to offer a broad or comprehensive range of material is another differentiator, but in an environment in which so much information already exists, it must be different from the traditional warehouse distributor types of businesses. To offer a large range of products in the online world, the service should provide a way to customize or personalize the information - they should add value to it by making it easier to use, adding more features, cross-referencing it with related things, cooperating with authors and manufacturers to make their services unique. The emphasis then falls on the kinds of personal and real-world services that can be provided. Is customer support helpful and friendly? Do I get my information on time in the format I want? Is there online help, a staff I can talk to that is knowledgeable about the products? Is the service fast and easy to use? Can I find what I'm looking for quickly, can I compare this to something else? The superior online services will have answered all of these questions affirmatively.

Businesses exist to propagate themselves; services exist to serve their clients and in doing so propagate themselves. Partnership for a specific task is fine; cooperation for a certain period is better. The topology of the large online dial-up services once resembled a star - users could connect to them and meet and share information in one place. As these traditional services merge with the Internet, their topology will resemble something more of a web and may one day make the transition to full-fledged, successful Internet services if they do the right thing. Their users will increasingly have the ability to be their own providers, hold or join their own meetings elsewhere, and use less restricted forms of communication.

There are a number of ways these large services can continue to differentiate themselves, in addition to the ways described above:

- They can offer better hardware and software resources
- They can offer better communities
- They can offer lower prices

In terms of equipment, there is certainly a lot of room for competition. Some services mirror popular Internet archives and have made them publicly available over fast networks with machines

with lots of storage space. This is a good thing to do, as it embraces the new users on their modems as well as the old on the Internet. It is a good step towards transition. Other ways to improve the services includes having convenient points of access, fast connect speeds, accessible, consistent interface, and the ability to interoperate with a number of network protocols, going both ways.

One of the chief ways that the services can become Net-friendly is to offer better community-based resources. With enough hardware it becomes possible to develop resources that foster large amounts of specialized communities in ways that has not been easily done over the Internet. The main gripe many long-time Internet users have about the modem services is that they are diluting the amount of good conversation and resources available. This is true, but it would have happened whether or not the large services existed. There are many millions left in the U.S. and billions in the world that have never touched the Internet.

What the services can do is create resources that foster the creation of limited communities - they might offer "sign-up" boards for special interest groups that are limited in size. They may appoint moderators, experts in particular fields to answer questions or delegate others. They may hold online classes, provide space to hold hands-on workshops, tutorials, manuals, galleries. Using these online services today is like seeing the world through a very small and dirty window; you may be able to make out human shapes, but you have no reason to talk to them or they to you. And should either of you communicate, you must do it in such a way that your voices pass through the glass. Who can create a community in this environment?

One of the most important things for any service that connects users to the Internet, or any network with people for that matter, must do is educate new users. Force them to read the guidelines, the local etiquette. Integrate this into the user's manual. Make the community histories easily available to everyone, freely; let them read about traditions so that they can create their own and write about them. Make it a *place* that people can respect, not an anti-septic environment crammed with buttons and advertisements. Would you want children to grow up in this environment? They are.

The Four-Fold Path

Many traditional companies as well as new ones trying to business on the Internet can do a number of things to help themselves and the online community.

Learn new business. The Internet encouraged the growth of a basic cycle of development which helped it to blossom into a wonderful ecology, much in the way modern science runs on the scientific hypothesis. The cycle is: one makes working code, it is given to as many people as possible, feedback is received, and a standard and perhaps a good product is produced. If one has

ended up with a standard, it is then used to build more working code, it is given to as many people as possible, and so on, and so on. This is the way of the net.

Businesses certainly need to make money, and it is difficult for traditional businesses to understand how to work within the new medium of cyberspace. It is possible to do so by doing any number of things: Work openly, establish yourself as a leader in the technology, differentiate your product, make unsupported, limited products available with wide distribution and sell supported ones. There are and will continue to be successful companies that do this. They have learned the new way of doing business because they know and understand the net. In short, they understand the future. Many stories have been told of the CEO in the traditional marketplace that hears about something on the net and decides they can make a quick profit by leaping into the business. But once they get started, they often have no idea what business model to pursue. They do not know the net. Learn the new way and make it your own.

Partner and cooperate. There are numerous core sets of general technologies and issues that organizations must be able to collaborate on openly in order for the Internet and related technologies to advance beyond the Dark Ages. Supercomputing has its “grand challenges” - in this age, where supercomputers are becoming desktop computers, one needs to work on the grand challenges for Internet development. The online ecology is innately global, extremely fast-paced, and full of small, nimble businesses that can dominate their fields by being able to coalesce, develop open standards, and then differentiate themselves, thereby raising the level of technology for everyone as well as ensuring market share. Large, slow, vertical corporations will find it very difficult to work in a similar fashion.

Understand the Internet. Follow current research, stay on top of developments, listen to new and experienced Internet users. Understand the culture, give credit where credit is due, gain a proper perspective on where things are. One will be more likely to understand where things are going. Don't listen to buzzwords. To create a lasting organization, listen to the developers, the artists, the designers, the kids in the garage, and the hackers in the universities. McLuhan succinctly wrote that it is always artists who discover how to enable one medium to use or release the power of another. These discoveries inevitably unleash great opportunity.

Follow the standards. Participate in their creation. Join consortia. Help to produce guidelines for developers and benchmarks for implementors. It is very tempting to produce products that fight current standards in hopes of “locking in” customers. Corporations usually do this on their own when they wish to crush the competition. If the strategy fails, however, they run the risk of ruining themselves or at least wasting a lot of money. More often

than not if they win they do so because of politics or marketing, in which case consumers lose, ending up with an inferior product that cannot operate with anything else but its own kind.

Organizations that do not wish to crush anyone yet still make conflicting standards do so either because they are dissatisfied with the standard and wish to change it or create a new one, or they fear that if they don't, the market will pass them by. These are valid reasons. Certainly in the first case they can work openly during the process, or propose a draft standard. But there is nothing a company can do if the standards process is too slow and/or unresponsive to market demands.

Certainly the phrase "standards are dictated by the marketplace" is obsolete. This is not because standards won't be dictated by the marketplace, but because *what the marketplace is* and how it makes itself known to companies is changing. When word of mouth becomes nearly instantaneous, companies will experience instant market saturation as well as decline. Standards can provide places to anchor in this turbulent ocean of supply and demand.

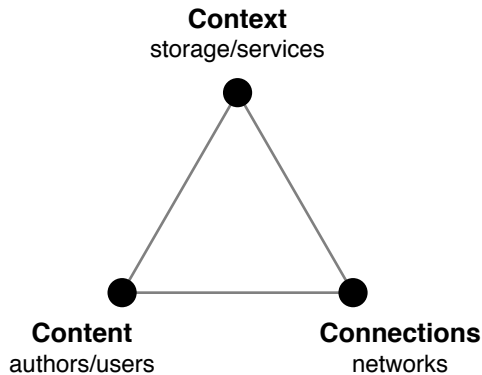
Because of the acceleration of the online environment, consortia themselves must change if they are to survive: they require dedicated members and must enable efficient information exchange, regular testing processes, and timely market analyses. It takes a good deal of energy to accelerate large bodies, but it is something that must be done for open systems to survive the coming decades.

If we wish to Balkanize, stratify, and segment the Internet; if there is a need to isolate users and their equipment; if it is in the public interest to transform the cyberspacial ocean into a collection of data puddles, then open standards processes should be ignored. Until then, all of us should endorse them, or if not, try to improve them.

Content and Context

Some people say that in the online world, content will be king. Others say that content is not important, rather, it is context - the environment in which the content is viewed - that will reign. I would say that it's not about what will dominate, but about how these things will coexist. You cannot have context without content. We cannot realistically live in a world of content without context, that is something television would have us do.

Content providers should work with context providers. Both of these groups should work with the people that provide the hardware necessary to get things done. Access providers can supply fast hardware, software, and storage; these are the basis of tomorrow's public squares. Telecommunications organizations provide the means to connect them and connect to them. The entertainment industry, the scientific and business communities, and the public at large will not find it difficult to create content.



71: A balanced triangle of context, content, and connectivity.

5.3 The Medium is Still the Message

We then end up with a triangle of equilibrium, with context, content, and connectivity. There are many large companies today that handle two or all of the areas, but this skews the balance. No area can exist without the others; without such balance the Internet becomes a wasted resource.

We can view using the Internet as going on a trip somewhere else. We often ask ourselves these questions:

- Where will I go?
- How will I get there?
- What will I do when I get there?

These questions illustrate the basic necessity of these services. Context helps you know where to go, connectivity helps you get there, and content is what you experience when you're there. The questions of when and why should always be left up to the end user.

Nicholas Negroponte has written that the medium is no longer the message, rather, it has become an embodiment of it. When the content we receive is computer generated and personally enhanced, what is the message? Will the medium have any bearing on how we think?

I say the medium is still the message: when the medium, computers, permeates our world and embodies all content, then the message becomes everything as well and reflects ourselves, everyone, and everything else. The content we receive can reflect ourselves, in which case we personalize it; it can reflect other people, in which case we listen to what the world has to say; and it can reflect the things we build - our computers, whether they are built into television sets or kitchen appliances, can talk to us. When the medium can show everything, we can become anything. That is the message.

But when every person can individualize everything they receive, when we can all filter out things we do not deem important to ourselves, then where does the common culture go? Newspapers used to serve as the common culture: gossip columns and daily events provided a core set of topics to talk about and ideas to rally for and against. It is this common culture that fosters community. Today we receive our common culture from television and movies. This would not be so bad if our common culture were not presented as disconnected entertainment. Unfortunately television and movies have predictably shaped American culture into random blocks of trivia - there is no common ground to hold onto.

One reason this has happened is because television and movies are a viewer medium; there are relatively few authors. It is not easy for people to provide feedback or become authors themselves. It is not very participatory. Services on the Internet, on the

other hand, can be highly participatory and authoring media can be made relatively easy. It is easier for people to create their own communities and common cultures with others.

This is good in the sense that the generation growing up today will be used to diversity and adaptation, but the problem is that today's mass media has made us all complacent, used to experiencing diversity and culture rather than creating it. We must learn to think differently and change ourselves rather than let the old media and television culture drag us into light-speed stupidity.

Just as writing made use into linear thinkers, and television made us into non-linear thinkers, cyberspace will encourage new forms of thought, new ways of thinking about ourselves and others.

We can diverge off into at least three extremes: we can each become our own universe, paying attention only to what we feel is relevant; we can enjoy a world view, listening to everyone else but risking the danger of information overload; and finally, we can think like machines, forming a close bond between ourselves and our devices.

The idea is, through education, we can help people strike a balance among these three extremes. None is healthy by itself, but all have benefits that we can take advantage of. Those who would focus only on the extremes and denounce new technology as useless would rather use no technology at all rather than work towards making it beneficial to all of us. To properly use technology and media for positive growth is the hardest path to take, because it requires knowledge, cooperation, tolerance, and self-control. Americans live in a society with no sense of moderation and a failing educational system; thus it is doubly hard to concentrate on the positive rather than the negative.

Newer Media

In early 1950's, the film industry felt threatened by television. This reaction has seemed pretty logical for some time, but today we can take a different perspective on things: movies are simply television with a larger budget, and television is movies with a smaller budget. The old way of providing an immersive experience was to increase the size of the display; the new way of providing immersive experiences will be to shorten the gap between the display and our senses. The old way of sharing an experience was to bring people together in the same room; one new way of sharing an experience is to bring people together in the same virtual environment.

In cyberspace there will always be room for old media and old representations, but they should not and cannot be the dominant forms. There will be media that we can passively experience, as humans need to be able to do this from time to time, but humans should not and cannot be passive creatures.

Today digital movies on the Internet are still somewhat in the silent era. Many are postage-stamp sized and have no audio track. Television displays once were as small also, yet today we can create displays as large as walls. Today we must stay a cer-

tain distance away from monitors to avoid ruining our eyes, but with microchip displays and lasers we will be able to beam signals directly onto the retina. The idea is that display-dependent media will be less emphasized - Negroponte has written about how silly it is for the HDTV standards parties to argue over such things as resolution, frame rate, and the shape of the screen. With a digital format, this is entirely unimportant. What becomes important is the quality of the content: Does it get the message across? Can I see the big picture? Can I find it easily? Is it worth my time? Is it trustworthy? Can I contact the author, see related media, see it in different contexts?

One new medium that is being proposed is interactive television. Is it a medium of any worth?

- **Many of the available boxes are not compatible.** A great number of proprietary systems with their own particular methods are being tested. Not only can they not talk to each other, but each interface will vary from system to system, making users learn every new control when they move or change boxes.
- **It is not a truly two-way medium.** Bandwidth going back to the cable company is typically very small, although for net applications this is may not necessarily be the case. For companies to set up such a lopsided connection is shortsighted; users can give as good as they get in the new scheme of things. If they do not have bandwidth both ways now, they will demand it in the future.
- **Most people watch TV to waste time.** Televisions, the equipment they connect to and the rooms they are put in, are miniature temples to leisure entertainment and are not the appropriate venue to introduce highly interactive applications. TV trains its viewers to expect passivity. The average computer user is more used to high levels of interactivity.
- **It is expensive to set up and maintain.** Some systems are costly to implement, with expensive boxes, equipment, and regular fees. Many systems offer complicated remote controls or controls that simply add to the clutter of equipment.
- **It is forcing old media to do the work of the new.** To retrofit old equipment and networks is costly and backwards. Television is not inherently interactive nor is it a computer. In time the average homeowner will have a fast computer in their house. Why introduce another one that is only good for shopping, movies, and games? Television as a medium is already being encompassed by the computer; it is time for a forward step and not a backwards stumble.

Then what would work for ITV? First, it should be something that viewers can watch passively, that they can reasonably evaluate, that they can offer feedback on, and that they can order in an exceedingly simple fashion. For these reasons, home shopping channels and infomercials have been wildly successful, no matter how banal they may be. One way to offer video-on-demand services with these features would be to allow people to read reviews and capsule summaries (say, one to four stars) by critics as well as other viewers and show the first hour of every movie free, to give an adequate amount of time for evaluation. After the first hour, the viewer would have to pay in order to see the rest of the movie. This provides much more of an impetus for payment, and users can still watch a good deal of movie portions passively as entertainment.

Many people today wish to digitize movies mostly for video on demand services. Today video on demand is not price-justified because it is costly to digitize movies. But for books digitization is easier - now that many books are written electronically, this bypasses the entire process. When movies are made entirely digitally, as they should be, with digital storage devices instead of film in cameras, they will be easier to distribute electronically. Digitized movies will allow directors to emulate any type of format: movies can be shown as Super-8 works, in different film grains and tints, in black-and-white, or in green-and-orange if they wish. It will open up the old medium into new and amazing areas.

An older medium that is increasingly being used today is CD-ROM. CD-ROM authoring used to be called "optical disk publishing" or "interactive video publishing". Phrases like these often occur in the early years of a new medium because people are highly fixated on the device rather than the ultimate function.

CD-ROM was an important evolution of media - it allowed access to large amounts of information at a time when popular mass storage devices were bulky and not as reliable as they are today. But in time we will go beyond this to a point where it will no longer matter much what new type of storage device is invented. A compact disk is a communications device only to the extent that the information on it can be experienced by somebody.

People will not necessarily have to care where their information is physically stored, if they can have access to any of it at any time. A networked device is a communications device. From today's point of view, CD-ROM is just another storage medium.

However, as long as portable storage mediums are used, manufacturers should decide on a common, generic format. It would be very convenient to be able to replace shelves of videocassettes, game cartridges, laserdiscs, and floppy disks with compact disks. Why should users bother any more to shuffle data in and out of their computer in one and a half megabyte chunks when they can make use of any number of devices that have superior price to performance ratios? As long as I can find them and use them securely, it will not matter where my storage or my CPU is.

When I have access to computing resources wherever I go, I also have the capability to capture real-world media and store it in my environment. One can take as an example the Media Lab's work with salient stills - a series of shots of an object from different views can provide enough information to produce a large, highly detailed composite.

Imagine all the outtakes and incidental shots from other people's camcorders. The generation growing up now is the most filmed, recorded generation on the planet. In time, much of this recorded material will be online and easily manipulated - I could grab all the incidental shots of someone walking in the background, create a 3D construct from them with high-resolution texture maps, tell a program to analyze their walking behavior and program in behaviors and facial expressions. The incidental material that thousands of people gather daily can be used to create virtual people, places, and things. In the same way that Andy Warhol wanted to recreate himself as an android, this will be possible to some extent virtually. Certainly we may all be famous for at least fifteen minutes, but how we define "famous" may change drastically.

In terms of music, many new possibilities present themselves. Brian Eno's *Music for Airports* brought about a new form of music - ambient - that could be written for particular places, to evoke particular moods. Virtual environments can have their own music tailored or generated in a customized fashion. Traditional media has often assembled separately: in making a music video, the music is recorded first, and then the video is shot. For movies, the film is shot first, and then the soundtrack is recorded. Online we will be able to do all of these things at once, where the creation of an environment includes the creation of an aural composition, where things blend together and objects create their own music to create virtual environmental noise.

Our technology defines the times - just as "black and white" brings memories of pre-1960, and faded yellow images suggest America's Old West, later decades are defined just as well by their media. MICR, an optical recognition font, was used widely throughout the 1970's to suggest images of the future. Entertainment of the day, such as *The Andromeda Strain* and *Rollerball* made use of similar fonts; now when one sees them, they are equated with the 70's. The 1980's and 90's is being defined by designers such as Neville Brody and Barry Deck, who has designed such fonts as Template Gothic, based on older forms of hand-made template lettering. Today's popular typefaces speak of postmodernism, imperfection, geometry, and world conscience. Rave art sees psychedelia through the computer and encourages us to expand our minds through the new medium. Much of what we see is coming from kids on their own working at their personal computers, using cheap scanners and pirated software; they have a need to create, and they are defining the look of today.

However, a lot of today's style borrows from the disco era of the 70's, leaving little to creativity. To recycle someone else's mainstream culture is not just nostalgia, it is saying that one has nothing new to contribute. With the new forms of communication that are evolving under our noses, artists have an incredible amount of opportunities to contribute something new. It is they that should be noticed; they inevitably help point the way into what's coming.

What might be the look of tomorrow? Just as hyperrealism was the rage in Japan as airbrush art took off a decade or so ago, we may see highly realistic artwork produced with computers, three-dimensional, responsive to interaction, high resolution. But the new style may very well have a number of global elements in it, incorporating multiple streams of media, showing stylized, clean lines or perfectly rendered versions of ancient technology.

Even when we are all online, there will always be an edge to look over as well as a place to look back. The idea is to take the best ideas of the old and use them something new that has never existed before, something that no one else can claim and that everyone can learn from.

5.4 New Society

Some people worry about the social divide between the information-rich and the information-poor, the haves and the have-nots, the First and the Third Worlds. But the real cultural divide is going to be generational.

Nicholas Negroponte
Being Digital

Something is Coming

A survey taken in the first half of 1995 by Stanford Research Institute shows that the reason that people are not on the Web is primarily because of education, not income. If we wish people to be online, they must learn what they can do, what they can become, and how they can handle new technology.

Perhaps not surprisingly, many surveys show an age gap. According to a recent Times-Mirror survey, 41% of people 50 and over have personal computers, but according to the Third World-Wide Web User Survey about 10% of this group uses the Web. SRI's study shows that a significant majority of World-Wide Web users are in the "Generation X" (their words) group.

So it appears likely that the average person using the Web today is young and fairly well educated. In time this will certainly better reflect the demographic makeup of the United States, and hopefully, one day the world. By the time the current American twenty-something generation is old, every rest home will likely have computers so that its physically infirm inhabitants can be productive or explore cyberspace freely. Today many senior citizens refuse to stop driving even if they are considered a hazard; to them the car represents freedom of movement and independence of thought. Perhaps one day the computer will represent freedom in this way as well to those who cannot be physically independent.

One indication that Web demographics are evening out is the decreasing gender gap. The Third User Survey has found that the percentage of females using the Web has increased from 6 to 15% in little over a year, and SRI has reported that the Web-using population is 30% female. Times-Mirror reports that 27% of females in the United States have personal computers.

The Web and the Internet as a whole will also mirror more of culture as time goes by. By definition this will make the net more average - instead of only reflecting the lifestyle of the smart young male, it is reflecting all of society. Like any media, people will find their own uses for cyberspace. Those who started the computer and Internet revolutions are smart enough to find new communities of their own and find other spaces that they feel comfortable in. The average computer owner will have to learn how to make their own media and communities online as well; this will prove to be two of the necessary skills of the future.

Today students of all ages are making home pages and assembling media on the World-Wide Web. They have learned that computer skills are necessary. They have Internet accounts and access through their schools and universities. When these students graduate, many will look for Internet access providers - think of it as acquiring new furniture and then having to move. In time, all students will be given net access, just as surely as they are given calculators, pens, and paper.

One can try to predict the kind of changes that will be taking place by looking at the rapid technological changes of the past few decades. The video game industry is an interesting case - in 1977, Atari released the VCS 2600, introducing video games into countless American homes. Up until the early 80's, the 2600 remained a mainstay of home entertainment and made Atari a very profitable company indeed. But around 1984, the gaming industry ground to a halt. Nobody seemed to want to buy games anymore. What happened?

If you were 10 to 20 years old in 1977, it is likely that you would have heard of the Atari system. Once introduced into the home, kids spread the information through word of mouth and played games with others. It was a cool, fun toy. Their parents, who would have grown up in the early 1940's to 1950's, would have found video games too radical a technology to deal with and would instead resign themselves to buying their children a system for Christmas. In seven to ten years, in the mid-1980's, all of these kids would have been ready to graduate either from high school or college. College students and graduates making it on their own would not necessarily have a television or time to play such games. And the generation before them would have been bored by the current crop of video game systems, finding them either too expensive or mediocre.

Enter Nintendo. In 1985, the NES system was released, featuring a quantum leap in home video gaming technology. This was exactly what the younger generation was looking for, and by 1989

Nintendo had captured 80% of the video game market, with sales more than twice as large as Microsoft's. Today that generation is graduating from high school and college, and a new leap in home entertainment technology is called for; thus great video game battles are occurring today, with players such as 3DO, Sega, and Sony.

A similar pattern occurred in the VR field - from the mid to late 1980's the technology made great progress, and by the end of 1991 the mass media had the idea that something great would happen. But in time investors became scarce, start-up companies folded, and the promise of virtual reality appeared to have been forgotten. What happened?

