Managing Data on the World Wide Web

"State of the Art Survey of Innovative Tools and Techniques"

by

Prasanth Duvvur

Submitted to the Department of Civil and Environmental Engineering in partial fulfillment of the requirements for the degree of

Master of Science
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1995

Author ....
Department of Civil and Environmental Engineering
July 28, 1995

Certified by ....
Professor John R. Williams
Thesis Supervisor

Accepted by ....
Joseph M. Sussman
Chairman, Departmental Committee on Graduate Studies

Managing Data on the World Wide Web
"State of the Art Survey of Innovative Tools and Techniques"

by

Prasanth Duvvur

Submitted to the Department of Civil and Environmental Engineering July 28, 1995, in partial fulfillment of the requirements for the degree of Master of Science.

Abstract

This thesis will focus on the latest, technological innovations and research being conducted on the World Wide Web. Today there are approximately 4, 000, 000 million nodes or servers and over 20 million users on the Internet, with the number of nodes and users growing exponentially. It is estimated that by 1998 the market for Internet services will be between $8-10 billion. This explosion in on-line users has been accompanied by rapid advancements in on-line technologies, especially in the areas of networking, multimedia and Internet security. New protocols for networking, that will provide for increased addressing, enhanced security and greater capacity are in development, some of which are described in the thesis. Video and audio conferencing, on-line publishing and shared workspaces are a few of the new and growing multimedia technologies discussed. The efforts being taken to improve net security as businesses move to sell their wares on-line are described in detail, along with possible monetary standards to perform on-line transactions. The topic of on-line data integrity, a methodology to measure the integrity of a datasource or web site is presented.

Thesis Supervisor: Professor John R. Williams
Title: Associate Professor, Civil and Environmental Engineering

Chapter 1

Introduction

The World Wide Web has captured our imagination and like a wide, wild open frontier beckons us for adventure. The thesis presents the latest trends in on-line data management. Special mention is made of web technologies that promise to have a major impact on the way we view and use the Internet. Advances and innovations in networking technology, network security tools, and multimedia applications, some of which are currently available on the web and some of which are in development are highlighted. The specification for a software tool that will enable the measurement of data integrity on the web is also presented.

1.1 Brief History of the World Wide Web

The World Wide Web a hypertext interface to the Internet was developed primarily by scientists at the CERN (European Particle Physics Laboratory). The web's hypertext interface (via hypertext links) gives Internet users quick in some cases instantaneous access to miscellaneous data residing in web servers across the globe. Prior to the World Wide Web, Internet users had to suffer through bland text-based interfaces, when the web was introduced in 1992, users were able to view graphics, listen to their favorite music, or even watch movies on-line from their computer terminals.

The Internet, the world's largest computer network, which forms the backbone of the web was first conceived about 20 years ago out of a US defense department initiative called the ARPAnet (Advanced Research Projects Agency). The ARPAnet was a network built to withstand natural catastrophes or bomb attacks, in the case of such events the ARPAnet was to function as the primary means of communication. In the ARPAnet model communication always occurred between a source and a destination computer. The network itself was assumed to be unreliable; any portion of the network could disappear at any moment. It was designed to require the minimum of information from the computer clients. To send a message on the network, a computer simply had to put it's data in an envelope, called an Internet Protocol (IP) packet, and "address" the packets correctly. The communicating computers not the network itself-were also given the responsibility for ensuring that the communication was accomplished. The philosophy was that every computer on the network could talk, as a peer, with any other computer. With these assumptions the US was able to develop a working network (the ancestor of the current Internet), and the academic and research users who had access to it were soon addicted.

Demand for networking quickly spread. Internet developers in the US, UK, and Scandinavia, responding to market pressures, began to put their IP software on every conceivable type of computer. It became the only practical method for computers from different manufacturers to communicate. This was attractive to the governments and universities, which didn't have policies saying that all computers must be bought from the same vendor. At about the same time Ethernet LAN (local area network) technology that connected several PC's and workstations became popular, users wanted to connect their entire LANs to the ARPAnet, allowing all computers in the LAN to access the ARPAnet.