Although promising at the time, the technology was not ready to be made so public. Early adopters who had grown up in the Atari age as well as older investors were wiser to computers and visualization technology by the early 1990's. The personal computer market was still growing, the more and more computers being introduced into the home. VR was not something readily cheap and available for computer users at the time.

All of changes have reflected changes in generation as well as in business. Today one can see a great many signs showing what is about to take place. Today's young generation is waiting for a quantum leap in technology - something cheap and of high quality. Many video game makers feel that they have the solution, but 1995 is not 1985; something new and wonderful has come onto the scene - the cheap, fast personal computer and the Internet.

Why buy a video game system when I can buy a computer that can rival today's older systems? Many game manufacturers today realize that their machines must provide a superior price-performance ratio in order to keep up to the PC. Others point to their profits and shrug off the matter. If they are complacent in the least, the hardware-based business will go the way of the dinosaur.

In order to stay competitive with the computer, other hardware-based media must differentiate or specialize themselves enough so that a superior advantage is maintained, otherwise they will be easily surpassed by the home computer. The questions that tomorrow's generation will be asking are: Why do I need to buy a TV when my computer will do just as well? Why do I need to buy a separate CD player or stereo system? Why do I need to subscribe to a paper-based magazine or newspaper when I have the online version? If I can find all the books I need over the Internet, why buy them in the bookstore?

This does not mean that the computer will swallow up everything. There will always be a place and a reason for certain media and products; little other than the demand will change. Today's Internet frenzy will even out within some years as millions of people undergo technological adjustment, as it all becomes integrated into the everyday lifestyle. For many it will be a quiet revolution, and will not appear very exciting. The true excitement is going to

happen in the generation in elementary school today, the generation in diapers. In ten to twenty years, if history repeats itself, yet another great quantum leap will occur:

The interfaces that we see today - the Macintosh Finder, or that of America Online, are the environments that millions of people will be using for some years to come. They are the environments that users are learning on, they are becoming their electronic world view. But once everyone becomes wired as a new generation grows up not ever having known a world without the Internet or computers, what then? By 2020 so many fields and technologies will have converged, it is hard to say what will happen. But it will likely cause another shift in our thinking and/or our lifestyle.

We often view change as coming in waves because it is such an appealing metaphor. We often think of browsing online as surfing. Why is this so? There has been a great deal of wave imagery in world culture, but the most timeless depiction comes from Katsushika Hokusai's woodblock print *The Great Wave off Kanagawa*, done in the early 19th century. It shows two long boats, its crews lying prostate under a monstrous wave ready to be submerged. The print appears to say to the viewer, something big is going to happen. Are you ready for it?

In a media-related context, wave and surfboard imagery appears in McLuhan's books in the late 60's. In *Computer Lib* of 1974 Ted Nelson quips: "If computers are the new wave, then monitors are the surfboards." The phrase "new wave" and Alvin Toffler's *The Third Wave* appear in the late 70's to early 1980's. Then, with the rise of Nielsen's rating system and the electronic remote control, the phrase "channel surfing" becomes popular. It is then a very short conceptual leap to talk about "net surfing", no matter how appropriate or inappropriate it may be. It is a visual image and a memorable phrase. What are we going to be surfing next?

Great waves occur periodically; the trick is in knowing where they occur. From a distance they appear as ripples and only prove to be large or small when they get much closer. Knowing that a wave is coming is only half the battle. You have to be in the right place in the right time, and once you're there, you must know how to ride it. I hope we're all ready by the time the next big one comes.

The Disappearance of Computers

Technological changes rarely happen as quickly as we would wish them to (and as others would report), but it is only after things happen that we realize that hardly took a blink of an eye in the scheme of things. Today we are witnessing a combination of evolution and revolution.

One day computers will be ubiquitous; we will take them for granted and they will drop out of sight. This happens to all common media. Neil Postman states the tell-tale signs:

- When we are no longer fascinated or perplexed by its machinery

- When we do not confine it to special rooms
- When we do not doubt the reality of what we experience through it

Our computing devices will take up less space, allowing people to pay more attention to the environment around them. You don't think, "There are radios all around us! They are taking over the planet!". Rather, they have become integrated into the environment: there are radios in cars, bars, and bicycles; there are designer radios for sport and leisure; we use radios in times of emergency as well as for pure entertainment. Speakers are incorporated into buildings and sometimes tastefully integrated into restaurant decor. All of these things will likewise happen to the computer.

The miniaturization and portability of computers will make things both cleaner and messier. I mean messier in the sense that they will have the capability to distract no matter where we go, and I mean cleaner in the sense that less materials will be used in their creation; they will have less physical requirements. Instead of constructing special-purpose machine rooms and reinforced computer desks and furniture, architects will be able to concentrate more on the humane side of designing environments rather than having to deal with real-world constraints brought on by various equipment.

But despite the benefits that miniaturization can bring, there will be more opportunities than ever for our environments to distract us as well as more opportunities for us to distract others. Many people today make calls from and answer their cellular phones when walking down the street, in a restaurant, and while driving. Although convenient, we will be carrying much more powerful devices in the future that we will have to maintain, listen to, and communicate through. How often will we be doing this? Shouldn't we all receive an etiquette guide when we purchase these devices? Or will they become such a common part of daily life that we won't care or notice?

When we see computers as simply being extensions of ourselves, as things we make, old phrases like "human-machine symbiosis" become redundant. We are not necessarily the machine, but the machine is definitely us. We make machines in our own image. When we realize this, the machine as a concept will disappear. Perhaps things will change when computers with no preconception of humans or human thought are able to wholly design machines themselves; only then will something truly foreign come into being.

But if machines are us, if they are supposed to work for us and with us, why do we fight them? One reason is because people cannot make them well enough, so that they do not do what people want. People vary widely and everyone has their own view of the world. When we make machines that can adapt to each person, people will find that the way they operate machines will be a reflection of themselves. Another reason why some people hate

machines is because people make machines that don't always work quite right. There will always be a little human error somewhere, and we can always hope to do better.

We distrust machines because other people create and use them in ways that others do not like. This is a human and societal problem, and it takes knowledge and understanding to work out. Those who would design machines must realize that people will use them as extensions of their own bodies, that they will be a part of living. Those who would use machines should reject and throw away or improve what they find hard to use. Technology is no longer something to fight against; it is nothing more than a human artifact that we can use to improve ourselves if we wish.

In the same way that our bodies are constantly changing, moving, thinking, and working, computers will also be on all the time, constantly listening for new information and sending it out. Computers must be made to withstand around-the-clock operation, and thus must be made energy and information efficient. Energy efficient computers will sleep when not used and consume little power; information efficient computers may wake only when there is something to be processed and may parcel out its computing power over the net, spreading its resources to those who need it.

When computers become part of our environment, when they are everywhere and anywhere, we will not care how we are connected, how we get the signal, as long as we get it. Whether we are joined via satellite, cable, phone lines, cellular, fiber, microwave or any number of other means is unimportant. No longer will we care about such things as the "phone company" or the "cable company", there will simply be the "signal provider". The provider of the signal, wherever I go in the world.

Cable companies, signal providers must make equal speed two-way. The digital generation will not stand for anything less. Non-equal speed forces us to become weaker providers and assumes that we have nothing to create or say. For those who are not interested in providing information, unequal access is adequate, but for the children growing up now, they will need it and will demand it in the years to come. And we must make sure our devices can receive what we are able to send out.

The amount of data our devices can hold will depend upon system bus speeds, the speed of interface cards and hardware, the speed of the network and the amount of storage connected to them. My storage does not have to be in the same place as my display - thus I may pay an access provider for a number of gigabytes to store movies for later viewing on my Newton. If I don't wish to store the movie, I may simply receive and display the data as it is streamed to me, although it would be better if I could keep it around or be able to view it a certain number of times.

When computers disappear, pointers and references in our environment will be highly important; they may very well help keep our daily lives together. For a while in film's early days, personal

credits were not mentioned in movies; the very concept of the movie star did not arise until 1910. Today we can assign credits and original references to every digital work that we make, but we do not always do so today. In some cases credit or comment areas are provided in different formats; in the case of bitmap images, references and credits can be encoded and sewn into the picture itself undetectably. Credits personalize our media; when we are able to talk to the authors, the actors, and get background information, we can get a much better sense that humans exist behind what we experience.

When we use computers, we should not see them as inanimate objects to fight with; we should instead view them as devices through which we can see humanity all around us. Computers are not simply a tool, any more than a Japanese garden is a tool. Computers reflect society; they give us something to experience and reflect upon.

Electric Speed

Fred Brooks told his readers to “plan to throw one away; you will anyhow”. We should always plan to throw one away. In today’s rapidly accelerating technological environment, we must learn to reinvent ourselves often and learn new ways of doing things. We can experience more and use tools to help us learn and remember.

Things are changing more quickly: as we use electronic communications and digital media, we will all be accelerated. In cyberspace new forms of communication will be born and destroyed daily in an ecological swirl. We will experience such things as media blitzes, instant fame, overnight lawsuits, and patent battles. Tomorrow’s revolutions can be instant, and their effects will be seen in government, in classrooms, and on the street. Just as British punks in the late 1970’s preferred to record and distribute singles instead of long-playing albums in order to get their messages across as fast as possible, tomorrow’s Internet generation will trade and manipulate quick media portions to effect change.

Internally, we think instantly, with no lag time between the suggestion of a thought and the thought itself. As our technology naturally approaches the speed of thought, we will begin to make more and more demands of it, expecting it to close the gap. Just as people have become impatient with television after the rise of the remote control, so are they tired of waiting for Web pages to arrive. It can be frustrating, but we must learn to be patient.

Today names come and go like yesterday’s news. They are temporary pointers to objects, like handles for MOO participants. Names are only so useful in that one understands what they point to and that others and the namer share this understanding. One culture believed that in the beginning, nothing had a name, so nothing could exist. But once something was named, the thing became real. In the future the labels of ideas and objects will come and go; their very existence will be transitory, instant. Just as we have instant noodles today, cyberspace will embody instant culture.


When a new domain is issued every time a new movie comes out or a new food product is introduced, what is going to happen to the name space? What if a movie takes a name that a corporation or school needs? In practice people rarely think of the global nature of the Internet. Today, assigning a domain name for a consumer product is like reserving a phone number so that nobody else in the world can use the same sequence of digits while dialing. Hopefully as the practice of finding places based solely on named IP address decreases this problem will be alleviated.

As culture and technology accelerates, we approach what Vernor Vinge calls the singularity, when change builds upon change so quickly that society one day will be completely different from the previous day, then an hour will be different from the previous hour, and at some moment humanity will become something we cannot fathom. Will this happen? Is there a limiting factor?

Acceleration does not have to destroy traditional cultures or leave poorer economies behind. At this moment, everything is globalizing into bits and media pieces and knowledge workers, but only the best world economies have them. The lesser economies have a more consistent culture, but are lower on the technology scale. It should be possible to hold on to a traditional culture while embracing cyberspace. If it is properly understood and handled, technology can even reinforce culture, but it is a slippery beast.

High technology is getting relatively cheaper so that low-tech countries can increasingly afford to grab channels into the information society, and a good, fast two-way channel becomes all that counts. So while India has great amounts of satellite dishes,

allowing the world's media to rain down on them, they are able to give as good as they get. When their next generation grows up, watch out.



72: The progress of media.

As the world population continues to explode, progress will march merrily on. In the chart above, look at the period you were born in and the things that have been created in your lifetime. How have they affected you? Do they reflect previous technology? Do they reflect the rising population? How can these creations help us deal with more people and more information?

Dangerous Nonsense

As things stand now, the geniuses of computer technology will give us Star Wars, and tell us that is the answer to nuclear war. They will give us artificial intelligence, and tell us that this is the way to self-knowledge. They will give us instantaneous global communication, and tell us this is the way to mutual understanding. They will give us Virtual Reality and tell us this is the answer to spiritual poverty. But that is only the way of the technician, the fact-monger, the information junkie, and the technological idiot.

Neil Postman

Informing Ourselves to Death
Speech given to the German Informatics Society, 1990

Whether you agree with Neil Postman or not, we must use the Internet for discourse and free thought, as a resource for finding and preserving knowledge, for promoting learning, for positive growth and understanding. None of the technology we create is an ultimate answer to anything; humans must be part of the solu-

tion. And in an age when information blows around us like a tornado, we must have solid ground to stand on. We must invest in people, not just technology.

In time the net will serve as a gigantic scratch pad for our memories - this is why it is important to ensure that it does not degenerate into a morass of lies and what Postman calls "dangerous nonsense".

What has made culture what it is is our ability to remember and transmit knowledge - from the oral traditions of ancient cultures onwards. Today we have technology that can serve as our memory, but it is crucial that we learn how to evaluate, analyze, compare, and learn from experience. Otherwise we will become incapable of any real understanding.

The written word changed the way people thought, putting emphasis on the linear. In some ways, we have been entrapped by writing. However, the online world can be completely free - thoughts may flow hapazardly and freely. This is both good and bad.

Media like USA Today and MTV are very good at chewing information up into small, easily digestible bits. This is a good strategy because it makes information easier to experience. Albums often have spaces between songs and paragraphs allow us to digest the written word. However, there are obvious limits to this technique. I cannot go into any amount of depth; I may end up trivializing the information or displaying it out of context. This is what infotainment does, and we must be careful to avoid this online.

We are at a time when we can make and view media without beginnings or ends. As long as we are able to see things in their original contexts, or in larger contexts, and know when something is out of its original context, and as long as we do not allow ourselves to become complacent, there is little problem.

Our machines will be able to make choices for us, find and choose the things that we like, listen to our preferences. The key is in being able to provide feedback, to evaluate our decisions and our machines' decisions. When we start passively accepting the choices other people and things make for us, that can be dangerous. In cyberspace we must exercise our freedom of media, and our freedom of choice. We must always be able to find the unexpected, to be able to be surprised. This is one of the places that creativity comes from.

Advertising has evolved from using reason to using distraction to order to achieve its goals. In a medium such as television, distraction is natural - everything on TV is basically a distraction, it's all trivia. But in an online environment in which the mind is actively pursuing a goal, such as acquiring new data and analyzing information, distraction is one of the most annoying and damaging things that can appear. We don't see commercials in libraries or advertisements in textbooks.

Learning should be fun, and it should be as fun as engaging entertainment. We learn in anything we do, consciously and unconsciously, and even mindless entertainment can help us learn a little. The danger lies in presenting education as entertainment for entertainment's sake. Students must be able to discern the difference between what they feel is useful and useless information. When they cannot, we will have gone too far and must rethink things over.

We will have failed when we can no longer tell the difference between advertisements and textbooks. It is now possible online for reasoned discourse to become distraction, and we see this happening now as once focused, informative newsgroups devolve into blurry piles of interruption. So many people have failed to understand why the traditional Internet community frowns upon irrelevant mass postings - they do not understand that people use the Internet as their library, as their memory, as their knowledge resource. When data can be manipulated, separated or structured in such a way that people can easily find and use what they are looking for, the potential for true learning increases. If data is corrupted by irrelevancies - as advertisements in textbooks - it loses its worth, and then everyone loses. The Internet must always be able to give people what they need, not just what they want.

How do we fix this? Education! This does not mean that the blind intake of facts will solve things, but rather a combination of knowledge, wisdom, experience, and self-realization will be necessary.

- **Learn about learning and evaluation.** Because our tools are changing so rapidly, perhaps tomorrow's business will put more emphasis on evaluating as a skill, and learning how to learn. Learn to evaluate different world views and see the overall picture.
- **Learn about media.** Demythologize media. Don't use the computer to control education; use education to control the computer.
- **Realize your potential.** You have one of the most versatile objects known to humankind. What do you want to do with it?

Real Reality

The question that people most frequently ask about virtual environments is: Will real reality still be worth it? The answer is it can be, but we cannot let it become the natural substitute for all experience. How we view things as being real is subjective in the long run, and high resolution alone may not necessarily count as true reality to somebody else.

While we race to design systems that are as real as possible, we should also ask ourselves: Is real reality worth it now? Do we make it worth living in now? Can we make reality better with the application of virtual environments?

People once thought of films as reality - they would jerk their heads when objects on the screen approached them, or believe that actors were actually getting bigger when the camera zoomed in. But in time people were able to recognize the difference, just as people do to some extent with television. As media takes on more of the visual qualities of reality, many will not be able to tell the difference at first, but in time they will be sensitive to the difference.

We can see this kind of sensitivity now in the media-savvy; people who have played video games, young people, are very sensitive to subtle differences in the feel and behavior of immersive environments. Ultimately, we must always be able to tell the difference between reality and what we have created, to be able to say, "Computer, stop simulation," or hit a switch, because if we truly could not tell the difference, why would we stop using the virtual thing? How easy would it be to mislead or control others?

The cyberspace that many developers wax eloquently over today is not the crude, blocky worlds we experience in daily computer life. They refer to as real as it gets, utilizing the bandwidth of human vision, estimated to be about 6 times 10 to the 12th megabits per second, or about four million times the bandwidth of the average T-1 Internet connection. Every month the Internet shuffles more than the data equivalent of the printed collection of the Library of Congress in North America alone. Human vision gives us a bandwidth of 60,000 times that per *second*.

It's not really until then, they say, that virtual environments can be all they want it to be, all they say it can be. For now, they have been saying, "trust us," and we have. For every person that experiences the Internet for the first time, that is one more individual trusting the best and brightest to come up with reality better than the real thing. Today with easily available (but expensive) commercial equipment it is possible to get nearly one million textured polygons per second, two million at the high end. If real reality is at least 50 to 80 million polygons per second, and one assumes that the number crunching power doubles every two years, then we will certainly get reality working in the labs within two decades. If not, then people can certainly wait a few more years before it happens. In any case, it is most likely going to happen before my peers are barely finished with their mid-life crises.

There is no way that this medium can be introduced to the human race unless it understands the media thoroughly and is prepared to deal with it. It has been said that it is morally bankrupt to give people technology without education. It is certainly not the wisest thing to do. So in this age, as technology grows and changes all around us at speeds the world has never seen before, people must be able to understand what's happening. The world needs an operator's manual for media before it destroys itself by pushing the wrong buttons.

Be Digital As a sign of how slowly ideas have propagated in the past, it is often pointed out that Gregor Mendel's laws of genetics were not discovered until over three decades after they were written. The was partly due to the fact that few copies of his published work existed. To add perspective to the topic of change, Alan Kay shows two pictures of Stone Age hand axes. One has a sharp edge around its entire perimeter; the other has a handle carved into the end. Other than that one difference, they look quite similar. The time difference between the two creations? About 200,000 years.

It may very well take decades to everything to be online. Until then, if we wish to explore an idea, we may have to go to a research library, gain access, search for the physical format, and manually scan for the information (if it is available). Progress has moved at this rate for thousands of years.

Today if any kind of research is to be performed, it must be put online, on the Internet. If the research is meant to be disseminated, then putting it in an obscure volume and releasing it to a small group of people is performing a terrible injustice. Today there is far more research in many fields than anyone can possibly read manually in a lifetime. But stored online, inventions, studies, and results can be compared and cross-referenced in an instant. Progress is summarized, people can view more material, and thus more possibilities for creation arise.

Some have argued that putting up information as traditional multimedia would leave little to the imagination are themselves not thinking correctly about the possibilities: to ensure that much is left to the imagination, why not just put up something as straight text? Just because something is on the Web does not mean that it must incorporate graphics, movies, and sounds. Alice in Wonderland is online - does that leave any less to the imagination than the physical book? The 1980's BBS hacker/cracker culture thrived on text-only files; data is data.

When a reader receives a book, it is as if they have received a brand new coloring book. The author has provided an outline, a story, but it is up to the reader to fill in the colors with their own imagination. To some extent, today's multimedia is the equivalent of coloring books that are already filled in. The audience is presented with little in the way of choice.

As we create virtual environments, we will temporarily have to fill in some of the colors - in many cases, users will not be able to change what the original author makes. But as technology allows the capability, in time users will be able to use their own imaginations to any extent they like. Environments then become coloring books in which one can add, delete, and change the outlines. Multimedia can become like books - only suggestions, frameworks, but it will take a while to evolve. We are still in the Carved in Stone Age.

For those who would refrain from digitizing information for fear of making it public: if one wishes to disseminate something digitally, it does not have to be sent to the entire world. Mailing lists, private networks, encryption, and portable media allow people to keep their digital media to themselves or certain individuals.

The U.S. Constitution calls for the establishment of patent and copyright methods "to promote the progress of science and useful arts". Whether or not one believes in the U.S. patent and copyrighting methods, progress can be made much swifter with fast, easy access to information. Think about mainframes and centralized databases. Now think about clients and servers, decentralization. Imagine the patent or copyright office, a centralized repository of claims to ownership. It can take months and even years to file, and even then it is sometimes done incorrectly and with damaging results. Now imagine a decentralized copyright scheme, where one can find the owner, and pay the owner directly within seconds. Centralization is a bottleneck, because today there are so many people making things, and so much media. It is a bottleneck in the progression of science and the useful arts.

Digitize your information. Make it available so that it can be found easily. Allow as many people to reach it as possible. Do not send it to people when you feel like it; allow people to get it for themselves at any time. Information by itself will not make progress any faster; it is its availability that counts, because once it can be found, it can be applied.

Some say that we are in the post-information age - we already have enough information! The truth is, humans can never have enough information or stop creating it. Think of all the new books and movies and artworks that are created every day in the world. If I can have instant access to any of this, isn't this much better? It is strange that people feel that there is too much information - if they don't want it, they can always ignore it.

What people are really concerned about is the amount of information they need to know and dig through in order to operate in modern society. This is a societal problem and a design problem: make things simpler and easier to use. Streamline processes. Distribute the burden. Make things easy to find.

We haven't yet passed the information age, as the information age can be split into at least two parts: acquisition and assimilation. We are nowhere near assimilating yet.

Online Freedom

Access to computers should be unlimited and total.
All information should be free.
Mistrust authority - promote decentralization.
Hackers should be judged by their hacking, not bogus criteria such as degrees, age, race, or position.
You can create beauty and art on a computer.
Computers can change your life for the better.

Steven Levy

The classic hacker ethic as reported in Hackers

And to the 25,000 people who want to call this Senator a barbarian, I simply say that, evidently, they are so selfish - at least their actions are so selfish, that they simply say: We do not want to give up anything. We want to be able to see what we want to see where we want to see it, any time we want to see it.

Senator James Exon (D-Nebraska)

From the Congressional Record of June 12, 1995

Cyberspace - live free or die. The United States supports the freedom of speech; for decades those that have been entranced with the Internet, the true hackers, have wished this to be a free world. But today legislation is introduced that would outlaw speech online that in a printed format would be perfectly legal. It would ask that non-commercial users be subject to state censorship and restriction. It would call for FCC jurisdiction over standards for Internet services, email, news, and the World-Wide Web.

In effect, I could examine and read all manner of uncensored international material at the click of a mouse, but I would not be able to create it on the World-Wide Web. The Internet is *world-wide*, not a closed network. Authoring is as essential an activity as viewing. With new forms of interactive media, they amount to nearly the same act.

Today a 17-year-old can walk into a bookstore and buy Playboy, or send money and get a subscription. I mail-order instructions on making bombs or do a little field work at any research library. The fact that I can get to the information does not change; the problem becomes one of temptation, ease of access to particular information. To help alleviate this problem, it is the cause that must be fixed, not the symptom: let people censor themselves, let parents decide what is best, help make a society worth living in rather than penalize those who would rebel against it.

Information that is made illegal then becomes similar to drugs and guns, with the exception that it is a lot easier to smuggle and use information than drugs or guns. Legislation can take things away, but we live in such a networked, global world, there are always places to find them. Children should be taught the facts about illegal materials so that they can become familiar with them, recognize their dangers, and evaluate them in an objective manner.

5.5 What Does Cyberspace Need?

In *The Mythical Man-Month*, Frederick Brooks analyzed the Tower of Babel project as an engineering project in a humorous yet informative way. He proposed that the core reason of Babel's failure was the lack of communication, and following that, the lack of organization. Perhaps it is time to take a look at the Internet as an engineering project. Does it have:

- *A clear mission?* Its original mission was to maintain military communications in the event of a nuclear war. However, it is now being used by the civilian population for no clear purpose but to act as a form of mass media.
- *Manpower?* There seems to be enough people working on the project, but they can barely keep up with demand.
- *Materials?* There is currently no shortage of silicon, glass, plastic, or wire. Enough content is available to keep people interested, at least for the time being.
- *Enough time?* The project is open-ended; there is no completion time.
- *Adequate technology?* Yes, the decentralization concept has worked well and progress is being made so that heavy and growing traffic can be accommodated with a minimum amount of problems in the near future.
- *Adequate communication?* There are a great many subdepartments working in different areas, so communication has been difficult. Much redundant work is being done in different areas.
- *A clear organization?* Not really, no.
- *Good documentation?* New users are not supplied with manuals, and hundreds of operating systems are implemented, so it is very confusing for people who are learning. Often they have no idea why or when to upgrade, and they do not understand the philosophy of the product.

What is needed? Better organization, more coordinated and timely development efforts, widespread access to new developments, and improved user access and education. Public freenets and forums, paid for by access providers. Coordinated open standards efforts.

The new online community needs to be educated about what having an open standard means, so they are aware when others try to force them into using a closed standard. Mike Azzara has written these characteristics of open systems:

- The specification is free or almost free.
- Multiple vendors developed the specification or support it.

- Multiple implementations exist.
- No single vendor can unilaterally decide to change the specification.
- Products are available from a relatively large number of vendors, and other vendors' products support the interface.
- Conformance testing assures users that products meet the agreed-upon standard specification.
- Bilateral testing assures users of interoperability among different vendors' products.

To help create new services and environments, standards organizations should foster interdisciplinary groups in order to discuss future growth and provide the things that users need as well as brainstorm new ways of doing things. This is one idea behind synectics, bringing together diverse people in brainstorming sessions in order to solve problems.

As an example, the World-Wide Web organization can foster communities of designers, protocol developers, and publishers into order to explore future applications. VRML discussion groups may consist of architects, musicians, programmers, and psychologists. Internet standards bodies should bring together people from many professions - sociologists, human interface designers, secretaries, executives - and ask, how can we make things better?

In terms of virtual environment technology, Roy Kalawsky mentions the establishment of virtual environment laboratories that are technology-independent and that allow rapid prototyping. These labs would do such things as investigate human factors issues, establish performance metrics, and demonstrate concepts.

For combat simulation (although these issues can be applied to general research), the Office of Technology Assessment has issued a number of challenges, including the development of:

- High-density, color flat-panel displays
- Fast head tracking
- Wideband networks with low latency
- Multilevel network security
- Automating object and world description for scene generators

Kalawsky notes that more work is needed in:

- The human perception of virtual environments
- Visual programming languages
- Database standards
- Environments with multiple participants

The OTA stresses such factors as validation, standardization, scalability, flexibility, effectiveness, cost-effectiveness, and infrastructure in looking at different solutions - developers would do well to examine each of them in making new creations.

Make all research public, put it in an easily accessible place, advertise the knowledge. In many scientific communities it is difficult to peek in to see what others are doing: the hacker culture started because students tried to enter a world they were never allowed to see. By showing the public what is happening and spreading the research, one then engages the CPU processing time of every person on the planet that comes into contact with that research. The world's population is the biggest massively parallel computer we may ever know, and the Internet is its data bus. Language is a virus; now media is itself a virus. Garbage in, garbage out.

When people are cut off from what's happening, they assume, and spread misinformation, and become ignorant. Thus Alexandria burned. The idea is to spread knowledge as far and widely and freely as possible, and to not let it become trivialized or twisted.

Gibson proposed free long-distance access and software for teachers over two years ago. Today this can be modified into: give teachers and students free Internet access! Make an EduNet, where teachers and students from all over can collaborate and make new technologies to help themselves, where needs can be stated and help can be offered. There are local EduNets around the world, but only in name. What is needed is something much larger in scope.

One way of looking at this new technology is as a gun. Depending on the person that holds it, this gun can shoot either knowledge or bullets. Would we choose to shoot ourselves, committing societal suicide, or would we choose to strive to make ourselves better than we are? Technology never takes sides, people do.

When we all speak different languages so that there is no understanding, we are ruined. When we cannot talk and cooperate, or exchange things freely, we lose the opportunity to create the greatest, largest technical achievement known to humankind. The achievement does not come from discovering new places or things, it comes from discovering ourselves and learning to act as one.

And the Lord said, Behold, the people is one, and they all have one language; and this they begin to do: and now nothing will be restrained from them, which they have imagined to do.

The King James Bible

Genesis 11:6

A History of Cyberspace

Everything is deeply intertwined.

Ted Nelson

Computer Lib/Dream Machines

In researching the ideas in this work, I followed a web of ideas: if an author, book, movie or other media was mentioned as a major influence, I looked at that, and followed its references, and so on. Perhaps not surprisingly, there are a great deal of common ideas and works that people in many converging fields have mentioned. I have noted many of these down in this timeline, a compendium of events in the last 500 years that have in some fashion contributed to the progress of cyberspace as a concept.

The selection of the events in the timeline are of my own choosing and reflects my subjective opinion of the relative importance of things. When I found different dates for the same event, I chose the earlier date.

This history merges five separate timelines:

- **World** - this timeline includes events in the popular media, including the release of movies, books, and other works; general technological facts; and events that encompass more than one timeline or did not seem appropriate for the other timelines.
- **Internet** - this timeline encompasses the history of the Internet.
- **VR** - this timeline encompasses the history of virtual environment technology, including VRML.
- **Hypermedia** - this timeline includes events related to the development of hypertext and hypermedia systems, including the World-Wide Web in particular.
- **Computers** - this timeline includes hardware achievements in the computer field as well as influential works.

There are a number of other timelines and histories that would have been included but for the lack of space, including the history of drugs, the late 1970's music revolution (including punk), impressionism, surrealism, futurism, situationists, the Bauhaus, beatniks, hippies, hackers, crackers, and today's twentysomething generation.

These timelines are presented vertically in parallel so that events in different fields can be compared and followed. One can see how many developments took place over a period of years; one thing often leads to other. How did the current environment of a time foster the development of a new gadget or idea? What media influenced future developers and artists?

One should note how many developments are generational: look at one event, then move forward twenty years or so and see if there is a relation. For instance, science fiction has often been an influential factor. Just as NASA workers in the late 60's grew up on Buck Rogers in the 1930's on, kids who saw Star Trek in the 60's are now grown up and changing the world.

After enough of this you will be able to see the present and picture what is going to happen twenty or so years from now. What you see will be a reflection of what we both believe.

All the people who have short-range goals will be forgotten.

Marvin Minsky

What would I recommend to a young visionary today? Very straightforward. Learn to deal with short-term goals, and not delegate.

Ted Nelson

June 1995

As you examine the timelines, you may notice that just as in real life, there are forgotten visionaries, dashed dreams, surprising twists, and wild successes. In the middle of it all I have delineated a number of ages, certain years that have encompassed movements and revolutions. In addition to the well-known eras, I have included a number of new ones:

1957 - 1965 The Age of the True Hackers

1969 - 1977 The Age of the Hardware Hackers

1977 - 1987 The Age of the Game Hackers

These three labels are from Steven Levy's *Hackers* and do not necessarily reflect how good or bad hackers were in these years but reflect a changing lifestyle and outlook. I chose 1957 as a beginning year, being the year ARPA started and just a few years after the introduction of the first UNIVACs and slightly before the development of FORTRAN. It is around this time that the classic hacker culture began, with punched tape and cards.

There was no definite end of the True Hackers but instead a beginning of a new era begun by the Free Speech Movement in late 1964. In 1969, with the creation of ARPANET and the release of Intel's 1KB RAM chip, the age of the Hardware Hackers was initiated. This age saw the development of home-built personal computer kits, the rise of Apple Computer, and the beginnings of Microsoft, and ends when Space Invaders is released in 1977, putting video games into the public eye.

1972 - 1984 The Golden Age of Video Games

The great age of video games begins with Nolan Bushnell's Computer Space game and ends with the great video game crash of 1984. The Game Hacker era ended around the late 1980's with the last-ditch computer efforts of Atari around 1987.

1984 - 1992 The Golden Age of VR

With the publication of *Neuromancer* and *True Names* and the creation of the first affordable HMD at NASA, 1984 was a seminal year for VR. This age saw the rise and fall of many companies such as VPL, which broke much new ground in the field. It ends after VPL's dissolution, a year after the VR craze hit a peak.

1984 - 1991 Cyberpunk

Neuromancer began this movement in part which began in spirit with the distribution of Bruce Sterling's informal *Cheap Truth?* newsletter in 1983. The core movement may have ended in 1988 with the last of Gibson's trilogy *Mona Lisa Overdrive* but a good number of events occurred until 1991 when the *Beyond Cyberpunk!* stack was released as well as a number of subsequent books on the topic. The peak was probably reached in 1989 to 1990 with *Tetsuo: The Iron Man*, *Mondo 2000*, and the debut of *Nine Inch Nails*. In any case, it was dead well before it made the front cover of *Time* and has since evolved into a mainstream trickle.

1985 - 1990 The Age of the BBS Hackers

This age was typified by the teenage male cracking from his parents' home with cheap equipment and homemade devices, using BBSes as central points of operation. After the rise of *2600* and the *Legion of Doom* in 1984, 1985 saw the first police sting BBSes: these kids were now considered a threat. The core movement ended around 1990 with the outrageous Sun Devil raids; after that, hackers in general became considerably more mainstream.

1993 - ? The Age of the Internet

There is no question that 1993 saw the beginnings of the great net migration, with the initial release of Mosaic, the First Internet World Conference, and the creation of the InterNIC. From this period on, traffic rates and all manner of growth statistics have been astounding. Where will it end?

? - ? The Information Age

Only time will tell when the Information Age truly began - was it with the release of the first commercial mainframes in the 50's, was it with the subsequent rise of the data processing industry, or was it with the creation of the ARPANET? Most likely people will simply say that it started in the second half of the twentieth century. For the last twenty years, a definite majority of Americans have felt that computers can improve their lives, and they are proving their faith by using and investing in all manner of equipment. When this age ends and the next begins, nobody knows.

? - ? The Post-Information Age

Some have talked about a Post-Information Age, when we have stopped acquiring information (and perhaps converting legacy data) and have begun to use it in daily life. Will this occur before the millennia, or will it take much longer?

This history is far from being complete; unfortunately there are to date no adequate tools to properly create structures like this to display cross references and relations. Perhaps one day this history will be online with references to original media and connected to other threads; until then it must remain confined to a paper-based medium in this format.

	World	Internet	VR	Hypermedia	Computers
1453	<ul style="list-style-type: none"> • Gutenberg prints the Bible at Mainz, Germany 				
1400's	<ul style="list-style-type: none"> • Giotto creates perspective 				
1536	<ul style="list-style-type: none"> • The <i>Venetian Gazzetta</i> is the first European newspaper 				
1623					<ul style="list-style-type: none"> • Wilhelm Schickard creates the first mechanical calculator
1755	<ul style="list-style-type: none"> • Dr. Johnson's <i>A Dictionary of the English Language</i> is published 				
1764	<ul style="list-style-type: none"> • James Watt invents the condenser 				
1768	<ul style="list-style-type: none"> • The <i>Encyclopaedia Britannica</i> is published 				
1770 to 1870 - The Industrial Revolution					
1775 to 1873 - The American Revolution					
1775	<ul style="list-style-type: none"> • James Watt perfects the condenser 				
1776 Jan 10	<ul style="list-style-type: none"> • Thomas Paine's <i>Common Sense</i> is published 				
1777					<ul style="list-style-type: none"> • The Stanhope Logic Demonstrator is invented
1784	<ul style="list-style-type: none"> • The first mail sent by coaches goes from London to Bristol 				
1788					<ul style="list-style-type: none"> • Robert Barker creates a 360-degree panoramic scene of Edinburgh
1791 Feb 22	<ul style="list-style-type: none"> • The First Amendment of the U.S. Constitution guarantees freedom of the press • Thomas Paine's <i>The Rights of Man</i> is published 				
1793	<ul style="list-style-type: none"> • The Mechanical Semaphore Network begins 				<ul style="list-style-type: none"> • The Stanhope Arithmetic Machine is invented
1794	<ul style="list-style-type: none"> • Thomas Paine's <i>The Age of Reason</i> is published 				
1800	<ul style="list-style-type: none"> • Conte Alessandro Volta produces steady electrical flow from his battery 				
1802	<ul style="list-style-type: none"> • Biology as a science begins • Atomic theory is invented 				
1804					<ul style="list-style-type: none"> • Joseph-Marie Jacquard creates the punch-card loom
1811	<ul style="list-style-type: none"> • The Luddite movement destroys industrial machines in Northern England 				
1814	<ul style="list-style-type: none"> • George Stephenson invents the first practical steam engine near Newcastle, England 				
1818	<ul style="list-style-type: none"> • Mary W. Shelley writes <i>Frankenstein</i> 				
1820's	<ul style="list-style-type: none"> • The term "slum" is coined 				
1822					<ul style="list-style-type: none"> • Charles Babbage develops his Difference Engine prototype
1825	<ul style="list-style-type: none"> • The first railway opens in London 				
1828	<ul style="list-style-type: none"> • The first U.S. railroad is built in Baltimore and Ohio 				
1830 to 1895 - The Scientific Revolution					
1833					<ul style="list-style-type: none"> • Wheatstone invents the stereoscope

	World	Internet	VR	Hypermedia	Computers
1834					• Charles Babbage conceives of the Analytical Engine
1836	• Wheatstone invents the electric telegraph				
	• Samuel F.B. Morse invents Morse Code and has a working telegraph prototype				
1839	• Daguerre invents the camera and takes the first photograph				
1844	• 500 stations of Claude Chappe's telegraph network connect 29 French cities				
	• Telegraph line joins Washington, D.C. and Baltimore; first words transmitted: "What hath God wrought"				
1845	• The Scientific American is established				
1851	• Reuters news service started				
1857	• The Atlantic Monthly is founded				
1858	• Transatlantic telegraph cable laid but soon breaks				
1859	• Charles Darwin's <i>On the Origin of the Species by Natural Selection</i> is published				
1860 to 1900 - Impressionism					
1861	• The U.S. introduces the passport system				
1863	• The U.S. Congress establishes free city mail delivery				
1865	• Electromagnetic theory of light is proposed by James Clerk Maxwell				
	• The Atlantic Cable is successfully laid				
	• The U.S. Secret Service is founded				
1868	• Christopher Sholes invents a typewriter that is as fast as handwriting				
	• Painters begin to paint in the Impressionist style				
1870	• Jules Verne's <i>Twenty Thousand Leagues Under the Sea</i> is published				
1871	• Charles L. Dodgson's <i>Through the Looking-Glass</i> is published				
	• Charles Darwin's <i>The Descent of Man</i> is published				
	• Cable to Australia is laid				
1872	• Eadweard Muybridge begins pioneering work in motion photography				
	• The duplex telegraph is perfected by Thomas Alva Edison				
	• Jules Verne's <i>Around the World in 80 Days</i> is published				
1874	• The first Impressionist exhibition is held in Paris				
1876	• Alexander Graham Bell invents the telephone				
1877	• Thomas Alva Edison invents the "sound writer", a cylindrical phonograph				
	• Thomas Alva Edison invents the phonograph				
	• The first public telephones appear				
1878	• First teenage males flung off phone system by enraged authorities				
	• David Hughes invents the microphone				

	World	Internet	VR	Hypermedia	Computers
1879	<ul style="list-style-type: none"> London installs its first telephone exchange 				
1884	<ul style="list-style-type: none"> The first regular comic strip, <i>Ally Sloper's Half Holiday</i>, begins in the UK George Eastman invents sensitized roll film Edwin Abbott's <i>Flatland</i> is published 				
1886	<ul style="list-style-type: none"> The Impressionists hold their eighth and last exhibition in Paris First International Copyright Convention is signed Over 1,400 strikes have occurred this year 				
May	<ul style="list-style-type: none"> Nationwide strikes for an 8-hour day 				
1888	<ul style="list-style-type: none"> George Eastman perfects the "Kodak" box camera Nikola Tesla constructs the A.C. electric motor 				
1889					<ul style="list-style-type: none"> A punch card system prototype is created by Herman Hollerith
1892	<ul style="list-style-type: none"> The San Francisco Examiner runs the first newspaper comic strip The first automatic telephone switchboard comes into use 				
1894			<ul style="list-style-type: none"> David Brewster refines the stereoscope 		
1895	<ul style="list-style-type: none"> Auguste and Louis Lumiere show the first moving picture in Paris, <i>La sortie des ouvriers de l'usine Lumiere</i>, call their system cinematographe Sir William Crookes invents the cathode ray tube Marconi invents the radio (wireless) telegraph H.G. Wells' <i>The Time Machine</i> is published 				
1898	<ul style="list-style-type: none"> A Danish engineer invents magnetic sound recording 				
1900 to 1929 - The Silent Film Era					
1900	<ul style="list-style-type: none"> R.A. Fessenden transmits speech by radio waves 				
1901	<ul style="list-style-type: none"> Guglielmo Marconi transmits telegraphic radio messages from Cornwall across the Atlantic 				
1902	<ul style="list-style-type: none"> Enrico Caruso makes his first phonograph recording Melies creates the movie <i>A Trip to the Moon</i> 				
1904	<ul style="list-style-type: none"> The first telegraphic transmission of photographs The first radio transmission of music in Austria 				<ul style="list-style-type: none"> Sir John Fleming invents the electronic vacuum tube
1906	<ul style="list-style-type: none"> The Haloid Corporation is founded (later Xerox) 				
1900 to 1929 - Cubism					
1907					<ul style="list-style-type: none"> Lee DeForest invents improved vacuum tube
1909	<ul style="list-style-type: none"> E.M. Forster's short story about a machine-bound society, <i>The Machine Stops</i>, is published 				
Feb 20	<ul style="list-style-type: none"> The word "futurism" is coined 				

	World	Internet	VR	Hypermedia	Computers
1913	<ul style="list-style-type: none"> • 500,000 telephones in U.S., 122,000 in Great Britain, 92,000 in Paris • Marcel Duchamp exhibits <i>Nude Descending a Staircase</i> 				
1914 to 1918 - World War I					
1915	<ul style="list-style-type: none"> • The first transcontinental phone call is made from Alexander Graham Bell to Thomas A. Watson • D.W. Griffiths' <i>Birth of a Nation</i> establishes the motion picture as art 				
1918	<ul style="list-style-type: none"> • Regular airmail service is established between New York and Washington 				
1919	<ul style="list-style-type: none"> • The Bauhaus is born out of World-War I in Weimar, Germany 				
1920	<ul style="list-style-type: none"> • Morris Fuller Benton reworks Century type as Century Schoolbook • The First International Dada Festival is held in Berlin, Germany • Karel Capek's play <i>R.U.R.</i> coins the word <i>robot</i>, later robot 				
1922	<ul style="list-style-type: none"> • Phonofilm containing both voice and picture invented by Lee DeForest • Radiophotography invented by Arthur Korn 				
1923	<ul style="list-style-type: none"> • Vladimir Zworykin invents iconoscope to transmit and kinescope to receive TV messages • Charles F. Jenkins transmits TV pictures from Washington to Philadelphia 				
1924	<ul style="list-style-type: none"> • Andre Breton's <i>The Surrealist Manifesto</i> is published • LeCorbusier's <i>The City of Tomorrow</i> debuts • The first singing commercial is broadcast 				
1925 to 1950 - The Golden Age of Radio					
1926	<ul style="list-style-type: none"> • Warner Brothers releases <i>Don Juan</i>, the first film with music synchronized on disk • <i>Metropolis</i> is released, directed by Fritz Lang • Kodak produces the first 16mm movie film 				
1927	<ul style="list-style-type: none"> • Warner Brothers produces <i>The Jazz Singer</i>, a silent film with singing sequences • I.P. Pavlov's <i>Conditioned Responses</i> is published • Transatlantic telephone service is inaugurated 				
1928	<ul style="list-style-type: none"> • <i>Buck Rogers</i> originates in Philip Francis Nowlan's "Armageddon 2149 A.D." in <i>Amazing Stories</i> • The first color motion pictures are exhibited by George Eastman • The first scheduled television broadcasts occur in New York • J.L. Baird demonstrates color television 				
1929	<ul style="list-style-type: none"> • A comic adaptation of <i>Buck Rogers</i> by Dick Calkins runs until June 1967 • The first true all-talking motion picture appears: <i>Lights of New York</i> • Philo Farnsworth develops the electronic television • Rene Magritte paints <i>The Betrayal of Images</i> 				