Observing the success of the ARPAnet many companies and organizations started building their own networks, of these the NSFNET (National Science Foundation net), made user of five super-computers stationed across the US. The main purpose of this network was to make supercomputers accessible to universities and researchers, an attempt was made to use the ARPAnet as the network for communication, connecting client sites to the supercomputers, but this strategy failed for bureaucratic and staffing problems. In response the NSF attempted to build it's own network, based on the

ARPAnet's IP technology. It connected the centers using 56, 000 bits/second telephone lines, however this strategy failed as well, for it was too expensive to connect every client to the network using telephone lines. Therefore they created regional networks, in each area of the country, universities and other research organizations would be connected to their nearest neighbor. Each chain was connected to a supercomputer at one point, and the centers were connected together. With this configuration, any computer could eventually communicate with any other by forwarding the conversation through it's neighbors.

This solution worked until the network again ran into data overloading, at this point in 1987 the contract to mange the Internet was given to a private party with funding from NSF, an offshoot of which is today's Internet (The Whole INTERNET - User's guide and catalogue - pg., 13). Recently NSF's contract with the network expired, as a result the network is now fully funded by private organizations.

Today the Internet boasts about 4,000,000 computer hosts or nodes (Byte magazine, July 1995, pg., 86) with this number growing every day, efforts are on to build descendants of the Internet that are faster and possess broader bandwidths, to provide for improved multi-media performance and speedier data transfer.

1.2 Goal of thesis

The goal of the thesis is to evaluate innovative World Wide Web technologies, providing the reader with a State of the Art report on the latest developments. The thesis will focus on current trends in

- networking technology
- multimedia
- web security and
- on-line data integrity evaluation.

1.3 Overview

In chapter 2, the latest trends in web browsing are presented, these include techniques to build web servers, from where the world can access user data, methods for writing home pages (with a home page one can piggy-back on other servers), a brief note on html versions 2 & 3 (hyper text markup language), which is the hyper-text language in which the web has been developed, and several web browsers (public domain and commercial) based on html.

In chapter 3, advances in Internet network technology, like the MBone (multicast backbone to the Internet) are presented, with emphasis on improvements in the network to enhance multi-media performance, i.e., video-conferencing etc., high speed networks and their contribution to improving network performance are also discussed.

In chapter 4, the developments in multimedia packages and tools are presented, with an emphasis on research being done at MIT's Laboratory of Computer Science on the VuSystem for distributed multimedia.

In chapter 5, tools and methodologies that enhance network security are discussed, a mention of on-line commerce and the efforts being taken to secure transactions on the Internet are also described.

In chapter 6, the specification for a software tool that will be used to evaluate data integrity at web sites is put forth, the tool will enable users to pick and choose data sources based on their overall data integrity, a scenario in which this tool may be useful is also presented.

Chapter 7, highlights the future areas of research for web computing, and reiterates some of the key conclusions of my research on the World Wide Web.

Chapter 8 contains the glossary, which readers might find useful, as the web is notorious for spewing out acronyms at alarming rates.

Chapter 9 contains a literature survey on select research articles, and chapter 10 closes with a list of references.

Chapter 2

Browsing the World Wide Web

Prior to 1992, the Internet was a haven to academicians, researchers and college students, who were accustomed to text-based, menuing systems and data retrieval. With the arrival of the World Wide Web and hypertext linking, the Internet became more maneuverable and accessible to the public at large. Mosaic the first hypertext browser designed to traverse ('surf') the web came into being in 1994. It was borne of a research project at the University of Illinois at Urbana Champaign. In what follows the different pieces of the puzzle that go toward making 'surfing the net' the pleasurable experience that it is are discussed.

2.1 HyperText Markup Language (HTML)

The hypertext markup language is a simple data format used to create hypertext documents that are portable from one platform to another. HTML has been in use by the World Wide Web global information initiative since 1990, when it was introduced by Tim Berners-Lee at CERN as part of the World Wide Web project. Three versions of the language have been released up-to date. HTML is an application of ISO Standard 8879:1986 information processing text and office systems, standardized markup language (SGML). The HTML document definition is a formal definition of the HTML syntax in terms of SGML. HTML refers to both the document type and the markup language for representing instances of this type. It can be used to represent, Hypertext news, mail, on-line documentation, and collaborative hypermedia, such as, menus of options, database query results, simple structured documents with inlined graphics and hypertext views of existing bodies of information.