	World	Internet	VR	Hypermedia	Computers
1929			• The Link Trainer is patented by Edwin A. Link		
1930	• A Britain to Germany picture telegraphy service opens				• Vannevar Bush et. al. make the differential analyzer at MIT
1931			• The Navy buys a fully featured Link Trainer for \$1,500		
1932	• Aldous Huxley's <i>Brave New World</i> is published, describes "feelies"				
1933	• Bell Telephone Labs develop stereophonic sound reproduction system • Hitler closes the Bauhaus due to Germany's new cultural policies • Henry Beck creates the London Underground subway map				
1934			• The Army buys six Model A Link Trainers for \$3,400 each		
1935	• Ambrose Bierce's <i>The Devil's Dictionary</i> is published in full • Print media captures 53% of U.S. advertising dollars • Technicolor films are invented			• Vannevar Bush conceives his first ideas on MEMEX	
1936	• Salvador Dail paints <i>Soft Construction with Boiled Beans: Premonition of Civil War</i> • First scheduled TV transmissions from BBC in London				• Alan Turing first describes the Turing Machine
Jul	• Piet Mondrian paints <i>Composition in Red and Blue</i>				
1937	• Maurits Cornelis Escher creates the first <i>Metamorphosis</i> image				
1938	• Orson Welles' <i>War of the Worlds</i> radio program causes panic • Albert Hofman first synthesizes LSD				
1939 to 1945 - World War II					
1939	• A "futurian" science-fiction group is raided by the Secret Service • Buster Crabbe stars in the first <i>Buck Rogers</i> movie				• Konrad Zuse creates the Z4, a computer that makes use of vacuum tubes • Howard Aiken begins work on first fully automatic large-scale calculator
Apr 30	• Roosevelt speaks at the first scheduled telecast at the New York World's Fair				
Oct					• George R. Stibitz and Samuel B. Williams develop the Complex Number Calculator at Bell Labs • John V. Atanasoff finishes the prototype of the first electronic digital computer, ABC
Late 30's					• Konrad Zuse invents the first electromechanical binary computer, the Z1, in Germany

	World	Internet	VR	Hypermedia	Computers
1940			<ul style="list-style-type: none"> • Robert Heinlein's short story <i>Waldo</i> coins the word 		
1941	<ul style="list-style-type: none"> • A television transmission is broadcast from Empire State Building, New York 				
Aug 1	<ul style="list-style-type: none"> • Isaac Asimov's <i>Foundation</i> series debuts 				
1942	<ul style="list-style-type: none"> • Isaac Asimov's "Runaround" story in <i>Astounding Science Fiction</i> explains the Three Laws of Robotics • The International Surrealist Exhibit, New York • Magnetic recording tape is invented 				
1943					<ul style="list-style-type: none"> • Development of ENIAC is begun, made to calculate artillery range tables
Apr 19	<ul style="list-style-type: none"> • Albert Hofman first synthesizes LSD 				<ul style="list-style-type: none"> • The Colossus Mark I is operational in the UK
Dec					<ul style="list-style-type: none"> • Howard Aiken's Harvard Mark I computer debuts
1944					<ul style="list-style-type: none"> • The Colossus Mark II computer debuts
1945	<ul style="list-style-type: none"> • George Orwell's <i>Animal Farm</i> is published 				<ul style="list-style-type: none"> • The ENIAC performs calculations for hydrogen bomb development and aspects of thermonuclear detonation
Mar					<ul style="list-style-type: none"> • Bell Labs introduces its Computer Model IV
Jul				<ul style="list-style-type: none"> • Vannevar Bush's influential "As We May Think" article is published in the <i>Atlantic Monthly</i>. First published proposal of a hypertext device. 	
1946 to 1991 - The Cold War					
1946	<ul style="list-style-type: none"> • Beginning of the explosion in television sales 				<ul style="list-style-type: none"> • John von Neumann, Herman Goldstine, and Arthur Burks issue influential report on stored-program computers • J. Presper Eckert and John W. Mauchly publicly announce ENIAC at the University of Pennsylvania, 30 tons and 18k vacuum tubes
Feb			<ul style="list-style-type: none"> • Jay Forrester at MIT begins work on Whirlwind computer, partly for digital flight simulation 		
1947			<ul style="list-style-type: none"> • The first teleoperated devices are made 		<ul style="list-style-type: none"> • The transistor is invented • The Association for Computing Machinery is formed

	World	Internet	VR	Hypermedia	Computers
1947					<ul style="list-style-type: none"> • Alan Turing publishes article on intelligent machinery which launches AI
1948	<ul style="list-style-type: none"> • E.E. Doc Smith's <i>Triplanetary</i>, first in the Lensman series, is published • George Orwell's <i>1984</i> (originally titled <i>1948</i>) is published • Polaroid creates instant photography • Dennis Gabor creates holography • The long-play record is invented 				<ul style="list-style-type: none"> • John Bardeen, Walter Brattain, and William Shockley at AT&T Bell Labs announce the transistor • Claude Shannon's <i>The Mathematical Theory of Communication</i> is published in the Bell System Technical Journal
1949			<ul style="list-style-type: none"> • Ole Kirk Christiansen, Lego's founder, begins making plastic interlocking bricks 	<ul style="list-style-type: none"> • Link Corporation wins Air Force contract to develop the first jet simulator for the F-80 	
1940's	<ul style="list-style-type: none"> • The <i>Quiz Kids</i> radio show airs 				
1950	<ul style="list-style-type: none"> • Alan Turing writes the paper <i>Computer Machinery and Intelligence</i>, in which he describes his test • Hermann Zapf designs the Palatino typeface • The military budget reaches \$12 billion • 1.5 million televisions exist in the U.S. • Isaac Asimov's <i>I Robot</i> is published 		<ul style="list-style-type: none"> • Ray Bradbury's <i>The Veldt</i>, one of the first stories to incorporate full VR, is published • The Whirlwind computer goes online at MIT 		
1951	<ul style="list-style-type: none"> • Maurits Cornelis Escher makes his <i>House of Stairs</i> • 15 million televisions exist in the U.S. • The color television is introduced 				<ul style="list-style-type: none"> • J. Presper Eckert and John W. Mauchly develop UNIVAC I, selling 46 machines in the U.S.
1952					
Nov 4					<ul style="list-style-type: none"> • The organization later known as the National Security Agency is born
1953	<ul style="list-style-type: none"> • B.F. Skinner's <i>Science and Human Behavior</i> is published • Alfred Bester's <i>The Demolished Man</i> is published • Robert Heinlein's <i>Revolt in 2100</i> is published • Watson and Crick discover DNA • The Ampex VTR is invented 				
Dec					<ul style="list-style-type: none"> • The UNIVAC 1102 is announced
1954	<ul style="list-style-type: none"> • The U.S. contains 6% of the world's population, but 58% of telephones, 45% of radios, 34% of railroads • Aldous Huxley's <i>The Doors of Perception</i> is published • Cinerama is developed using three-sided screens • The Eurovision network is established 				

	World	Internet	VR	Hypermedia	Computers
1954	<ul style="list-style-type: none"> The TV dinner is invented The microchip is invented 		<ul style="list-style-type: none"> Ray Goertz develops the first electric force-reflecting teleoperator Ralph Mosher at General Electric creates "Handyman", a robotic arm exoskeleton 		
1955	<ul style="list-style-type: none"> The Universal Copyright Convention comes into being Adrian Frutiger designs the Univers typeface <i>MAD</i> magazine makes its debut 				
1955	<ul style="list-style-type: none"> TAT-1, first transatlantic cable capable of carrying telephone calls (52 at a time) <i>Invasion of the Body Snatchers</i> is released, directed by Don Siegel Marshall McLuhan begins writing drafts for <i>Understanding Media</i> Rudolph Flesch's <i>Why Johnny Can't Read</i> is published Bell Laboratories develops the "visual telephone" Morton Heilig writes <i>Cinema of the Future</i> 				
1956			<ul style="list-style-type: none"> A.N. Stanton receives a U.S. patent on a "headgear mounted CRT and binocular viewing device" Morton Heilig begins developing <i>Sensorama</i>, a virtual environment entertainment ride 		
1957 to 1965 - The Age of the True Hackers					
1957	<ul style="list-style-type: none"> The U.S.S.R. launches Sputnik, the first artificial Earth satellite Max Miedinger of Switzerland designs the Helvetica typeface 	<ul style="list-style-type: none"> The Advanced Research Projects Agency (ARPA) is established by President Eisenhower in response to Sputnik 	<ul style="list-style-type: none"> Morton Heilig files patent for "Stereoscopic-Television Apparatus for Individual Use" 		
1958	<ul style="list-style-type: none"> The U.S. artificial Earth satellite Explorer I is launched The U.S.S.R. launches the Sputnik III satellite Hermann Zapf designs the Optima typeface The magnetic OCR typeface is created 45 rpm single records appear Ampex creates the color VTR 		<ul style="list-style-type: none"> Philco Corporation builds the first visual telepresence system with a CRT and camera 		<ul style="list-style-type: none"> FORTRAN, the 1st high-level programming language, is created
1959	<ul style="list-style-type: none"> William Burroughs' <i>Naked Lunch</i> is published in Paris 				<ul style="list-style-type: none"> Texas Instruments announces the first integrated circuit, by Jack S. Kilby Robert Noyce at Fairchild Semiconductor develops an idea for integrated circuits
Nov	<ul style="list-style-type: none"> <i>Communications of the ACM</i> launches 				

	World	Internet	VR	Hypermedia	Computers
1960	<ul style="list-style-type: none"> The U.S. has 85 million televisions, Britain 10.5 million, West Germany 2 million, France 1.5 million The U.S. military budget reaches \$45.8 billion, 49.7% of the U.S. budget Xerox markets the first commercial copier, the 914 xerographic copier Maurits Cornelis Escher makes <i>Ascending and Descending</i> Suzanne Langer's <i>Philosophy in a New Key</i> is published Lasers are invented 				
Mar					<ul style="list-style-type: none"> J.C.R. Licklider writes his "Man-Computer Symbiosis" paper
Apr	<ul style="list-style-type: none"> 50,000 students have participated in sit-ins 				
Jun 13	<ul style="list-style-type: none"> <i>Newsweek's</i> cover story is titled <i>Apathy of Students</i> 				
Oct			<ul style="list-style-type: none"> Morton Heilig receives a U.S. patent on Sensorama 		
Nov					<ul style="list-style-type: none"> Digital Equipment Corporation unveils the PDP-1 (1960 to 1964) Ivan Sutherland works on Sketchpad
1961	<ul style="list-style-type: none"> Newton N. Minow, the FCC's chairman, calls television "a vast wasteland" Steve Russell et. al. begin work on Spacewar at MIT on a PDP-1 Robert Heinlein's <i>Stranger in a Strange Land</i> is published Dolby creates the compact cassette Lego bricks first appear in the U.S. The Berlin Wall is first erected 				
Aug 12			<ul style="list-style-type: none"> Norbert Wiener's <i>Cybernetics; or, Control and Communication in the Animal and the Machine</i> is published 		
1962	<ul style="list-style-type: none"> Arthur C. Clarke's <i>Profiles of the Future</i> is published, with the "indistinguishable from magic" quote Charles Tandy buys the Radio Shack chain and fosters its growth U.S. has 92 million televisions, Europe 26 million, Africa 100,000 <i>How the West Was Won</i> is shot and shown in Cinerama U.S. has 182 million radios, Europe 91 million Rachael Carson's <i>Silent Spring</i> is published 				
			<ul style="list-style-type: none"> Morton Heilig proposes Sensorama as a commercial venture 		
May	<ul style="list-style-type: none"> Spacewar is shown to the public at the annual MIT Open House 				
Jul 19	<ul style="list-style-type: none"> The Telstar satellite transmits its first TV picture across the Atlantic 				
Sep					<ul style="list-style-type: none"> DEC unveils the PDP-4
1963	<ul style="list-style-type: none"> Timothy Leary et. al. write <i>The Psychedelic Experience</i> Color Polaroid film appears 				
					<ul style="list-style-type: none"> The Compatible Time Sharing System is up and running at MIT

	World	Internet	VR	Hypermedia	Computers
1963					<ul style="list-style-type: none"> Ivan Sutherland's Sketch-pad paper is presented at the Spring Joint Computer Conference
Jul 1	<ul style="list-style-type: none"> Project MAC begins at MIT (later: Laboratory for Computer Science) 				
1964	<ul style="list-style-type: none"> Marshall McLuhan's seminal <i>Understanding Media: The Extensions of Man</i> is published AT&T introduces the Picturephone at the New York World's Fair The start of videoconferencing Sony makes a consumer VTR 				<ul style="list-style-type: none"> John Kemeny and Thomas Kurtz develop the BASIC programming language The first word processor: the IBM magnetic Tape Selectric Typewriter
Oct	<ul style="list-style-type: none"> The Free Speech Movement is founded at U.C. Berkeley 				
1965	<ul style="list-style-type: none"> <i>Alphaville</i> is released, directed by Jean-Luc Godard The Moog synthesizer appears 		<ul style="list-style-type: none"> Ivan Sutherland's short paper, The Ultimate Display, is published 	<ul style="list-style-type: none"> Ted Nelson coins the word "hypertext" 	
May					<ul style="list-style-type: none"> IBM unveils the IBM 360/30 and 360/40
1966	<ul style="list-style-type: none"> Ivan Sutherland and Dave Evans form the Evans and Sutherland Computer Company Francois Truffaut directs <i>Fahrenheit 451</i>, the first movie with spoken credits <i>Fantastic Voyage</i> is released, starring Stephen Boyd and Raquel Welch <i>Star Trek</i> starts as a TV series and lasts for the next three years Robert Heinlein's <i>The Moon is a Harsh Mistress</i> is published Harry Harrison's <i>Make Room! Make Room!</i> is published The Kodak Carousel projector debuts 	<ul style="list-style-type: none"> Bob Taylor at DARPA conceives of building a computer network 	<ul style="list-style-type: none"> Ivan Sutherland begins work on the first head-mounted display at MIT's Lincoln Lab, later at University of Utah 		<ul style="list-style-type: none"> William Sutherland, Ivan's brother, makes first visual programming language The SDS-940 is created
Apr					
Oct 31	<ul style="list-style-type: none"> Ken Kesey holds the Acid Test Graduation 				
1967	<ul style="list-style-type: none"> Marshall McLuhan's <i>The Medium is the Massage</i> is published, with graphics by Quentin Fiore <i>The Velvet Underground with Nico</i> is released, produced by Andy Warhol The Architecture Machine Group (ArchMac) is started at MIT Jann Wenner drops out of U.C. Berkeley and starts <i>Rolling Stone</i> David Kahn's <i>The Codebreakers</i> is published 				

	World	Internet	VR	Hypermedia	Computers
1967		<ul style="list-style-type: none"> • SIGGRAPH is formed 	<ul style="list-style-type: none"> • Fred Brooks et. al. work on the GROPE-I system at the University of North Carolina (UNC) 	<ul style="list-style-type: none"> • Andy Van Dam et. al. build the Hypertext Editing System 	<ul style="list-style-type: none"> • DEC unveils the PDP-10
Jun	<ul style="list-style-type: none"> • The first Consumer Electronics Show is held in New York City 				
Jul	<ul style="list-style-type: none"> • Time cover: "The Hippies" 				
1968	<ul style="list-style-type: none"> • <i>Planet of the Apes</i> debuts, starring Charlton Heston, spawns four movies and TV show • Philip K. Dick's <i>Do Androids Dream of Electric Sheep?</i> is published • Jascia Reichardt gives the exhibition <i>Cybernetic Serendipity: the computer and the arts</i> at London's Institute for Contemporary Art • <i>Fantastic Voyage</i> as an animated series runs for two years • Tom Wolfe's <i>The Electric Kool-Aid Acid Test</i> is published • Stanislaw Lem's <i>Solaris</i> is published in Russia • Stanley Kubrick directs <i>2001</i> 	<ul style="list-style-type: none"> • A proposal for a packet-switching network is made to ARPA • Charlie Herzfeld gives Bob Taylor \$1 million to build a small experimental network • Larry Roberts designs a packet-switching, four-node network • Wes Clark invents the Interface Message Processor (IMP) • Larry Roberts and Berry Wessler send out requests to build the IMP • BBN builds the first IMP, team lead by Frank Heart, includes Bob Kahn, Dave Walden, and Severo Ornstein 			<ul style="list-style-type: none"> • Douglas Englebart gives the first presentation of NLS/Augment at the Fall Joint Computer Conference
Jan	<ul style="list-style-type: none"> • (Jan to May) 40,000 students participate in 221 major demonstrations on 101 campuses 				
Nov	<ul style="list-style-type: none"> • <i>The Whole Earth Catalog</i>, started by Stewart Brand, is published 				
<h2>1969 to 1977 - The Age of the Hardware Hackers</h2> <h3>? to ? - The Information Revolution</h3>					
1969	<ul style="list-style-type: none"> • Buckminster Fuller's <i>Operating Manual for Spaceship Earth</i> is published • <i>Colossus: The Forbin Project</i> is released, with scenes filmed at the Lawrence Hall of Science, Berkeley, California • Michael Crichton's <i>The Andromeda Strain</i> is published • Kurt Vonnegut's <i>Slaughterhouse-Five</i> is published 	<ul style="list-style-type: none"> • ARPANET is commissioned by the Department of Defense for networking research • Jon Postel writes the first telnet program • Vint Cerf and Bob Kahn develop TCP/IP 			<ul style="list-style-type: none"> • Intel announces a 1KB RAM chip

	World	Internet	VR	Hypermedia	Computers
1969		<ul style="list-style-type: none"> • Len Kleinrock runs the Network Measurement Center • A data network is built in the U.K., by a team including Roger Scantlebury 			
Jul 20		<ul style="list-style-type: none"> • Neil Armstrong steps on the moon 			
Sep		<ul style="list-style-type: none"> • (Labor Day weekend) The first ARPANET node is established at UCLA on a Xerox DSS 7 (may be Honeywell), which is flown into the campus • The ARPANET UCLA team includes Vint Cerf, Jon Postel, Steve Crocker, and Bill Naylor • The second ARPANET node is established at Stanford Research Institute (SRI) using a SDS-940 • Doug Englebart works on the second node and runs the Network Information Center • The third ARPANET node is established at UCSB using an IBM 360, run by Roland Bryan • The fourth ARPANET node is established at the University of Utah using a PDP-10 			
1970	<ul style="list-style-type: none"> • Richard Buckminster Fuller's <i>I Seem to Be A Verb</i> is published • Nicholas Negroponte's <i>The Architecture Machine</i> is published • <i>THX-1138</i> is released, George Lucas' first full-length effort • Herb Lubalin and Tom Carnase design Avant Garde Gothic • <i>The Andromeda Strain</i> debuts, directed by Robert Wise • Eugene Youngblood's <i>Expanded Cinema</i> is published • Kraftwerk releases its first album, <i>Kraftwerk</i> • <i>Future Shock</i> by Alvin Toffler is published • Datapoint releases its first terminal 	<ul style="list-style-type: none"> • ARPA has a \$238 million budget and changes its name to DARPA (Defense ARPA) • ALOHAnet is developed by Norman Abrahamson at the University of Hawaii 			
Apr					<ul style="list-style-type: none"> • DEC unveils the PDP-11/20
Oct	<ul style="list-style-type: none"> • <i>Scientific American</i> describes John Conway's LIFE program 				
1971	<ul style="list-style-type: none"> • At Xerox, Garry Starkweather's laser printer prototype prints one page per second, 500 dpi • The Berne Convention for the Protection of Literary and Artistic Works is created • Douglas Trumbull directs <i>Silent Running</i>, first film to portray robot doctors • Nolan Bushnell's Computer Space is the first upright computer video game • Donald E. Knuth's <i>The Art of Computer Programming</i> is published • A <i>Time</i> survey shows 75% feel computers will improve lives, 55% feel people are too dependent, 76% encourage computer career • A video display-based word processor is introduced by Lexitron • Paula Mosaides creates "Fuzzy Fonts", anti-aliased fonts, at MIT but do not see the light of day until 1987 • <i>The Omega Man</i> debuts, starring Charlton Heston 				

	World	Internet	VR	Hypermedia	Computers
1971	<ul style="list-style-type: none"> • Yippie phone phreaks start <i>YIPL/TAP</i> magazine • Paul Erlich's <i>The Population Bomb</i> is published • Stanley Kubrick directs <i>A Clockwork Orange</i> • Omnimax cinema is created 	<ul style="list-style-type: none"> • ARPANET has 15 nodes 	<ul style="list-style-type: none"> • The GROPE-II system is completed at UNC 		
May	<ul style="list-style-type: none"> • <i>The Last Whole Earth Catalog</i> is published and wins the National Book Award 				
Jun	<ul style="list-style-type: none"> • The Demise of the Whole Earth Catalog Party is held 				
Nov					<ul style="list-style-type: none"> • Intel introduces the 4004 chip, the first microprocessor, at \$200 and .006 MIPS
1972 to 1984 - The Golden Age of Video Games					
1972	<ul style="list-style-type: none"> • Nicholas Negroponte begins work on his book <i>Soft Architecture Machines</i> • Alan Kay popularizes his DynaBook concept at Xerox PARC • Michael Crichton's <i>The Terminal Man</i> is published • <i>Solaris</i> is released, directed by Andrei Tarkovsky • Philips introduces the first consumer VCR • Sun Microsystems, Inc. is formed 	<ul style="list-style-type: none"> • Bob Kahn organizes a demonstration of 40 machines on ARPANET at the International Conference on Computer Communications • Ray Tomlinson at BBN invents one of the first email programs 		<ul style="list-style-type: none"> • Ted Nelson writes up the original full design for Xanadu • General Electric builds the Advanced Development Model to measure effectiveness of computer graphics for training • Ted Nelson writes ideas for graphic revision trees, virtual desktops, and web visualization techniques 	
Apr 6	<ul style="list-style-type: none"> • Cray Research, Inc. is incorporated 				<ul style="list-style-type: none"> • The personal computer name and concept comes into use
Jun					<ul style="list-style-type: none"> • DEC unveils the PDP-8/M
Nov	<ul style="list-style-type: none"> • Nolan Bushnell founds Atari and ships Pong, its first product 				<ul style="list-style-type: none"> • Intel releases 8008 chip, 8 bits, .3 MIPS, 16KB of memory
1973	<ul style="list-style-type: none"> • <i>Westworld</i> is released, directed by Michael Crichton, starring Yul Brynner and Richard Benjamin • Magnavox releases the Odyssey, the first home video game system • Video word processor with floppy diskettes introduced by Vydec • <i>Soylent Green</i> is released, starring Charlton Heston 				

	World	Internet	VR	Hypermedia	Computers
1973	<ul style="list-style-type: none"> Ivan Illich's <i>Tools for Conviviality</i> is published The modern fax machine comes into being 	<ul style="list-style-type: none"> The ARPANET gains its first international connections to England and Norway Bob Metcalfe describes ideas for Ethernet in his Harvard Ph.D. thesis 	<ul style="list-style-type: none"> Computers render scenes of 200 to 400 polygons per second, 20 frames per second <i>Welt am Draht</i>, a TV movie dealing with simulated reality, is shown in Germany 		
May					<ul style="list-style-type: none"> The French-built Micral becomes the first non-kit computer based on a microprocessor (the 8008)
Jun					<ul style="list-style-type: none"> The word "microcomputer" first appears in print, in reference to the Micral
1974	<ul style="list-style-type: none"> <i>The Conversation</i> is released, directed by Francis Ford Coppola <i>The Terminal Man</i> debuts, starring George Segal The New York Dolls release <i>Too Much, Too Soon</i> The first answering machines are unveiled The first digital watches are announced H.R. Giger paints <i>Biomechanoid I</i> 		<ul style="list-style-type: none"> Myron Krueger coins term "artificial reality" 		<ul style="list-style-type: none"> Gary Kildall develops the CP/M operating system Brian Kernighan and Dennis Ritchie develop the C programming language
Apr					<ul style="list-style-type: none"> Intel releases its 8080 chip, 8 bits, 3 MIPS, 64KB memory
Summer	<ul style="list-style-type: none"> The first SIGGRAPH is held in Boulder, Colorado 				
Jul					<ul style="list-style-type: none"> <i>Radio Electronics</i> publishes a DIY microcomputer article - the Mark 8
Aug				<ul style="list-style-type: none"> Ted Nelson distributes <i>Computer Lib/Dream Machines</i>, coins "hypermedia", among other words 	
Sep					<ul style="list-style-type: none"> Motorola introduces the 6800 chip, 8 bits <i>Creative Computing</i>, the first magazine for home computerists, is founded by Dave Ahl
Oct	<ul style="list-style-type: none"> Fred Brooks finishes writing <i>The Mythical Man-Month</i> about his experience with the IBM 360 OS 				
1975	<ul style="list-style-type: none"> <i>Three Days of the Condor</i> with Robert Redford premieres, introduces phone phreaking to many Gordon Moore predicts that the number of transistors on chips will double every two years The CBGB club fosters the Ramones, Talking Heads, Blondie, Television, et. al. 				