An HTML instance is like a text file, except that some of the characters are interpreted as markup. The markup gives structures to the document. The instance represents a hierarchy of elements. Each element has a name, some attributes, and some content. Most elements are represented in the document as a start tag, which gives the name and attributes, followed by the content, followed by the end tag. There are also tags to mark text as headings, paragraphs, lists, quotations, emphasized and so on, there are also tags to include images within documents, for including fill-in forms that accept user input, and, most importantly, for including hypertext links connecting the documents being read to other documents or Internet resources, such as WAIS (Wide Area Information Service) databases and anonymous FTP sites. It is this last feature that allows users to click on a string of highlighted text and access a new document, an image, or a movie file from a computer thousands of miles away. The HTML specifies where the document file is to the client computer by means of a URL (Uniform Resource Locator). The following is a sample HTML document:

```
<HTML>
        <TITLE>
            sample HTML
        </TITLE>
        <H1>
            An example of structure
        </H1>
            Here's a typical paragraph.
        <P>
```

2.1.1 Hyperlinks

In addition to general purpose elements such as paragraphs and lists, HTML documents can express hyperlinks. A hyperlink is a relationship between two anchors, called the head and the tail of the hyperlink. An anchor is a resource such as an HTML document, or some fragment of, i.e., view or portion of a resource. Typically, the user activates a link by indicating the tail of the link with the head of the link being presented as a result. Anchors are addressed by Uniform Resource Locators (URL).

2.1.2 HTML Graphics

This section describes graphical tools that can be used to manipulate HTML graphics, and how external graphical files can be used to create and use interactive maps.

- *Mapedit for developing Active maps*:

An active map is also sometimes referred to as a clickable image. It is an in-line image (an image merged with displayed text) in an HTML document. An area of the image can be selected, usually by clicking with a mouse. The coordinates of the image that have been selected are sent to a program which can then process the information. An active map can be used to provide a graphical menu, in which selecting a menu option will retrieve a specified HTML document. Active maps can also be used in developing teaching and learning software. For example a medical student could be asked to click on an area of an x-ray which shows a cancerous growth. If an incorrect area is selected a HTML document giving further information can be displayed.

Mapedit is an editor for creating image map files. Image map files are a feature of NCSA and CERN servers which enable users to turn a GIF image into a clickable map by designating areas using polygons and circles within the GIF and specifying a destination URL for each area.

- *Paintshop Pro:*

This tool can be used to convert file formats, to reduce color depth and to convert colors.

2.1.3 Forms

A form is a template for a form data set and an associated method and action URL. A form data set is a sequence of name/value pair fields. Names associate variable names to the data associated with this input element. A name consists of a letter followed by letters, digits, periods or hyphens, the length of a name is limited to 72 characters. The names are specified on the name attributes of form input elements, and the values are given initial values by various forms of markup and edited by the user. The resulting form data set is used to access an information service as a function of the action and the method. An HTML user agent begins processing a form by presenting the document with the fields in their initial state. The user is allowed to modify the fields, with constraints on the field type etc. When the user indicates that the form should be submitted, the form data set is processed according to it's method, and action URL. Forms can be used to query data bases residing on server machines, or post email to web server masters in standard formats etc.

2.1.4 HTML tools

Several tools exist to convert from various text formats to HTML, following are a few samples, LaTex to HTML, Framemaker interfaces, mail archives to HTML (make mail archives available on the web). Following are a list of HTML editing and authoring tools,

- *tkHTML*, based on tk/tcl
- *tkWWW*, also based on tk
- *Phoenix*, WYSIWYG editor for the X window system
- *ASHE*, a simple editor in c for Motif
- *asWedit*, a text editor for the X window system and Motif
- *curl*, a tool for automatically generating links between HTML pages based on a contents file
- *webify*, a tool running on UNIX to convert postscript to browsable web pages
- *Imagizer*, a tool to help develop web pages with images
- *Symposia*, web editor and browser for MAC, PC and UNIX

Tools to check for dubious HTML are as follows, *HTML Check Toolkit*, to validate HTML, *weblint*, a perl script to check HTML, *htmlchek*, an awk script to check HTML, *Arena*, comments bad HTML. For generating standards from HTML, plain text, LaTeX etc. filters are available. For analyzing and maintaining HTML databases, the HTML analyzer can be used. For setting up public World Wide Web access, the telnet server (for anonymous users to log into a web client) and mail robot (a program to return any information in the web by electronic mail) are available.