	World	Internet	VR	Hypermedia	Computers
1975	<ul style="list-style-type: none"> • The People's Computer Company (PCC) releases its first newsletter • The U.S. music scene is introduced to Germany's Kraftwerk • Benoit Mandelbrot's paper on his fractal discoveries, <i>Les objets fractals; form, hazard et dimension</i>, is published • John Brunner's <i>The Shockwave Rider</i> is published • Sony announces the Betamax VCR 	<ul style="list-style-type: none"> • The PLATO system starts at the University of Illinois 	<ul style="list-style-type: none"> • Myron Krueger exhibits VIDEOPLACE, which influences future telepresence work, at the University of Connecticut • LEEP Systems, Inc., formerly known as Pop-Optix Labs, creates the LEEP Optical System • Ken Knowlton proposes a virtual push-button system, a virtual keyboard overlay 	<ul style="list-style-type: none"> • ZOG (now KMS), a distributed hypermedia system, debuts at Carnegie-Mellon 	<ul style="list-style-type: none"> • Microsoft is founded • MITS sells BASIC for the Altair 8800, written by Bill Gates and Paul Allen
Jan					<ul style="list-style-type: none"> • DEC unveils the PDP-8/A • <i>Popular Electronics</i> publishes a DIY article on the MITS Altair 8800 computer
Mar 5					<ul style="list-style-type: none"> • The first Homebrew Computer Club meeting occurs in Menlo Park, California
May	<ul style="list-style-type: none"> • Martin Hellman and Whitfield Diffie invent public-key cryptography at Stanford 				
Sep					<ul style="list-style-type: none"> • The first issue of <i>Byte</i> magazine is published
1976	<ul style="list-style-type: none"> • <i>Futureworld</i>, <i>Westworld's</i> sequel, is the first movie to contain a computer-generated sequence • Members of Throbbing Gristle form Industrial Records, coin "industrial music" term • Atari is sold to Warner Communications for \$26 million, Ray Kassar takes over • Fairchild introduces the Fairchild Video Entertainment System • The RCA Studio II home video game system is released • <i>Eraserhead</i> premieres, directed by David Lynch • The Ramones release Ramones, their first album • Richard Dawkins' <i>The Selfish Gene</i> is published • Coleco releases the Telstar Arcade • <i>Dr. Dobbs' Journal</i> launches • JVC debuts the VHS VCR 				<ul style="list-style-type: none"> • The first Cray-1 is announced
Mar	<ul style="list-style-type: none"> • The First World Altair Computer Convention is held 	<ul style="list-style-type: none"> • Mike Lesk at Bell Labs creates UUCP 			

	World	Internet	VR	Hypermedia	Computers
1976					
Apr 1					<ul style="list-style-type: none"> • Apple Computer is formed
Jul					<ul style="list-style-type: none"> • The first Apple is sold in kit form and delivered to stores by Steve Jobs and Steve Wozniak
Aug					<ul style="list-style-type: none"> • iCOM advertises an 8" floppy drive in Byte for \$1200 • Steve Wozniak proposes personal computer to Hewlett-Packard, Steve Jobs proposes to Atari, both rejected
Dec	<ul style="list-style-type: none"> • Steve Wozniak shows the Apple II to the Homebrew Computer Club • (Late 1976) <i>Never Mind the Bollocks Here's the Sex Pistols</i> is released in England 				<ul style="list-style-type: none"> • Michael Shrayner writes Electric Pencil, the first popular microcomputer word-processing program • Shugart announces a 5.25" floppy drive for \$390

1977 to 1987 - The Age of the Game Hackers

1977	<ul style="list-style-type: none"> • Ron Rivest, Adi Shamir, and Leonard Adleman at MIT found RSA Data Securities, Inc. • <i>Demon Seed</i> is released, directed by Donald Cammell, starring Julie Christie • The Spatial Data Management System is conceived at MIT's ArchMac • The first interactive-fiction game, ADVENT (Colossal Cave Adventure), is created by Don Woods and Willie Crother at Stanford's AI lab • Ted Nelson's <i>The Home Computer Revolution</i> is published • Jim Blinn begins work on the computer-generated Voyager fly-by animations, some made on a PDP-11/55 • Space Invaders is introduced as an upright video game • <i>Star Wars</i> incorporates cutting-edge computer graphics • Atari introduces the 2600 VCS 		<ul style="list-style-type: none"> • Dan Sandin, Richard Sayre, and Thomas Defanti create the bend-sensing Sayre Glove at the University of Illinois, Chicago under an NEA grant 		
Feb	<ul style="list-style-type: none"> • Computer Shack opens its first store and later changes name to ComputerLand 				
Apr					<ul style="list-style-type: none"> • Commodore Business Machines, Inc. releases the Commodore PET (6502, 4kb RAM, 14kb ROM) for \$600 • The Apple II (6502, 4kb RAM, 16kb ROM) sells for \$1300

	World	Internet	VR	Hypermedia	Computers
1977					
Apr					<ul style="list-style-type: none"> The First Annual West Coast Computer Faire at San Francisco's Civic Center, organized by Jim Warren
Aug					<ul style="list-style-type: none"> Tandy/Radio Shack sells Steve Leninger's TRS-80 for \$600
	<ul style="list-style-type: none"> (Late summer) Apple ships \$250k worth of equipment 				
1978					
	<ul style="list-style-type: none"> Brian Kernighan and Dennis Ritchie's <i>The C Programming Language</i> is published Brian Eno's <i>Ambient 1: Music for Airports</i> is released, Eno creates ambient music <i>Invasion of the Body Snatchers</i> is remade, directed by Philip Kaufman Devo's <i>Q: Are We Not Men? A: We Are DEVO!</i> is released <i>Germfree Adolescents</i> by X-Ray Spex is released Throbbing Gristle's <i>D.O.A.</i> is released Philips launches the laserdisc Bally releases the Astrocade 			<ul style="list-style-type: none"> Atari announces the Atari 400 and 800 	
				<ul style="list-style-type: none"> Andrew Lippman, Scott Fisher, et. al. at MIT's ArchMac create the Aspen Movie map 	
Feb					<ul style="list-style-type: none"> Digital releases the VAX 11/780
Jun					<ul style="list-style-type: none"> Ward Christenson and Randy Suess create the first personal computer bulletin board system Intel releases the 8086 chip, 16 bits
1979					
	<ul style="list-style-type: none"> <i>Star Trek: The Motion Picture</i> debuts, directed by Robert Wise, spawns at least six sequels Richard Bartle and Roy Trubshaw create the first multi-user dungeon (MUD) <i>Being There</i> is released, directed by Hal Ashby, starring Peter Sellers Milton Bradley releases the Microvision hand-held game system Dick Bolt and Chris Schmandt create "Put That There" at MIT Jenny Holzer begins posting sayings around New York City Douglas Hofstadter's <i>Godel, Escher, Bach: An Eternal Golden Braid</i> is published, wins Pulitzer Prize Gary Numan's <i>The Pleasure Principle</i> is released Devo's <i>Duty Now For the Future</i> is released Sony releases the Walkman Philips announces CD audio Cellular telephones debut 	<ul style="list-style-type: none"> Meetings between DARPA, NSF, et. al. are held to create a Computer Science Department research computer network Tom Truscott and Steve Bellovin create the beginnings of Usenet at Duke University and UNC 		<ul style="list-style-type: none"> Marvin Minsky coins "telepresence" in a ten-year R&D proposal to the U.S. government 	

	World	Internet	VR	Hypermedia	Computers
1979					
Feb					<ul style="list-style-type: none"> Intel releases the 8088 chip
Apr	<ul style="list-style-type: none"> MicroPro releases WordStar word processor 				
May	<ul style="list-style-type: none"> Software Arts, Inc. releases VisiCalc, the first commercial spreadsheet program 				
Jun	<ul style="list-style-type: none"> The Source telecommunications service goes online 				
Aug	<ul style="list-style-type: none"> CompuServe goes online under the name MicroNet 				
Sep					<ul style="list-style-type: none"> Motorola releases the 68000 chip, 16 bits
Dec	<ul style="list-style-type: none"> The first Comdex show is held in Las Vegas 				
1970's		<ul style="list-style-type: none"> Queen Elizabeth becomes the first leader of a country to send an email message 			
1980	<ul style="list-style-type: none"> The Supreme Court rules that organisms created by genetic engineering can be patented Video artists Kit Galloway and Sherri Rabinowitz create "Hole-in-Space" The Minitel (Teletel) network is started in France by French Telecom William Gibson's short story <i>Johnny Mnemonic</i> appears in <i>Omni</i> The biotechnology field experiences a number of breakthroughs The Dead Kennedys release <i>Fresh Fruit for Rotting Vegetables</i> Frederick Pohl's <i>Beyond the Blue Event Horizon</i> is published Mattel releases the Mattel Intellivision nationally Seymour Papert's <i>Mindstorms</i> is published The group Einstuerzende Neubauten forms Alvin Toffler's <i>The Third Wave</i> is published Magnavox releases the Odyssey^2 Apple grosses \$139 million 	<ul style="list-style-type: none"> McDonnell-Douglas makes a see-through HMD with magnetic tracking for its VITAL F-18 simulator 			<ul style="list-style-type: none"> The number of computers in the U.S. exceed one million Roughly two dozen firms selling personal computers The Acorn Atom computer is released
Feb					<ul style="list-style-type: none"> Sinclair announces the ZX80
Mar	<ul style="list-style-type: none"> Microsoft Corporation announces its first hardware product, the Z80 SoftCard for the Apple II 				
May	<ul style="list-style-type: none"> An ad for Mystery House, the first graphic adventure game, is placed in MICRO magazine 				
Jun					<ul style="list-style-type: none"> Shugart begins selling Winchester disk drives Commodore sells the VIC-20 for \$300
Jun 3	<ul style="list-style-type: none"> (To June 6) Computer errors trigger two false nuclear alarms 				
Jul					<ul style="list-style-type: none"> Tandy/Radio Shack introduces the TRS-80 Model III and the Color Computer
Dec	<ul style="list-style-type: none"> The adventure game Zork is taken from a MIT mainframe by Infocom to commercialize it 				
1981	<ul style="list-style-type: none"> Alan Kay becomes Atari's Chief Scientist, with an R&D budget of \$100 million per year Tracy Kidder's <i>The Soul of a New Machine</i> is published and wins Pulitzer Prize 				

	World	Internet	VR	Hypermedia	Computers
1981	<ul style="list-style-type: none"> • H.R. Giger finishes the more popular of his <i>New York City</i> series of works • Vernor Vinge publishes <i>True Names</i> as a story in <i>Binary Star No. 5</i> • The U.S. suffers the highest unemployment rate since 1940 • Nintendo releases Donkey Kong as an upright video game • Silicon Graphics Incorporated is formed • U.S. divorce rates peak 	<ul style="list-style-type: none"> • BITNET comes into being 		<ul style="list-style-type: none"> • Ted Nelson publishes <i>Literary Machines</i>, describing the Xanadu project 	<ul style="list-style-type: none"> • 18% of U.S. schools have computers • Roughly 900,000 personal computers are shipped worldwide • Hayes Microcomputer Products releases the Smartmodem 300 • Franklin Computer Corporation sells the Franklin Ace 100, the first legal Apple II clone • Adam Osborne of Osborne Books introduces the Osborne 1 portable computer (64kb RAM) for \$1800 • After a decade of research, Xerox sells the Xerox Star 8010 for \$50,000 • Microsoft buys all rights to DOS from Seattle Computer Products and introduces MS-DOS • IBM announces the IBM Personal Computer (PC), 8088 chip, 64kb RAM, 40kb ROM, for \$3,000 • Epson releases the HX-20, the first laptop computer • National Semiconductor announces the 32000 chip, the first commercial 32-bit microprocessor
Apr					
May					
Jul					
Aug					
Nov					
Dec					
1982	<ul style="list-style-type: none"> • Pac Man, programmed by Toru Iwatani and 14 others at Namco Limited, grosses \$500 million • <i>Tron</i>, starring Jeff Bridges, is the first movie to make extensive use of computer graphics • The TV series <i>Square Pegs</i> starts, runs for a year, starring Sarah Jessica Parker • <i>Star Trek II: The Wrath of Khan</i>, introduces a groundbreaking computer graphics sequence by Alvy Ray Smith at Lucasfilms • Jerry Buckner and Gary Garcia's "Pac-Man Fever" is released 				<ul style="list-style-type: none"> • Over one hundred firms are selling personal computers

	World	Internet	VR	Hypermedia	Computers
1982	<ul style="list-style-type: none"> Mattel announces the Aquarius home computer <i>Blade Runner</i>, directed by Ridley Scott, debuts <i>Max Headroom</i> debuts on British Channel 4 <i>Android</i> is released, starring Klaus Kinski John Naisbitt's <i>Megatrends</i> is published Video game mania reaches its peak Atari makes \$100 billion this year Atari unveils the Atari 5200 VCS Neil Young's <i>Trans</i> is released 	<ul style="list-style-type: none"> The INWG names TCP/IP (Transmission Control Protocol/Internet Protocol) as the standard protocol suite for ARPANET The Defense Data Network (DDN) is established 	<ul style="list-style-type: none"> Thomas Zimmerman receives a U.S. patent for a optical flex sensing glove 		<ul style="list-style-type: none"> Roughly 1.4 million computers have been shipped worldwide, with 1 million in the U.S. alone Commodore announces the Commodore 64 (6510 chip, 64kb RAM, 20kb ROM) with Microsoft BASIC for \$600
Jan 8	<ul style="list-style-type: none"> The AT&T lawsuit is settled; 22 regional Bell System phone companies are given up 				
Jul	<ul style="list-style-type: none"> Boston's Computer Museum is incorporated 				<ul style="list-style-type: none"> (Summer) Compaq debuts the Compaq portable PC (8088 chip, 128kb RAM) for \$3,000
Oct	<ul style="list-style-type: none"> Lotus announces the Lotus 1-2-3 spreadsheet program 				
Nov	<ul style="list-style-type: none"> The Vectrex is made available to the U.S. public for \$200 				
Dec	<ul style="list-style-type: none"> Wizard vs. Wizards, a video game competition among programmers, is televised from Las Vegas 				
1983	<ul style="list-style-type: none"> <i>Brainstorm</i> is released, directed by Douglas Trumbull, starring Christopher Walken and Natalie Wood <i>Whiz Kids</i> starts as a TV series, continues for a year, stars Matthew Laborteaux <i>WarGames</i> is released, directed by John Badham, starring Matthew Broderick Edward R. Tufte's <i>The Visual Display of Quantitative Information</i> is published <i>Brainstorm</i> is the first major motion picture to have scenes played at 60 fps Thinking Machines Corporation is founded in Cambridge, Massachusetts TAT-7, copper transatlantic cable, carries 9k telephone conversations <i>Automan</i> debuts as a shot-lived series about a virtual superhero Dan Levine at Lucasfilms programs Ballblazer for the Atari 800 Bruce Sterling begins the informal newsletter <i>Cheap Truth?</i> Ithiel de Sola Pool's <i>Technologies of Freedom</i> is published <i>The Day After</i>, a nuclear disaster TV movie, is shown <i>Testament</i> is released, directed by Lynne Littman The font foundry and magazine <i>Emigre</i> starts up The Kurzweil K-250 electronic keyboard debuts 				<ul style="list-style-type: none"> The price of the Commodore 64 drops to \$200 Every tenth book sold in America is a computer book Bjarne Stroustrup creates C++ at Bell Labs

	World	Internet	VR	Hypermedia	Computers
1983	<ul style="list-style-type: none"> • AT&T is still dismantling itself • Philips launches CD audio 	<ul style="list-style-type: none"> • MILNET splits from ARPANET and becomes integrated with the DDN 	<ul style="list-style-type: none"> • DARPA begins the Simulator Network (SIMNET) project • Myron Krueger's <i>Artificial Reality</i> is published 		
Jan	<ul style="list-style-type: none"> • Commodore's sales of the VIC-20 reach 1,000,000 				<ul style="list-style-type: none"> • <i>Time</i> magazine selects the computer as 1982's Machine of the Year • Apple announces the Lisa computer for \$10,000
Jan 1		<ul style="list-style-type: none"> • TCP/IP becomes the standard for ARPANET 			
Jun	<ul style="list-style-type: none"> • Apple Computer has shipped 1,000,000 computers • Coleco announces the Adam computer 				
Jul	<ul style="list-style-type: none"> • William Gibson finishes writing <i>Neuromancer</i> 				

1984 to 1992 - The Golden Age of VR

1984 to 1991 - Cyberpunk

1984	<ul style="list-style-type: none"> • <i>The Fundamentals of Human-Computer Interaction</i> is published, edited by Andrew Monk • <i>Revenge of the Nerds</i> debuts, starring David Carradine, spawns at least three sequels • James Cameron's <i>The Terminator</i> is released, shows Apple assembly code in movie • <i>The Last Starfighter</i> is the first movie to do all special effects via computer (a Cray) • IBM makes \$4 billion, Apple, \$1.9 billion, Commodore, \$1.1 billion, HP, \$510,000 • Apple's 1984 commercial is shown over the Super Bowl, directed by Ridley Scott • <i>Lensman</i> debuts as an animated movie with computer-generated effects • Steven Levy's <i>Hackers: Heroes of the Computer Revolution</i> is published • Richard Saul Wurman hosts the first elite TED conference • The Comprehensive Crime Control Act give the Secret Service jurisdiction over credit card and computer fraud • The year of the great video game crash • <i>The Whole Earth Software Catalog</i> is published • <i>2600: The Hacker Quarterly</i> is founded • The Legion of Doom (LOD) is formed 	<ul style="list-style-type: none"> • Domain Name Server (DNS) service begins on ARPANET 	<ul style="list-style-type: none"> • Michael McGreevy, Stephen Ellis, and Humphries at NASA Ames build the Virtual Visual Environment Display, the first \$2k opaque HMD 		<ul style="list-style-type: none"> • Digital Research announces the GEM icon/desktop user interface for 8086-based computers • Several companies announce 2400 baud modems at Comdex for \$800-900 • SGI releases its first workstation at \$80k, .33 MIPS
-------------	---	--	---	--	---

	World	Internet	VR	Hypermedia	Computers
1984			<ul style="list-style-type: none"> David Zeltzer at MIT's Media Lab creates George, a virtual skeleton based on constraints 	<ul style="list-style-type: none"> Telos introduces Filevision, a hypermedia database for the Macintosh 	
Jan	<ul style="list-style-type: none"> Jack Tramiel, founder and president of Commodore, leaves the company 				<ul style="list-style-type: none"> Apple introduces the Macintosh for \$2500
Jan					<ul style="list-style-type: none"> Seiko Instruments U.S.A. Inc. displays the first wrist-watch computer, with a 10-character, 4-line LCD
Mar	<ul style="list-style-type: none"> The <i>National Geographic</i> runs an eagle hologram on its cover, the first hologram in a magazine 				<ul style="list-style-type: none"> IBM announces the IBM PCjr
May	<ul style="list-style-type: none"> A.K. Dewdney writes about Core War in <i>Scientific American</i> 				<ul style="list-style-type: none"> Microsoft announces its first mouse, "The Microsoft Mouse", for \$200
Jun	<ul style="list-style-type: none"> Tom Jennings creates the FidoNet BBS network 				<ul style="list-style-type: none"> Motorola announces the 68020 chip, 32 bits
Jul	<ul style="list-style-type: none"> William Gibson's <i>Neuromancer</i> is published and coins the word "cyberspace" Jack Tramiel buys Atari from Warner Communications 				
Aug	<ul style="list-style-type: none"> Commodore purchases Amiga Corporation 				<ul style="list-style-type: none"> IBM announces the PC AT (80286 chip, 256kb RAM) for \$4,000
Oct		<ul style="list-style-type: none"> The number of ARPANET hosts breaks 1k 			
Nov	<ul style="list-style-type: none"> Vernor Vinge's <i>True Names</i> is published as a book, with a Marvin Minsky afterward Stuart Brand hosts The First Hackers' Conference Microsoft announces Windows 				