2.2 HTTP (HyperText Transfer Protocol)/Gateways

The HyperText Transfer Protocol, a new Internet protocol is designed specifically for the rapid distribution of hypertext documents. HTTP is a client-server protocol, in that, the client program running on the user's machine issues a message, requesting services from a program running on the server machine, which then responds by sending back a message to the client. At the simplest level, HTTP servers act much like anonymous FTP servers, delivering files when clients request them. However, HTTP servers support additional features: i) The ability to return to the client not just files, but also information generated by programs running on the server. ii) The ability to take data sent from the client and pass this information on to other programs on the server for further processing.

These special server-side programs are called gateway programs, because they usually act as a gateway between the HTTP server and other local resources, such as databases. Just as an FTP server can access many files, an HTTP server can access many different gateway programs; in both cases, users can specify which resource (file or program) they want through a URL. The interaction between the server and these gateway programs is governed by the Common Gateway Interface (CGI) specifications. Using the CGI specifications, a programmer can easily write simple programs or scripts to process user queries, query databases, make images that respond to mouse clicks, and so on.

2.3 Web Servers

A web server, like the FTP daemon, is a program which responds to an incoming TCP connection and provides a service to the caller. There are many varieties of World

Wide Web servers to serve different forms of data. In what follows World Wide Web CERN UNIX server will be described.

The benefits to using a UNIX machine for a server as opposed to a Macintosh or PC compatible machine are as follows, the UNIX operating system is designed to run many simultaneous process so that a well-designed UNIX server can almost effortlessly respond to many simultaneous HTTP service requests. Also UNIX is designed to isolate the server and user processes from the management level of the operating system, which ensures reliable operation even under heavy loads. Finally, the powerful HTTP servers from CERN and NCSA (National Center for Supercomputing Applications) were originally written for UNIX machines and are generally the most full-featured servers available. The source code for these servers is also publicly available.

2.3.1 CERN server

CERN WWW Server HTTPD daemon is a generic, full featured server for serving files using the HTTP protocol. This is a TCP/IP based protocol running by convention on port 80. Files can be real or synthesized, produced by scripts generating virtual documents. It handles clickable images, fill-out forms, and searches etc. The CERN server is fully CGI/1.1 compliant, and comes with several useful CGI applications, including a program for handling active images. It also supports file and directory access control and encrypted user-authentication.

The CERN server can also act as a caching proxy server, which is ideal if users want to run a server and also access other servers through an Internet firewall. A firewall is used to separate a local network from the outside world. In general, a local network is connected to the outside world by a gateway computer. This gateway machine can be converted into a firewall by installing special software that does not let unauthorized TCP/IP packets pass from inside to outside, and vice versa. Users on the local network, and inside the firewall can be given access to the outside world by using the SOCKS package or by installing the CERN HTTP proxy server on the firewall machine. Other available UNIX servers are,

- *CL-HTTP*, which is an HTTP server written using the Common Lisp Object System (CLOS). It runs on Symbolic LISP machines under Genera 8.3. CMS HTTPD, is written for IBM mainframe servers

- *DECTHREAD HTTP* server, is a native VAX/VMS HTTP server that uses the DECthreads model for multithreaded processing. As a result, this server should be significantly faster than the straightforward port of a UNIX server.

- *GN Gopher/HTTP* server, is a single server supporting both the Gopher and HTTP protocols. It allows both protocols to access the same data so that GN is extremely useful for Information Services managers making a transition from a Gopher based system to an HTTP-based one. GN, is a by product of the GNU foundation.

- *GWHIS HTTP* server, is a commercial server package distributed by Quadralay Inc. It includes special components for data encryption, user authentication, and gateways to special database software.

- *Jungle*, is a server written in tcl/tk.

- *NCSA HTTPD*, the NCSA public domain server is fully CGI-compliant, and supports access control through encrypted user-authentication and domain name restrictions. The distribution package comes with several useful gateway programs, including a program for handling active images. It also permits executable HTML documents which allows users to create parsable HTML documents that include other documents on the fly. This is often called the server-side include feature. The server can also be configured to allow users to have executable CGI programs in their own home directories.

- *PLEXUS*, is written in perl language. It is a public domain server designed to be fast, extensible, and easy to use.

2.4 Web Clients

Web clients allow users to access the Internet from their personal computers, usually all that is needed is a high-speed Internet connection and the client software. Outlined below are sample clients, with emphasis on web browsers. The Mosaic browser developed by the University of Illinois at Urbana-Champaign, will be highlighted.