1985 to 1990 - The Age of the BBS Hackers

1985	<ul style="list-style-type: none"> Dire Straits' <i>Money for Nothing</i> is the first music video to use computer generated images Richard Stallman forms the Free Software Foundation, Inc. and begins the GNU project "Dogfight", a short story by William Gibson and Michael Swanwick, appears in <i>Omni</i> Douglas Adams' <i>The Hitchhiker's Guide to the Galaxy</i> is published Nintendo introduces its 8-bit Nintendo Entertainment System Xerox PARC begins work on their Media Spaces project <i>The Quiet Earth</i> is released, directed by Geoff Murphy The first police "sting" bulletin boards are established Print media captures 34% of U.S. advertising dollars "The New Science Fiction" appears in <i>Interzone #14</i> 	<ul style="list-style-type: none"> Aldus introduces Page-Maker for the Macintosh, starts desktop publishing Intermedia, a hypermedia system, is conceived by Norman Meyrowitz et. al. in IRIS at Brown University Janet Walker creates the Symbolics Document Examiner Xerox NoteCards is created by Frank Halasz and Randy Trygg
-------------	---	---

	World	Internet	VR	Hypermedia	Computers
1985	<ul style="list-style-type: none"> Howard Rheingold's <i>Tools for Thought</i> is published The Whole Earth 'Lectronic Link (WELL) is formed <i>Threads</i>, a nuclear disaster TV movie, is shown Orson Scott Card's <i>Ender's Game</i> is published Peter G. Neumann starts the RISKS mailing list Bruce Sterling's <i>Schismatrix</i> is published Greg Bear's <i>Blood Music</i> is published Microsoft ships Windows 1.0 Holographic film is invented Sony debuts Video 8 		<ul style="list-style-type: none"> Scott Fisher et. al. work on the Virtual Interactive Environment Workstation (VIEW) project at NASA Ames The U.S. Air Force develops the Super Cockpit Jaron Lanier and Jean-Jacques Grimaud found VPL Research, Inc. 		
Jan					<ul style="list-style-type: none"> Commodore sells the Commodore 128
Feb	<ul style="list-style-type: none"> Steve Jobs resigns from Apple 				
Apr	<ul style="list-style-type: none"> IBM abandons production of the IBM PCjr 				
Mar	<ul style="list-style-type: none"> <i>First, Last, and Always</i> is released by Sisters of Mercy 				
Jul					<ul style="list-style-type: none"> Commodore announces the Amiga 1000, with a multitasking, windowing OS, for \$1,300
Oct					<ul style="list-style-type: none"> Intel releases the 80386 chip, 32 bits
Nov 17	<ul style="list-style-type: none"> <i>Phrack</i> releases its first issue 				
Dec	<ul style="list-style-type: none"> Nintendo NES is test-marketed in New York City 				
1986	<ul style="list-style-type: none"> <i>CD-ROM: The New Papyrus</i> is published, edited by Steve Lambert and Suzanne Ropiequet <i>Mirrorshades: The Cyberpunk Anthology</i> is published, edited by Bruce Sterling K. Eric Drexler's <i>The Engines of Creation</i> is published, coins "nanotechnology" Richard Dawkins' <i>The Selfish Gene</i> is republished, coins the word "meme" Lucasfilms' Pixar generates <i>Andre the Android and Wally Bee</i> The Electronic Communications Privacy Act is passed Richard Dawkins' <i>The Blind Watchmaker</i> is published Susan Lammers' <i>Programmers at Work</i> is published Thomas Pynchon's <i>Gravity's Rainbow</i> is published Marvin Minsky's <i>The Society of Mind</i> is published The Computer Fraud and Abuse Act is passed William Gibson's <i>Burning Chrome</i> is published William Gibson's <i>Count Zero</i> is published 			<ul style="list-style-type: none"> The word "hypertext" comes into common use OWL introduces Guide, a hypermedia document browser 	<ul style="list-style-type: none"> Motorola announces the 68030 chip 95% of the software for home computers is in the form of video games The number of computers in the U.S. exceeds 30 million

	World	Internet	VR	Hypermedia	Computers
1986		<ul style="list-style-type: none"> The NSFnet is created and establishes five supercomputing centers - Princeton, Pittsburgh, UCSD, UIUC, and the Theory Center at Cornell NNTP (Network News Transfer Protocol) is introduced Dozens of universities come online via NSFnet NSFnet's backbone speed is 56kbps The Cleveland Freenet is started 	<ul style="list-style-type: none"> VPL creates the first finished working version of the DataGlove, by Tom Zimmerman and Young Harvill Work on the GROPE-III system is begun at UNC The Air Force adds voice command recognition and response with 3D audio to VCASS 		
Jan	<ul style="list-style-type: none"> The Mentor writes "The Conscience of a Hacker" in <i>Phrack #7</i> 	<ul style="list-style-type: none"> The first Internet Engineering Task Force meeting is held in San Diego, with 15 attendees 			<ul style="list-style-type: none"> Apple introduces the Macintosh Plus
May			<ul style="list-style-type: none"> SIMNET goes online 		
Aug					<ul style="list-style-type: none"> Intel ships the 80386 chip
Sep					<ul style="list-style-type: none"> Compaq introduces its 386-based PC
1987	<ul style="list-style-type: none"> Microsoft Press publishes a revised edition of <i>Computer Lib/Dream Machines</i> <i>Star Trek: The Next Generation</i> runs until 1994 and introduces the holodeck France's Minitel network has 2 million displays in use with over 3k services Chicago prosecutors form the Computer Fraud and Abuse Task Force The Association of Shareware Professionals (ASP) is formed James Gleick's <i>Chaos: Making a New Science</i> is published Fisher-Price debuts the PXL-2000 camcorder for under \$100 Paul Verhoeven's <i>Robocop</i> is released; Robocop runs on a PC, spawns two sequels and a TV series The Nielsen Company creates the People Meter Stewart Brand's <i>The Media Lab</i> is published <i>Max Headroom</i> premieres on U.S. television DAT recorders appear 	<ul style="list-style-type: none"> The NSF agrees to manage the NSFnet backbone with Merit Network, Inc. The number of ARPANET hosts breaks 10k 	<ul style="list-style-type: none"> SIMNET has about 250 simulators online British Aerospace creates the Virtual Cockpit Jon Waldern founds W Industries VPL introduces the DataSuit 		<ul style="list-style-type: none"> Atari introduces the Atari 7800 DEC sells the microVAX 2000

	World	Internet	VR	Hypermedia	Computers
1987				<ul style="list-style-type: none"> The Hypertext '87 Workshop is held in North Carolina 	
Apr	<ul style="list-style-type: none"> Max Headroom appears on the cover of <i>Newsweek</i> VGA graphics are available for PCs 				
Aug					<ul style="list-style-type: none"> IBM and Microsoft announce OS/2
Nov	<ul style="list-style-type: none"> The Future of Freedom Conference is held (To 1988) Acid House develops in the U.K. 				<ul style="list-style-type: none"> NewTek's Video Toaster is unveiled at Comdex
1988					
		<ul style="list-style-type: none"> The <i>Mississippi Review</i> publishes its "Cyberpunk Issue", edited by Larry McCaffery The magazine <i>boING boING</i> is founded by Mark Frauenfelder and Carla Sinclair TAT-8, the first fiber optic cable installed, can carry 40k conversations Donald Norman's <i>The Psychology of Everyday Things</i> is published AT&T's fiber-optic transatlantic cable carries 40k telephone calls Timothy May writes the Crypto Anarchist Manifesto William Gibson's <i>Mona Lisa Overdrive</i> is published Sonic Youth's <i>Daydream Nation</i> is released Bruce Sterling's <i>Islands in the Net</i> is published <i>Akira</i> debuts, directed by Katsuhiro Otomo The term "rave" enters the culture 	<ul style="list-style-type: none"> Jarkko Oikarinen writes Internet Relay Chat (IRC) VPL makes the first commercial VR system, RB2 (Reality Built for Two) Virtuality, a virtual video game by W Industries, is released 		
Apr				<ul style="list-style-type: none"> Ted Nelson joins Autodesk as an Autodesk Fellow 	
		<ul style="list-style-type: none"> (Summer) The first multicast (MBONE) tunnel is made between BBN and Stanford University 			
May	<ul style="list-style-type: none"> Cliff Stoll writes "Stalking the Wily Hacker" in <i>Communications of the ACM</i> 				
Jul	<ul style="list-style-type: none"> The Secret Service secretly tapes the SummerCon hacker convention 				
Aug					<ul style="list-style-type: none"> NeXT unveils the NeXT computer (68030 chip, 8MB RAM, 17" monitor, 256MB MO drive) for \$6,500
Sep 1			<ul style="list-style-type: none"> John Walker at Autodesk writes <i>Through the Looking Glass</i>, an internal white paper 		
Nov		<ul style="list-style-type: none"> Robert Tappan Morris unleashes his Internet worm, disables about 6,600 computers 			
1989					
	<ul style="list-style-type: none"> <i>Tetsuo: The Iron Man</i> premieres, directed by Shinya Tsukamoto, spawns a sequel in 1992 An entire UNIX file system is viewed using Xerox PARC's Information Visualizer 				

	World	Internet	VR	Hypermedia	Computers
1989	<ul style="list-style-type: none"> Marc Najork drafts Cube, the first 3D programming language Bill Landreth's <i>Out of the Inner Circle</i> is published Nine Inch Nails' <i>Pretty Hate Machine</i> is released Clifford Stoll's <i>The Cuckoo's Egg</i> is published Esther Dyson begins <i>Release 1.0</i> Nintendo releases the Game Boy Sony launches the CD walkman 	<ul style="list-style-type: none"> RIPE (Reseaux IP Europeens) is formed to coordinate the operation of the pan-European IP network NSFnet's backbone upgrades to T1 (1.544kbps) 	<ul style="list-style-type: none"> Jaron Lanier gives keynote speech at SIGGRAPH Jaron Lanier coins the term "virtual reality" Mattel introduces the PowerGlove for use with the Nintendo Entertainment System 	<ul style="list-style-type: none"> Tim Berners-Lee writes Enquire at CERN, a personal organizer 	<ul style="list-style-type: none"> Sun announces the 12.5 MIPS SPARCstation 1 for \$9,000
Mar				<ul style="list-style-type: none"> Tim Berners-Lee submits the first proposal for the World-Wide Web 	
Apr	<ul style="list-style-type: none"> The VESA standard provides a uniform method of accessing Super VGA chipsets 				
Jun	<ul style="list-style-type: none"> Nintendo captures about 80% of the video game market, with sales more than twice as large as Microsoft's The "NuPrometheus League" distributes proprietary system software from Apple Computer 				
Jul		<ul style="list-style-type: none"> The number of Internet hosts breaks 100,000 			
Jul 7			<ul style="list-style-type: none"> Autodesk gives the first in-house demonstration of Cyberia 		
Jul 20	<ul style="list-style-type: none"> Wes Thomas creates "Friday the 13th" virus hoax, reported by CNN (Fall) <i>Mondo 2000</i> prints its first issue, earlier named <i>Reality Hackers</i> 		<ul style="list-style-type: none"> (Fall) Broderbund releases U-Force for the Nintendo 		
Nov	<ul style="list-style-type: none"> (Late 1980's) Xerox PARC begins work on "electronic co-presence" systems The East German government opens the country's borders (Third quarter) The first real-time holographic 3D images are produced by the Media Lab's Spatial Imaging Group 		<ul style="list-style-type: none"> (Late 1980's) Chip Morningstar and F. Randall Farmer create Habitat for Lucasfilm Games, which makes use of the word "avatar" Eric Gullichsen and Patrice Gelband leave Autodesk to form Sense 8 		
1990	<ul style="list-style-type: none"> <i>The Art of Human-Computer Interface Design</i>, edited by Brenda Laurel, is published The Electronic Frontier Foundation (EFF) is founded by Mitch Kapor <i>Hardware</i>, one of the first B-grade cyberpunk films, is released Richard Stevens' <i>UNIX Network Programming</i> is published 				<ul style="list-style-type: none"> Gilbert Hyatt is granted a basic patent on the micro-processor, 20 years after his first patent application

	World	Internet	VR	Hypermedia	Computers
1990	<ul style="list-style-type: none"> Edward R. Tufte's <i>Envisioning Information</i> is published Sega releases the Sega Master System II Thomas Ray creates the Tierra Simulator 	<ul style="list-style-type: none"> Hytelnet is released by Peter Scott at the University of Saskatchewan The ARPANET ceases to exist 	<ul style="list-style-type: none"> British Aerospace creates Virtual Environment Training Aids W Industries' Virtuality System is released 	<ul style="list-style-type: none"> ECHT, the European Conference on Hypertext, is held 	<ul style="list-style-type: none"> Sun sells its SPARCstation 2 series for \$20,000 NewTek sells the VideoToaster for \$1,600
Jan	<ul style="list-style-type: none"> Four hackers are raided in separate busts AT&T's long distance network crashes 	<ul style="list-style-type: none"> Sense 8 is founded 			<ul style="list-style-type: none"> Motorola announces the 68040 chip, with 1.2 million transistors
Feb	<ul style="list-style-type: none"> Five hackers are arrested in separate busts 				
Mar	<ul style="list-style-type: none"> The Chicago Task Force raids Steve Jackson Games in Austin, Texas 				
May	<ul style="list-style-type: none"> The U.S. Secret Service engages Operation Sun Devil, which nets 28 raids in 14 days 	<ul style="list-style-type: none"> First Conference on Cyberspace in Austin, Texas 			
Aug	<ul style="list-style-type: none"> The first BattleTech Center opens on Chicago's North Pier 				
Oct				<ul style="list-style-type: none"> Tim-Berners Lee and Robert Cailliau submit a revised Web proposal at CERN 	
Nov	<ul style="list-style-type: none"> Pavel Curtis opens up LambdaMOO, the first text-based collaborative object-oriented environment 			<ul style="list-style-type: none"> The first prototype Web browser runs on a NeXT machine Nicola Pellow begins work on the CERN line-mode browser 	
1991	<ul style="list-style-type: none"> William Strauss and Neil Howe's <i>Generations: The History of America's Future, 1584-2069</i> is published James Cameron's <i>Terminator 2: Judgement Day</i> is released, breaks new ground in computer graphics Douglas Coupland's <i>Generation X: Tales For an Accelerated Culture</i> is published David Blair's <i>WAX or the discovery of television among the bees</i> premieres William Gibson and Bruce Sterling's <i>The Difference Engine</i> is published The Computer Lab releases the <i>Beyond Cyberpunk!</i> HyperCard stack <i>Storming the Reality Studio</i> is published, edited by Larry McCaffery <i>Cyberspace: First Steps</i>, edited by Michael Benedikt, is published <i>Until the End of the World</i> is released, directed by Wim Wenders <i>The Hacker's Dictionary</i> is published, edited by Eric S. Raymond 	<ul style="list-style-type: none"> WAIS (Wide Area Information Server) is released by Thinking Machines Corporation The Commercial Internet Exchange (CIX) Association, Inc. is formed by the owners of CERFnet, PSInet, and AlterNet 			<ul style="list-style-type: none"> Motorola's 68040 chip becomes available Tandy introduces its low-cost PC CD-ROM drive for \$400

	World	Internet	VR	Hypermedia	Computers
1991		<ul style="list-style-type: none"> The U.S. High Performance Computing Act establishes the National Research and Education Network (NREN) 	<ul style="list-style-type: none"> Howard Rheingold's <i>Virtual Reality</i> is published Myron Krueger's <i>Artificial Reality II</i> is published The University of North Carolina renders 1 million untextured polygons per second, 200k textures polys/sec The Advanced Robotics Research Center (ARRC) creates Teletact I, the first tactile feedback glove British Aerospace's Virtual Environment Training Aid (VECTA) is displayed at the Paris Airshow 	<ul style="list-style-type: none"> Gopher is released by the University of Minnesota The first HTTP servers outside of CERN are set up at Stanford's SLAC and NIKHEF 	
Jan					
Mar	<ul style="list-style-type: none"> San Francisco is host to the Computers, Freedom, and Privacy Conference 			<ul style="list-style-type: none"> CERN's line-mode browser is released on a Sun, VAX and RS/6000 	
May	<ul style="list-style-type: none"> Apple releases System 7.0, two years late 				
Jun	<ul style="list-style-type: none"> Phil Zimmerman releases Pretty Good Privacy (PGP) 		<ul style="list-style-type: none"> Sense 8 releases its WorldToolkit 		
Jun 12				<ul style="list-style-type: none"> Tim Berners-Lee, Robert Cailliau, and Nicola Pellow hold a Web seminar at CERN 	
Jul					<ul style="list-style-type: none"> Apple and IBM sign agreement to work on multimedia standards, make new OS, and use PowerPC chips
Aug	<ul style="list-style-type: none"> (Third quarter) The Media Lab generates full-color holographs 			<ul style="list-style-type: none"> The World-Wide Web is announced in alt.hypertext, comp.sys.next, comp.text.sgml, and comp.mail.multi-media 	
Sep	<ul style="list-style-type: none"> An AT&T phone crash affects New York City and three airports 				
Oct				<ul style="list-style-type: none"> The www-talk mailing list is started CERN enables their anonymous telnet Web browser 	
Dec	<ul style="list-style-type: none"> Nirvana's <i>Nevermind</i> is released 		<ul style="list-style-type: none"> (Late 1991) VR appears to have reached megahype proportions by this time 	<ul style="list-style-type: none"> Tim Berners-Lee gives a poster and demonstration of his Web browser at Hypertext '91, San Antonio, Texas 	
1992	<ul style="list-style-type: none"> Bill Cheswick of Bell Labs writes <i>An Evening with Berferd In Which a Cracker is Lured, Endured, and Studied</i> Texas Instruments projects the first image from a Digital Micromirror Device (DMD) chip in Dallas, Texas <i>Sneakers</i> is released, directed by Phil Robinson, starring Robert Redford 				

	World	Internet	VR	Hypermedia	Computers
1992	<ul style="list-style-type: none"> • Bruce Sterling's <i>The Hacker Crackdown</i> is published • Terence McKenna's <i>Food of the Gods</i> is published • <i>Slacker</i> is released, directed by Richard Linklater • Kodak debuts Photo CD 	<ul style="list-style-type: none"> • The NSFnet backbone is upgraded to T3 (44.736Mbps) • The Internet Society is chartered • Internet World launches 	<ul style="list-style-type: none"> • British Aerospace's Real and Virtual Environment Training Aid (RAVECTA) is displayed at the Farnborough International Air Show • Disney's Imagineering Labs begins work on the virtual <i>Aladdin</i> ride • The Virtual Cockpit project begins at the U.S. Air Force's Wright Labs 		
Jan				<ul style="list-style-type: none"> • Pei Wei gives early version of Viola, the first scriptable, GUI Web browser for X, to CERN 	
Mar		<ul style="list-style-type: none"> • The MBONE multicasts its first audio media 		<ul style="list-style-type: none"> • The possibility of adding graphics to the Web is discussed 	
Mar 6	<ul style="list-style-type: none"> • The PC virus Michaelangelo causes a nation-wide stir, is reported on CNN 		<ul style="list-style-type: none"> • <i>The Lawnmower Man</i> is released, starring Pierce Brosnan and Jeff Fahey, introduces VR to many 		
Jul				<ul style="list-style-type: none"> • Pei Wei releases his Viola browser 	
Jul 25				<ul style="list-style-type: none"> • Joseph Wang at MIT releases tkWWW 0.1 	
Jul 26				<ul style="list-style-type: none"> • Ari Lemmke releases Erwise 0.1, a Motif/X Web client 	
Aug			<ul style="list-style-type: none"> • SGI's Inventor is announced 	<ul style="list-style-type: none"> • Autodesk drops the Xanadu project 	
Sep				<ul style="list-style-type: none"> • Tony Johnson begins work on MidasWWW at Stanford's Linear Accelerator Center (SLAC) 	
Oct		<ul style="list-style-type: none"> • The number of Internet hosts breaks 1 million 			
Nov	<ul style="list-style-type: none"> • The first completely digital telephone network is launched in the U.S. 	<ul style="list-style-type: none"> • The MBONE multicasts its first video media 	<ul style="list-style-type: none"> • VPL ceases operation 		
Nov 16				<ul style="list-style-type: none"> • Tony Johnson releases MidasWWW 1.0 	
1993 to ? - The Age of the Internet					
1993	<ul style="list-style-type: none"> • Katie Hafner and John Markoff's <i>Cyberpunk: Outlaws and Hackers on the Computer Frontier</i> is published • <i>Jurassic Park</i>, directed by Steven Spielberg, breaks new ground in computer-generated imagery 				

	World	Internet	VR	Hypermedia	Computers
1993		<ul style="list-style-type: none"> • Roy Kalawsky's <i>The Science of Virtual Reality and Virtual Environments</i> is published • TAT-11 installed from U.S. to France; carries over 22k conversations • Neil Howe's <i>13th Gen: Abort, Retry, Ignore, Fail?</i> is published • Vernor Vinge's <i>A Fire Upon the Deep</i> is published • William Gibson's <i>Virtual Light</i> is published • Broderbund releases <i>Myst</i> on CD-ROM • AT&T's "You Will" ad campaign begins 			
		<ul style="list-style-type: none"> • Carl Malamud founds Internet Talk Radio and begins broadcasting • NSF creates the InterNIC to provide directory, registration, and information services on the Internet • The United Nations and the World Bank come online • The U.S. White House comes online 	<ul style="list-style-type: none"> • The IEEE issues DIS (Distributed Interactive Simulation) as a standard protocol • Mark Pesce and Tony Parisi begin work on Labyrinth, a 3D space browser 		
Jan				<ul style="list-style-type: none"> • <i>A Hard Day's Night</i> becomes the first full-length movie to be transcribed into hypertext and distributed on a CD • The growth of Web traffic in 1993 reaches 944% by NSFnet statistics • The first alpha version of Mosaic for the X Window System is released • About 50 HTTP servers are known to exist • First discussions of forms on www-talk 	
	<ul style="list-style-type: none"> • Wired's first issue goes on the bookshelves 				
Feb		<ul style="list-style-type: none"> • Babylon 5's pilot TV movie shows, displaying the Amiga's graphic power 			
Feb 8		<ul style="list-style-type: none"> • Time cover: "Cyberpunk" 			
Feb 25				<ul style="list-style-type: none"> • Marc Andreessen proposes the IMG tag 	
Mar		<ul style="list-style-type: none"> • Bill Clinton becomes the first U.S. president to send an email message 			
Mar				<ul style="list-style-type: none"> • Web traffic makes up one tenth of one percent on the NSFnet backbone • Marc VanHeyningen finishes the first Perl HTTP server 	
Mar 9	<ul style="list-style-type: none"> • Eric Hughes writes <i>A Cypherpunk's Manifesto</i> 				
Mar 11				<ul style="list-style-type: none"> • Lou Montulli releases Lynx 2.0a, with Web support for the first time 	
Mar 14				<ul style="list-style-type: none"> • Marc Andreessen releases XMo-saic 0.10, with IMG support 	
Mar 30				<ul style="list-style-type: none"> • Tim Berners-Lee attends Web BOF at the 26th IETF meeting in Columbus, Ohio 	

	World	Internet	VR	Hypermedia	Computers
1993					
Apr					
Apr 6					
Apr 21					
Apr 22					
Apr 24					
May					
May 1					
May 4					
May 7					
May 13					
Jun					
Jun 14					
Jun 15					
Jul					
Jul 5					
Jul 6					
Jul 26					
Aug					
Sep					
Sep 27					
Oct					

	World	Internet	VR	Hypermedia	Computers
1993					
Oct				• Over 500 known Web servers exist	
Oct 7				• Chris MacRae organizes the first SIGWEB meeting at UCSF	
Nov				• Hypertext '93 is held in Seattle, Washington, with Ted Nelson as guest of honor	
Nov 10				• Mosaic for X 1.0 is released, which supports most form elements	
Nov 17				• Rob McCool proposes the Common Gateway Protocol (CGI)	
Nov 19				• SIGWEB #2 is held at Xerox PARC	
Nov 23				• The World-Wide Web and Mosaic is mentioned in <i>Open Systems Today</i>	
Nov 29				• Kevin Hughes organizes a WWW BOF at Hypertext '93 in Seattle, Washington	
Dec		• The First Internet World Conference is held in New York City		• An article on the Web is written in the UK's Guardian	
Dec 8				• John Markov writes a Web and Mosaic article in the <i>New York Times</i>	
Dec 10	• Doom is released by Id Software				
Dec 14				• Robert Cailliau announces the First International WWW Conference	
Dec 27	• NBC's series <i>Almost 2001</i> offers an email address for feedback				
		• (Late 1993) Cornell University releases CU-SeeMe			
1994	• Kevin Kelly's <i>Out of Control: The Rise of Neobiological Civilization</i> is published				
	• TekWar is the first TV movie to portray a 360-degree virtual cyberspace				
	• The Knightmare's <i>Secrets of a Super Hacker</i> is published				
	• Bruce Schneier's <i>Applied Cryptography</i> is published			• The growth of Web traffic in 1994 reaches 1,713% by NSFnet statistics	
	• <i>Clerks</i> is released, filmed for \$27k, grosses millions				
	• Richard Preston's <i>The Hot Zone</i> is published				
	• <i>Reality Bites</i> debuts, starring Winona Ryder				
		• O'Reilly and Associates and Spry announce Internet in a Box			
			• Dave Raggett coins the term VRML, originally "Virtual Reality Markup Language"		
			• 1,000 entities are involved in a Synthetic Theater of War (STOW) exercise		
			• The first virtual marriage takes place		
			• Kaiser Electro-Optics, Inc. offers an HMD with 1,024 scan lines for \$165k		

	World	Internet	VR	Hypermedia	Computers
1994					
Jan					<ul style="list-style-type: none"> • MCI's "information highway" ad campaign starring Anna Paquin begins • <i>Microserfs</i> is printed as a short story in <i>Wired</i>
Jan 11					<ul style="list-style-type: none"> • Vice-President Al Gore gives his information superhighway speech at UCLA
Feb					<ul style="list-style-type: none"> • Palo Alto, California becomes the first U.S. city to have its own Web site
Feb 3					<ul style="list-style-type: none"> • SIGWEB #4 is held at the Amdahl Corporation in Sunnyvale, California
Feb 14					<ul style="list-style-type: none"> • Mark Pesce and Tony Parisi finish a working version of a 3D viewer, <i>Labyrinth</i>, that fetches objects using HTTP
Feb 19					<ul style="list-style-type: none"> • Lou Montulli, Kevin Altis, and Ari Luotonen's lib-<i>WWW</i> proxy gateway code is released
Mar					<ul style="list-style-type: none"> • Bruce Tognazzini publically shows <i>Starfire</i>, Sun Microsystem's conceptual interface, at Stanford • Sense 8 releases its <i>WorldToolKit</i> for Windows • <i>WWW</i> traffic surpasses that of <i>Gopher</i> on the <i>NSFnet</i>
Apr					<ul style="list-style-type: none"> • Work on Vangelis' <i>Blade Runner</i> soundtrack is completed • <i>CommerceNet</i> is launched from San Jose, California • Jim Clark and Mark Andreessen form Mosaic Communications Corporation (later <i>Netscape</i>) • A SIGWEB meeting is held in the SLAC Auditorium
Apr 19					
May					<ul style="list-style-type: none"> • Nearly one of three U.S. households has a personal computer and 12% have modem, 46% of teenagers have a PC in-house • The First International World-Wide Web Conference is held at CERN in Geneva, Switzerland, 380 participants • The World-Wide Web Hall of Famers: Tim Berners-Lee, Marc Andreessen, Eric Bina, Rob Hartill, Kevin Hughes, Lou Montulli • Mark Pesce presents a paper on <i>Cyberspace Protocol</i> at the First International <i>WWW</i> Conference in Geneva • Mark Pesce and Brian Behlendorf convene a <i>VRML BOF</i> at the First International <i>WWW</i> Conference
May 4					<ul style="list-style-type: none"> • Brandon Plewe opens the <i>Best of the Web '94</i> to votes
Jun					<ul style="list-style-type: none"> • Over 1500 registered HTTP servers exist
					<ul style="list-style-type: none"> • The Second Internet World Conference is held in San Jose, California
Jun					<ul style="list-style-type: none"> • World Conference on Educational Multimedia and Hypermedia in Vancouver, Canada • The <i>www-vrml</i> mailing list opens, over 2,000 people subscribe in the first week

	World	Internet	VR	Hypermedia	Computers
1994					
Jun 8				<ul style="list-style-type: none"> • SIGWEB #8, U.C. Berkeley's Dwinelle Hall, Berkeley, California 	
Jul				<ul style="list-style-type: none"> • MIT and CERN announce the formation of the World-Wide Web Consortium 	
				<ul style="list-style-type: none"> • The U.S. has 193 million TV sets in use, 1k TV broadcast systems, and about 9k TV cable systems 	
August			<ul style="list-style-type: none"> • SIGGRAPH '94 holds a discussion group for selecting the basis of VRML 	<ul style="list-style-type: none"> • Gavin Bell announces a SIGWEB for the UK and Ireland 	
				<ul style="list-style-type: none"> • 1,600 copies of Mosaic are downloaded per day from the NCSA 	
Sep				<ul style="list-style-type: none"> • ReBoot premieres in U.S. households as the first fully computer generated animation series 	
				<ul style="list-style-type: none"> • European Conference on Hypermedia Technology (ECHT '94) in Edinburgh, Scotland; Berners-Lee gives keynote 	
Sep 27				<ul style="list-style-type: none"> • CERN httpd 3.0, the first proxy server, is released 	
Oct			<ul style="list-style-type: none"> • Steven Spielberg, Jeffrey Katzenberg, and David Geffen form DreamWorks SKG 	<ul style="list-style-type: none"> • The Second International WWW Conference is held in Chicago, Illinois, with over 1,500 participants 	
Nov			<ul style="list-style-type: none"> • Gavin Bell of SGI, Tony Parisi, and Mark Pesce write the VRML 1.0 draft • Nintendo announces "Virtual Boy", a head-mounted display 		
Nov 17		<ul style="list-style-type: none"> • The IETF recommends IPNG as the solution to shrinking Internet address space 			
Dec		<ul style="list-style-type: none"> • The Sony PlayStation, which offers over 500 MIPS, debuts • The America Online "Good Times" virus hoax begins • The Third Internet World Conference is held in Washington, D.C., with 11,000 participants 		<ul style="list-style-type: none"> • Netscape Communications Corporation releases version 1.0 of its browser 	
Dec 5		<ul style="list-style-type: none"> • <i>Time</i> names Marc Andreessen one of America's future leaders 			
Dec 9		<ul style="list-style-type: none"> • <i>Disclosure</i> is released in the U.S., incorporates many VR/GUI scenes 			
Dec 14				<ul style="list-style-type: none"> • The World-Wide Web Consortium (W3C) holds its first meeting at MIT's Laboratory for Computer Science (LCS) 	
Dec 22				<ul style="list-style-type: none"> • AT&T buys Interchange Network Corporation from Ziff Communications for \$50 million 	
1995				<ul style="list-style-type: none"> • Michelle Slatalla's <i>Masters of Deception: the gang that ruled cyberspace</i> is published • France's Minitel network has over 6.5 million displays in use with over 25k services • U.S. penetration of TV: 95%, Cable: 60%, VCRs: 91%, answering machines: 57%, CD players: 43%, video game systems: 42%, home computers: 36% 	

	World	Internet	VR	Hypermedia	Computers
1995		<ul style="list-style-type: none"> • Work on TAT-12/13 begins; each cable can carry over 120k conversations • The U.S. video game business reaches \$5 billion • Nicholas Negroponte's <i>Being Digital</i> is published • Douglas Coupland's <i>Microserfs</i> is published • VR5 debuts as a short-lived series 		<ul style="list-style-type: none"> • The Web server market reaches \$5 million 	
Jan		<ul style="list-style-type: none"> • A search of media turns up 1,205 mentions of the word cyber; 57% of Americans don't know what cyberspace means 	<ul style="list-style-type: none"> • Paul Strauss of SGI writes and releases QvLib, a public domain VRML parser 		
Feb		<ul style="list-style-type: none"> • (Spring) Time runs special Welcome to Cyberspace article 		<ul style="list-style-type: none"> • CERN passes its Web responsibilities to INRIA 	
Mar		<ul style="list-style-type: none"> • 14 CompuServe buys Spry for \$100 million • 239 Internet books are commercially available in the U.S. 			
Apr		<ul style="list-style-type: none"> • The Fourth Internet World Conference is held in San Jose, California 		<ul style="list-style-type: none"> • The World-Wide Web becomes the largest service on the NSFnet backbone • The Third International WWW Conference is held in Darmstadt, Germany 	
Apr 3			<ul style="list-style-type: none"> • SGI announces WebSpace, the first commercial VRML browser • WaxWeb is the first Web site to integrate VRML models 		
Apr 24		<ul style="list-style-type: none"> • Worlds, Inc. launches WorldsChat, a collaborative networked environment 			
Apr 30		<ul style="list-style-type: none"> • The NSFnet is decommissioned 			
May		<ul style="list-style-type: none"> • <i>Casper</i> is released, with the first digital performer in a movie 			
May 15			<ul style="list-style-type: none"> • <i>Newsweek</i> runs an article on VRML 		
May 23				<ul style="list-style-type: none"> • Hot Java, a transportable language, is introduced at SunWorld '95 	
May 26		<ul style="list-style-type: none"> • <i>Johnny Mnemonic</i> is the first motion picture to portray cyberspace on the Internet, screenplay by William Gibson 	<ul style="list-style-type: none"> • The official 1.0 specification of VRML is released 		
Jun			<ul style="list-style-type: none"> • The Interactive Media Festival Gallery Tour in Los Angeles makes extensive use of VRML • Tony Parisi and Mark Pesce win Meckler's Market Impact Award for Virtual Reality 		
Jun 1		<ul style="list-style-type: none"> • America Online acquires the Global Network Navigator from O'Reilly and Associates for \$11 million 			
Jun 13		<ul style="list-style-type: none"> • The FCC begins to ration scarce 1-800 numbers 	<ul style="list-style-type: none"> • InterVista releases WorldView 		
Jun 20		<ul style="list-style-type: none"> • AT&T announces \$120 million alliance with BBN 			
Jun 22			<ul style="list-style-type: none"> • Virtus announces support for VRML in its Walkthrough products 		

	World	Internet	VR	Hypermedia	Computers
1995					
Jul		<ul style="list-style-type: none"> WorldsAway, Habitat's sequel, is available on Compuserve Network Solutions, Inc. begins to enforce trademarks on Internet addresses 			
Jul 3		<ul style="list-style-type: none"> <i>Time</i> runs misleading Cyberporn cover story: Over 83% of net images are pornography 			
Aug		<ul style="list-style-type: none"> Pixar's <i>Toy Story</i> is released, the first full-length fully digital movie 			
Dec			<ul style="list-style-type: none"> Mark Pesce's <i>VRML: Browsing and Building Cyberspace</i> is released at SIGGRAPH '95 VRML '95 Symposium, San Diego, California 		<ul style="list-style-type: none"> Fourth International World-Wide Web Conference, Boston, Massachusetts
		<ul style="list-style-type: none"> (Summer) <i>Hackers</i> premieres (Fall) Sandra Bullock stars in <i>The Net</i>, directed by Irwin Winkler 			
1996		<ul style="list-style-type: none"> Clerks TV series debuts 			
			<ul style="list-style-type: none"> (First quarter) Abrams Gentile Entertainment reintroduces the PC PowerGlove for \$120 		
Jan		<ul style="list-style-type: none"> Microserfs TV series debuts 			
May				<ul style="list-style-type: none"> Fifth International World-Wide Web Conference, Paris 	
? to ? - The Post-Information Age					
1997			<ul style="list-style-type: none"> Stewart Brand predicts that computers will disappear into common environments around this time The Synthetic Theater of War's goal to include 100,000 entities in an exercise The National Academy of Sciences predicts artificial realities in classrooms 		
1998			<ul style="list-style-type: none"> VR industry predicts fully immersive systems for \$10,000 		
1999					<ul style="list-style-type: none"> Dataquest predicts 100 millions PCs shipped world-wide, \$185 billion market
2000		<ul style="list-style-type: none"> The Media Lab's Stephen Benton predicts holographic storage, the holophone, and holoprinter Forrester Research predicts a \$644 million Web server market, \$10 billion Internet market, \$6.9 billion interactive market 			<ul style="list-style-type: none"> Intel predicts chip with 400 million transistors running at 250 MHz, over 2 BIPS
2003		<ul style="list-style-type: none"> GTE predicts 60% of American households will have interactive television 			
2005		<ul style="list-style-type: none"> The UN predicts that 50% of the world's population will live in urban areas Given 1995 statistics and growth, there will be more machines than people on the Internet at this time Nicholas Negroponte predicts Americans will spend more hours on Internet than television 			
2010					<ul style="list-style-type: none"> If Moore's law is followed, Eric Drexler says we'll be storing bits as atoms

	World	Internet	VR	Hypermedia	Computers
2010			<ul style="list-style-type: none"> • ATR's goal for the creation of communications with sensory feedback • SRI estimates that VPL technology will be readily publically available 		
2017	<ul style="list-style-type: none"> • Eric Drexler predicts the nanocomputer 				
2019	<ul style="list-style-type: none"> • The year of <i>Blade Runner</i> 				
2020	<ul style="list-style-type: none"> • The estimated year of <i>Snow Crash</i> 				
2021	<ul style="list-style-type: none"> • The UN predicts that 61% of the world's population will live in urban areas 				
2025	<ul style="list-style-type: none"> • Date of <i>Johnny Mnemonic</i> cyberspace 				
2100			<ul style="list-style-type: none"> • UNC at Chapel Hill's estimate for putting whole populations into cyberspace 		

References

Particular pieces of information and/or ideas contained in these media have been reflected in this work, in addition to the timeline references. Extensive virtual reality bibliographies can be found at <http://www.cms.dmu.ac.uk/~cph/VRbib.html> and http://www.hitl.washington.edu/projects/knowledge_base/meta/

A glossary of virtual reality terms can be found at either one of these sites: <http://198.97.199.60/vrterms.html> and <http://alex-immersion.army.mil/vrterms.html>

Adler, Jerry, "Bye-Bye, Suburban Dream", *Newsweek*, May 15, 1995

Apple Computer, *The Basics of Color Design*, Apple Computer, 1992

Aukstakalnis, Steve, *Silicon Mirage*, Peachpit Press, 1993

Azzara, Mike, "Let's Zero In On What 'Open' Really Means", *Open Systems Today*, September 29, 1994, Issue 159, under <http://techweb.cmp.com/>

Barlow, John Perry, "Being In Nothingness", *Mondo 2000*, Number 2

Benedikt, Michael, *Cyberspace: First Steps*, MIT Press, 1991

Bly, Susan A., Harrison, Steve R., and Irwin, Susan, "Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment", *Communications of the ACM*, January 1993, Volume 36, Number 1

Brand, Stewart, *The Media Lab: Inventing the Future at MIT*, Viking, 1987

Bricken, Meredith, and Byrne, Chris M., *Summer Students in Virtual Reality: A Pilot Study on Educational Applications of Virtual Reality Technology*, Human Interface Technology Lab, <http://www.hitl.washington.edu/projects/education/psc/psc.html>

Brooks, Frederick, *The Mythical Man-Month: Essays on Software Engineering*, Addison-Wesley, 1975

Brutzman, Donald P., Macedonia, Michael R., and Zyda, Michael J., *Internetwork Infrastructure Requirements for Virtual Environments*, May 1995, under ftp://taurus.cs.nps.navy.mil/pub/NPSNET_MOSAIC/npsnet_publications.html

Bush, Vannevar, "As We May Think", *The Atlantic Monthly*, July 1945, Volume 176, Number 1, <http://www.csi.uottawa.ca/~dduchier/misc/vbush/as-we-may-think.html>

Catledge, Lara D., and Pitkow, James E., "Characterising browsing strategies in the World-Wide Web", *Computer Networks and ISDN Systems*, Elsevier, April 1995, Volume 27, Number 6

Cavendar, Sasha, "Making the Web Work For All", *San Francisco Examiner*, July 16, 1995

Clickery, Roy, *Data Powers of Ten*, <http://www.ccsf.caltech.edu/~roy/dataquan/>

Cunningham, Donald J., Duffy, Thomas M., and Knuth, Randy A., "The Textbook of the Future", *Hypertext: A Psychological Perspective*, Ellis Horwood, 1993

Curtis, Pavel, *LambdaMOO Help System*, June 1993, <ftp://parcftp.xerox.com/pub/MOO/HelpSystem-02Jun93.txt>

Curtis, Pavel, *MUDs Grow Up: Social Virtual Reality in the Real World*, May 1993, [ftp://parcftp.xerox.com/pub/MOO/papers/MUDsGrowUp.\(ps,txt\)](ftp://parcftp.xerox.com/pub/MOO/papers/MUDsGrowUp.(ps,txt))

Curtis, Pavel, *Mudding: Social Phenomena in Text-Based Virtual Realities*, March 1992, [ftp://parcftp.xerox.com/pub/MOO/papers/DIAC92.\(ps,txt\)](ftp://parcftp.xerox.com/pub/MOO/papers/DIAC92.(ps,txt))

Doerschler, Jeffrey S., and Freeman, Herbert, "A Rule-Based System for Dense-Map Name Placement", *Communications of the ACM*, January 1992, Volume 35, Number 1

Dyson, Esther, "Peripheral Vision", *Release 1.0*, EDventure Holdings, May 26, 1995

Farber, Jerry, *The Student as Nigger*, Pocket Books, 1972

Fish, Robert S., Kraut, Robert E., Root, Robert W., and Rice, Ronald E., "Video as a Technology for Informal Communication", *Communications of the ACM*, January 1993, Volume 36, Number 1

Gelernter, David Hillel, *Mirror Worlds, or The day software puts the universe in a shoebox - how it will happen and what it will mean*, Oxford University Press, 1991

Gibson, William, *Neuromancer*, Ace Books, 1984

Goldfarb, Charles, *A Brief History of the Development of SGML*, SGML Users' Group, June 1990, <ftp://ftp.ifi.uio.no/pub/SGML/Newsire/027.1993-07-21>

Haan, Bernard J., Kahn, Paul, Riley, Victor A., Coombs, James H., and Meyrowitz, Norman K., "IRIS Hypermedia Services", *Communications of the ACM*, January 1992, Volume 35, Number 1

Interview with Ted Nelson, *Cyberspace Report*, June 1995, <http://www.ics.uci.edu/~ejw/csr/cyber.html>

"Is It Live...Or Is It Autodesk? Cyberspace Rising!", *Mondo 2000*, Number 7

Ishii, Hiroshi, Kobayashi, Minoru, and Arita Kazuho, "Iterative Design of Seamless Collaboration Media", *Communications of the ACM*, August 1994, Volume 37, Number 8

Klemets, Anders, "The Design and Implementation of a Media on Demand System for WWW", *Advance Proceedings of the First International World-Wide Web Conference*, May 1994

Lambert, Steve, and Suzanne Ropiequet, editors, *CD-ROM: The New Papyrus*, Microsoft Press, 1986

Landauer, Thomas, Egan, Dennis, Remde, Joel, Lesk Michael, Lochbaum, Carol, and Ketchum, Daniel, "Enhancing the Usability of Text Through Computer Delivery and Formative Evaluation: the SuperBook Project", *Hypertext: A Psychological Perspective*, Ellis Horwood, 1993

Laurel, Brenda, *Computers as Theatre*, Addison-Wesley, 1993

Levy, Steven, *Hackers*, Anchor Press/Doubleday, 1984

McLuhan, Marshall, *The Medium is the Massage: An Inventory of Effects*, Bantam Books, 1967

McLuhan, Marshall, *Understanding Media: The Extensions of Man*, McGraw-Hill, 1964

The Mentor, "The Conscience of a Hacker", *Phrack*, September 25, 1986, Number 7, <http://www.fc.net/phrack/files/p07/p07-3.html>

Michalski, Jerry, "Kids, Education, and Technology", *Release 1.0*, EDventure Holdings, April 24, 1995 and May 26, 1995

Morse, Alan, and Reynolds, George, "Overcoming Current Growth Limits in UI Development", *Communications of the ACM*, April 1993, Volume 36, Number 4

Mukherjea, Sougata, and Foley, James D., "Visualizing the World-Wide Web with the Navigational View Builder", *Computer Networks and ISDN Systems*, Elsevier, April 1995, Volume 27, Number 6

Mullet, Kevin, and Sano, Darrell, *Designing Visual Interfaces: Communication Oriented Techniques*, SunSoft Press, 1995

Najork, Marc-Alexander, *Programming in Three Dimensions*, 1994, http://www.research.digital.com/SRC/personal/Marc_Najork/thesis/index.html

Negroponce, Nicholas, *Being Digital*, Alfred A. Knopf, 1995

Nelson, Theodor H., *Computer Lib; Dream Machines*, Microsoft Press, 1987

Nelson, Theodor H., "A File Structure for the Complex, the Changing, and the Indeterminate", *ACM Proceedings of the 20th National Conference*, 1965

Nelson, Theodor H., "Getting it Out of Our System", *Information Retrieval: A Critical View*, Thompson Book Company, 1967

Nelson, Theodor H., *The Transclusion Paradigm*, May 1995

Nielsen, Jakob, "The Art of Navigating Through Hypertext", *Communications of the ACM*, March 1990, Volume 33, Number 3

Nielsen, Jakob, *Interface Design for Sun's WWW Site*, June 1995, under <http://www.sun.com/>

Nielsen, Jakob, "Noncommand User Interfaces", *Communications of the ACM*, April 1993, Volume 36, Number 1

Norderhaug, Terje, *The Effect of the Media User Interface on Interactivity and Content*, April 1994, http://www.ifi.uio.no/~terjen/pub/MediaUI_Interactivity/940420.html

Norman, Donald A., *The Design of Everyday Things*, Doubleday, 1988

NSFnet Traffic Distribution Highlights, April 1993, <ftp://ftp.merit.net/statistics/nsfnet/1995/nsf-9504.highlights>

The Office of Technology Assessment, *Virtual Reality and Technologies for Combat Simulation*, February 1995, <ftp://ota-bbs.ota.gov/pub/vr.military/hypertext/contents.htm>

Pitkow, James E., and Recker, Margaret M., *Results From The First World-Wide Web Survey*, January 1994, <http://www.gatech.edu/pitkow/survey/survey-1-1994/survey-paper.html>

Pitkow, Jim, and Kehoe, Colleen, *GVU's 3rd WWW User Survey Home Page*, April 1995, http://www.cc.gatech.edu/gvu/user_surveys/survey-04-1995/

Postman, Neil, *Amusing Ourselves to Death: Public Discourse in the Age of Show Business*, Penguin Books, 1985

Postman, Neil, "Informing Ourselves to Death", *Telecom Digest*, January 26, 1994, under <ftp://lcs.mit.edu/telecom-archives/back.issues/> and http://cec.wustl.edu/~cs142/articles/MISC/informing_ourselves_to_death--postman

The Progress and Freedom Foundation, *Cyberspace and the American Dream: A Magna Carta for the Knowledge Age*, August 1994, Release 1.2, <http://www.pff.org/position.html>

Roadbush, Judson, "Bandwidths", *Pixel Vision*, 1994, Number 12

Robertson, George G., Card, Stuart K., and Mackinlay, Jock D., "Information Visualization Using 3D Interactive Animation", *Communications of the ACM*, April 1993, Volume 36, Number 4

Roa, Ramana, Pedersen, Jan O., Hearst, Marti A., Mackinlay, Jock D., Card, Stuart K., Masinter, Larry, Haivorsen, Per-Kristian, and Robertson, George G., "Rich Interaction in the Digital Library", *Communications of the ACM*, April 1995, Volume 38, Number 4

Rucker, Rudy, *The Hacker and the Ants*, Avon Books, 1994

SRI International, *Exploring the World-Wide Web Population's Other Half*, June 1995, <http://future.sri.com/vais/vais-survey.results.html>

Sterling, Bruce, "Computer as Furoshiki", *Mondo 2000*, Number 9

Stephenson, Neal, *Snow Crash*, Bantam Books, 1992

Tognazzini, Bruce, *Tog on Interface*, Addison-Wesley, 1992

Tufte, Edward R., *Envisioning Information*, Graphics Press, 1990

Tufte, Edward R., *The Visual Display of Quantitative Information*, Graphics Press, 1983

Walker, John, *Through the Looking Glass*, September 1988, under <http://www.fourmilab.ch/autofile/www/autofile.html>

Whalley, Peter, "An Alternative Rhetoric for Hypertext", *Hypertext: A Psychological Perspective*, Ellis Horwood, 1993

White, Jan V., *Color for the Electronic Age*, Xerox Press, 1990

Timeline References

Particular timeline events and statistics have been gleaned from these sources, as well as from the ones mentioned above.

Abrams Gentile Entertainment, *PC PowerGlove*, <http://age-inc.com/age/PC.html>

American Federation of Information Processing Societies and Time Magazine, *A National Survey of the Public's Attitudes Toward Computers*, Time, 1973

Apple, Inc. *Corporate Timeline*, June 1995, <http://www.info.apple.com/aboutapple/timeline.html>

Aukstakalnis, Steve, *Silicon Mirage*, Peachpit Press, 1993

Batson, Jay, and DePalma, Donald A., and Trenkle, Ann, *Sizing the Internet*, Forrester Research, April 1995, <http://www.forrester.com/research/ns/1995-ao/apr95nse.html>

Bunnell, David, "The Roots of Personal Computing", *Mondo 2000*, Number 7

Burr, Simon, "Good Times Virus Hoax: The FAQ!", <http://www.tcp.co.uk/tcp/good-times/>

Chandler, Daniel, *Video Games and Young Players*, December 1994, <http://www.aber.ac.uk/~dgc/vidgame.html>

Cotton, Bob, and Oliver, Richard, *Understanding Hypermedia*, Phaidon, 1992

Cringely, Robert X., *Notes From the Field*, InfoWorld, July 19, 1993, <http://www.fc.net/defcon/infoworld.html>

Dataquest Incorporated, *Dataquest Predicts PC Market will Double Over Next 5 Years*, June 1995, under <http://www.dataquest.com/>

Elmer-DeWitt, Philip, "Welcome to Cyberspace: What is it? Where is it? When can we go?", *Time*, Spring 1995, Volume 145, Number 12, under <http://www.pathfinder.com/>

"First Laser Phone Call Zips Across the Atlantic!", *Telecom Digest*, December 1988, Volume 8, Issue 202, <ftp://lcs.mit.edu/telecom-archives/history/tat-8.fiber.optic>

FOIA Jewel: Original Charter of the National Security Agency, February 1990, <ftp://lcs.mit.edu/telecom-archives/history/nsa.original.charter-1952>

Friedrich, Otto, "Machine of the Year: The Computer Moves In", *Time*, January 3, 1983, under <http://www.pathfinder.com/>

Gauvin, Peter, "Palo Alto Wired", *Palo Alto Weekly*, June 7, 1995, Volume XVI, Number 71

Goldsmith, Judith, *A Timeline of the Counter-Culture*, 1993, <gopher://gopher.well.sf.ca.us/11/Community/60sTimeline>

Green, Emily Nagle, and Eichler, Sara H., *Retail's Interactive Future*, Forrester Research, May 1995, <http://www.forrester.com/research/pt/1995/may95pte.html>

H.R. Giger's Necronomicon, Morpheus International, 1991

H.R. Giger's Necronomicon II, Morpheus International, 1985

Hardenbergh, Jan, *timeline.txt*, April 1995, <http://www.oki.com/vrml/timeline.txt>

Hall, Devra, and net.Genesis, *Build a Web Site: The Programmer's Guide to Creating, Building, and Maintaining a Web Presence*, Prima Publishing, 1995

Hayes, Brian, *The Numbering Crisis in World Zone 1*, July 1993, <ftp://lcs.mit.edu/telecom-archives/reports/number.crisis.in.zone.1>

Hauben, Ronda, *Cybernetics, Human-Computer Symbiosis and On-line Communities: The Pioneering Vision and the Future of the Global Computer Network*, January 1995, <ftp://lcs.mit.edu/telecom-archives/reports/pioneer.vision.behind.net>

Historical Computer Society, *A Chronology of Computer History*, <http://www.cyberstreet.com/hcs/museum/chron.htm>

Holaxis Corporation, *A New Dimension*, February 3, 1995, <http://www.hmt.com/holography/holaxis/holaxis.html>

Hughes, Kevin, *Entering the World-Wide Web: A Guide to Cyberspace*, May 1994, Version 6.1, <http://www.eit.com/web/www.guide/>

The Internet Movie Database, <http://www.cm.cf.ac.uk/Movies/>

"Java: The Inside Story", *SunWorld Online*, July 1995, Volume 8, Number 7, <http://www.sun.com/sunworldonline/swol-07-1995/swol-07-java.html>

Johnson, Tony, "Spinning the World Wide Web", *Beam Line*, Stanford Linear Accelerator Center, Fall/Winter 1994, Volume 24, Number 3/4

Kalawsky, Roy S., *The Science of Virtual Reality and Virtual Environments*, Addison-Wesley, 1993

Kaplan, David A., "How to Make a Ghost", *Newsweek*, February 27, 1995

Kumar, Vinay, *The MBONE Information Web*, March 1994, <http://www.eit.com/techinfo/mbone/>

"The Kurzweil Story", http://www.musicpro.com/kurzweil/KZW_Story.html

Lee, John A. N., *Konrad Zuse*, September 1994, <http://info.cs.vt.edu/~history/Zuse.html>

Lendell, Allan, "U-Force Takes Control of Nintendo", *Mondo 2000*, Number 7

Lexikon Services, *Who Was First? A Brief Chronology of Historical Firsts in Computers 1931-1981*, 1995, <http://www.apollo.co.uk/a/Lexikon/who-was-first.html>

Lucente, Mark, *Holographic Video: Timeline of Major Firsts*, <http://www.media.mit.edu/people/lucente/holo/holovideo-timeline.html>

Malkin, Gary, *The Tao of IETF*, RFC 1718, <http://www.ietf.cnri.reston.va.us/tao.html>

"Mandelbrot", February 1995, <http://www-groups.dcs.st-and.ac.uk/~history/Mathematicians/Mandelbrot.html>

McMillan, Kate, *Virtual Reality: Architecture and the Broader Community*, May 1994, <http://www.arch.unsw.edu.au/subjects/arch/specres2/mcmillan/vr-arch.htm>

Meyer, Michael, "Surfing the Internet in 3-D", *Newsweek*, May 15, 1995

The National Academy of Sciences, *Revolution in the U.S. Information Infrastructure*, 1995, <http://xerxes.nas.edu/nap/online/newpath/>

The National Academy of Sciences, *Reinventing Schools: The Technology is Now*, 1995, <http://xerxes.nas.edu/nap/online/tech-gap/welcome.html>

Petry, James, *Buck Rogers in the 25th Century*, <http://www.attes-tel.co.uk:8000/jnp/br25c.html>

Pimentel, Ken, *Virtual Reality: Through the New Looking Glass*, Intel/McGraw-Hill, 1995

Plotkin, Hal, "Invasion of the Cybercensors", *Santa Clara Metro*, June 22, 1995, Volume 11, Number 17

Polsson, Ken, *Chronology of Events in the History of Microcomputers*, June 1995, <http://www.islandnet.com/~kpolsson/comphist.htm>

SCEA Debut at E3, Sony Computer Entertainment, May 1995, http://www.sepc.sony.com/SEPC/E3/PSX_E3Annc/E3Overview.HTML

Sirus, R.U., and Gleason, George, "Do G-Men Dream of Electric Sheep?", *Mondo 2000*, Number 2

Sterling, Bruce, *The Hacker Crackdown*, Bantam Books, 1992, <gopher://gopher.well.sf.ca.us:70/11/Publications/authors/Sterling/hc>

Speeches by William Gibson and Bruce Sterling at the National Academy of Sciences, May 10, 1993, gopher://gopher.well.sf.ca.us/11/Publications/authors/Sterling/Sterling_and_Gibson_at_the_Academy

Sun Corporate Overview, May 1995, http://www.sun.com/cgi-bin/show?corporateoverview/SMI_overview.html

Sutherland, Ivan E., *A head-mounted three dimensional display*, AFIPS Conference Proceedings of the 1968 Fall Joint Computer Conference, 1968, Volume 33, Part 1

Sutherland, Ivan E., *The Ultimate Display*, Proceedings of the IFIP Congress '65, 1965, Volume 2

Thomas, Wes, "How I Created a Media Virus", *Mondo 2000*, Number 2

Times-Mirror Center for People and the Press, *The Role of Technology in American Life*, Times-Mirror, May 1994

United Nations Population Division, Department For Economic and Social Information and Policy Analysis, *World Population Growth From Year 0 to Stabilization*, June 1994, under <gopher://gopher.undp.org/>

Vice-President Gore Speech to Television Academy, UCLA, January 1994, <ftp://lcs.mit.edu/telecom-archives/reports/info.highway.speech-gore>

Watson, Bruce, *Every Computer Type Ever Made*, May 1992, wats@isis.cs.du.edu

White, Jan V., *Graphic Design for the Electronic Age*, Xerox Press, 1988

Wired, Issues 1.1 to 3.08

WWW-Talk and WWW-HTML Mail Archives, From October 1991 to July 1995, <http://gummo.stanford.edu/html/hypermail/archives.html>

Zakon, Robert H'obbes, *Hobbes' Internet Timeline 1.3a*, June 1994, <http://www.amdahl.com/internet/events/timeline.html>

Other Media

In addition to the various media mentioned in the timeline, below are a number of works that are related to the topics discussed in this document.

A number of bibliographies exist that deal with cyberculture; some ones of note are Bruce Sterling's in the directory **[gopher://gopher.well.sf.ca.us/00/Publications/authors/Sterling/](http://gopher.well.sf.ca.us/00/Publications/authors/Sterling/)** - the files are **cyberpunk_library.txt**, **biblio**, and **sfra_june_92.txt**; Anne Balsamo's at **[gopher://gopher.zilker.net/00/bruces/balsamo_biblio](http://gopher.zilker.net/00/bruces/balsamo_biblio)**, and Meng Weng Wong's at **<http://www.seas.upenn.edu/~mengwong/cyber/cgvc.html#1.2>**

Burke, James, *Connections*, Little Brown, 1978

Chandler, Daniel, *The Media and Communications Studies Page*, <http://www.aber.ac.uk/~dgc/media.html>

Marcus, Greil, *Lipstick Traces*, Harvard University Press, 1989

McLuhan, Marshall, *The Global Village*, Oxford University Press, 1989

National Research Council, *Virtual Reality: Scientific and Technological Challenges*, National Academy Press, 1994

Nielsen, Jakob, *Hypertext and Hypermedia*, Academic Press, 1990

Nielsen, Jakob, *Multimedia and Hypertext*, AP Professional, 1995

Norman, Donald, *Things That Make Us Smart*, Addison-Wesley, 1993

Norman, Donald, *Turn Signals Are the Facial Expressions of Automobiles*, Addison-Wesley, 1992

Potter, Beverly, *The Way of the Ronin: Riding the Waves of Change at Work*, Ronin Publishing, Inc., 1988

Rheingold, Howard, *Tools For Thought*, Simon and Schuster, 1985

Rucker, Rudy, Sirius, R.U., and Queen Mu, editors, *A User's Guide to the New Edge*, HarperCollins, 1992

Index

A

Adobe Acrobat 19, 21
 InternetLink 19
 Advanced Design Research Program
 91
 Alvin Toffler 181
 America Online 13, 181
 artificial reality
 origin of phrase 37
 Atari 179
 avatar 65
 avatars 77

B

BADGER 60
 Bar Code Hotel 49
 BioMOO 25
 blossoming file structure 132
Brazil 125
 Brenda Laurel 113, 147
 browser
 modular 18
 Bruce Tognazzini 51
 Bryan Lyles 142
 Butterfly 129

C

Cabri 55
 CCI 94
 CD 134
 CD-ROM 10, 40, 55, 121, 134, 176
 CERN 9, 163
 Clay Graham 152
 collaborative systems 13
 chart 17
 needs 13
 Collage 17
 color
 characteristics 115
 contrasts 117
 cultural meanings 116
 elements 115
 shade model 117
 computer environments
 editing 139

input/output devices 120
Computer Lib 181
 computer monitors 121
 computers
 history 197
 computing environment
 layouts 127
 computing environments
 characteristics 102
 common problems 158
 fonts 130
 merging and division 140
 navigation 136
 page widget 138
 scenarios 143
 Conceptual Design Space 43
 constructs 74
 event hooks 75
 example 76
 content and context 172
 coordinate systems 69
 cross-pollination 39
 cultural acceleration 184
 CutPlane 43
 cybernetics 36
 cyberspace 35
 definition 42
 distribution 45
 encapsulation diagram 39
 history 197
 myths 43
 needs 193
 original definition 35
 protocol 94
 scenarios 97
 second sense 37
Cyberspace - First Steps 35, 79

D

Dark Forces 85
 DataGlove 123
 David Zeltzer 147
 differentiation 169
 Digital Micromirror Device 125
 DIVE 78, 83
 Doom 17
 Douglas Englebart 141

E

ECIT 57
education 53
 goals 56
EduNet 55
encapsulation 37
Enquire 9
Enterprise Integration Technologies
 153
 collaboration space 154

F

Fairmount 156
feely-blank 126
fonts 130

G

Geosphere Project 60
Global Schoolhouse 53
Globewide Network Academy 53
Gopher 20, 36, 90
Graphics Visualization and Usability
 Center 43
Great Wave 181

H

Habitat 44, 78
harpoon 83
head-mounted display 44
 color considerations 118
heads-up display 83
Home Space Builder 83
HotJava 141
htMUD 25
Human Interface Technology Labora-
 tory 43, 55, 62, 124
hybridization 39
HyperCard 10, 21
hypermedia
 history 197
 navigation 109
HyperTalk 76
hypertext 109
 scale of extremes 113

I

information encoding 119

Institute for Research on Learning 55
IntelligentPad 19
interactive television
 problems 175
Internet
 business guidelines 170
 business models 165
 censorship 192
 definition 42
 generation gap 178
 history 197
Internet Movie Database 22
InterNIC 45
Inventor 70
IRC 13, 28, 54, 78

J

J.C.R. Licklider 111
Jakob Nielsen 102
Jaron Lanier 37
Jerry Farber 53
Johnny Mnemonic 151

K

Karl Sims 52
Katsushika Hokusai 181
keyboards 122

L

LEGO 55
Library of Congress 189
light pens 122
liquid architecture 90
LiveBoard 49

M

Marcos Novak 91
Mark Heyer 110
Marshall MacLuhan 181
Marshall McLuhan 37
Marvin Minsky 198
MBONE 142
Media Lab
 Sociable Web project 26
 spherical browser 67
Media Research Lab 135
Media Spaces 49
Meme 76

metaspace 40
 Michael Benedikt 35
 MIME 75
 mirror worlds 62
 Mitsubishi Electric Research Lab 77
 MOO 19
 mouse 123
 MTV 57, 187
 MUD 19
 multimedia 38
 Myron Krueger 37

N

National Academy of Science 57
 navigation and orientation 82

- indicators 85
- maps 85
- modes 84

 Neil Postman 181, 186
Neuromancer 35
 new media 174

- terminology 41

 Nicholas Negroponte 173
 Nintendo 179
 Norbert Wiener 36
 NTT Human Interface Laboratory 49
 nv 17

P

Pad++ 135
 Phoenix 25
 Ping project 23
 pointers 105
 portals 93
 progress acceleration 190

R

Ray Dream Designer 70, 71
 rear-view mirror thinking 40
 Rudy Rucker 126

S

semiotics 107
 Sesame 26
 SGML 162
 SIMNET 44
 simulator sickness 124

- display mismatch 124

field of view 124
 six degrees of freedom 82
 six degrees of separation 110
Snow Crash 50
 Sociable Web project 26
 Sprawl 25
 Starfire 51
 surfing

- origin of phrase 181

T

Ted Nelson 139, 181, 197, 198
 teleport 65
 TerraVision 61
 The 94
The Student as Nigger 53
 three user types 110
 Tim Berners-Lee 9, 140
 transclusion 139
 TurboGopherVR 21

U

Ubique 26
 ubiquitous computing 181
 USA Today 145, 187

V

Vannevar Bush 3
 VIMS 124
 virtual environment systems 37

- applications 47
- avatars 77
- clip objects 71
- color and shadow 87
- coordinate systems 69
- definition 42
- design guidelines 99
- generic objects 90
- history 197
- input/output devices 123
- navigation and orientation 82
- portable set 125
- privacy and permissions 72
- scale of extremes 113
- teleportation 87

 virtual paper 120
 virtual reality

- definition 42
- origin of phrase 37

 Virtual Retinal Display 124

- virtual tool box 79
- Virtual Tourist 60
- Virtus VR 65
- Virtus Walkthrough 71
- vision 114
 - blurriness 120
 - field of view 124
 - recognition in reading 118
 - recognition in virtual environments 119
 - rods and cones 117
 - space and depth perception 119
 - static parallax 120
 - texture gradient 119
- visual programming 111
- VRML
 - browser capabilities 95
 - camera options 87
 - camera paths 88
 - client processing 80
 - editor and modeler features 70
 - floors and gravity 81
 - general information 65
 - home view 87
 - level of detail 89
 - links 80
 - portals 93
 - prototype VRML/HTML system 94
 - proximity and location 80
 - scribble attribute 72
 - texture maps 73
 - tiling 93
 - tools 66
- VRML browsers
 - navigation and orientation 82
- VRweb 82

W

- WAIS 90
- WAXweb 22
- WebMap 68
- WebSpace 83
- webspaces
 - definition 42
- WebViz 68
- WebWorld 21
- William Gibson 35
- WIMP 102
- WOO 21
 - future systems 27
 - general information 25
 - scenarios 31

- world design 144
 - basic styles 146
 - examples 151
 - guidelines 148
- world population 186
- worldmaster 65
- WorldsAway 78
- WorldsChat 79
- World-Wide Web
 - browser navigation 137
 - definition 42
 - history 9
 - problems 162
 - reasons for popularity 15

X

- Xanadu 45
- Xerox PARC 49, 129, 142
 - cone tree 67
- Xerox Star 130

Y

- Yahoo 23

Z

- Zeltzer Cube 147