2.4.1 Browser development/line-mode browsers

The team at CERN implemented a line-mode browser. This is the lowest common denominator among browsers, and can be used from almost any kind of terminal. In this browser, the screen is formatted in ASCII text. Links are numbered and appear inside brackets. To follow a link, users enter the number of the link at the prompt at the bottom of the screen. With these early browsers, World Wide Web had reached proof-of-concept stage when the first versions of Mosaic became available in the spring of 1993, but it had not achieved widespread use. While multiple clients existed, none of them suggested an easy-to-use interface that lets users click on a link to navigate the Web, as well as the ability to display graphics. Mosaic made the Internet accessible to a broader group of users.

2.4.2 Mosaic

Mosaic is an Internet navigation and data retrieval tool. It allows users to access network information with the click of a mouse button. Mosaic is capable of accessing data from World Wide Web servers (HTTP), Gopher servers, FTP servers and Usenet News servers (NNTP). Mosaic can also access other data services through gateway servers.

These services provide search capabilities in database environments such as Archie, WAIS and Veronica. Mosaic provides transparent access to these information sources and services. Mosaic also includes the following features, display of plain text, rich text (postscript), and hypermedia, in-line graphics, images, customizable graphical user interface, global history of information space navigation - tracking where a user has been, quick access to important or frequently used documents via a personal "hotlist", search capabilities within a documents and over the Internet, text and voice annotation for documents anywhere on the Internet, full TCP/IP based communications support, easily extendible to arbitrary viewers or other data formats.

2.4.3 Multimedia through Mosaic

Mosaic can be used to play movies, music, look at full color graphics, run scientific animations and models, display 3D graphics and more. But Internet multimedia can sometimes be frustrating and is almost always time consuming. The problem is that these media files can be quite large. A one minute MPEG (Motion Pictures Expert Group, MPEG denotes a standards committee, a method of file compression, and a graphics file format.) movie can be a megabyte or more, a three minute song might be four or five megabytes, large full color graphics are typically half a megabyte or so. In order for Mosaic to play multimedia, certain software programs are needed, there are different programs for audio, video, graphics etc.

When Mosaic attempts to launch a media file that it can't execute by itself it calls on one of these programs. When a user clicks on a link to a document that isn't HTML or plain text, like an image or a sound file, Mosaic tries to use an external program to display the image or play the sound. The default configuration expects to find external programs like,

- *xv*, to display graphical images
- *showaudio*, to play audio files
- *mpeg_play*, to display movies
- *ghostview*, to display postscript files
- *xdvi*, to display DVI files
- *xwud*, to display Xwindow dumps
- *metamail*, to display messages

2.4.4 Other UNIX Web Clients

Discussed below are other web clients with which users can browse the Internet,

- *Batch mode browser*, is a perl program that operates in a batch-mode. It allows users to retrieve a document via any of the World Wide Web protocols simply by specifying the document URL. Useful for downloading web documents, or for use in automated processing, such as UNIX cron (UNIX clock daemon) jobs.

- *Chimera*, is a UNIX based, X-windows based browser. Chimera uses the public domain Xaw Athena widget set, so it can be compiled almost anywhere. A C compiler and the standard X11 libraries are needed for this.

- *EMACS-W3*, is a browser for the emacs text editor.

- *LYNX*, is a full-screen hypertext browser that can run on any dumb terminal. It uses the arrow and tab keys and single letter keycodes to navigate around the document and for executing program commands.

- *MIDASWWW*, is based on the Midas hypertext library, it is an elegant graphical browser supporting inlined images and clear marking of hypertext links.

- *TKWWW*, is a browser based on tcl/tk toolkits, it's most interesting use is as an HTML editor.

- *VIOLAWWW*, is an X-Windows based browser based on Viola interactive media scripting language/toolkit.

- *Agora*, an email based browser, modeled on the line-mode browser.

2.5 Searching the Web

Several searching interfaces exist on the web, when using these search interfaces it is useful to note that the user is not searching the documents on the web itself but databases of these documents. These databases may include information for each document such as, title, location (URL), links, keywords. All this data is gathered off the web and organized in a database, with links pointing back to the original documents. When the user searches for a word or phrase, a program is initiated on a remote computer that knows how to query and extract information from these databases. If the search is successful, the program on the remote computer returns a web document that contains links to the original document. Many of the searching databases are constructed and updated by programs called "robots" or "spiders" that continually travel the web finding documents. As they encounter a document, they record information about it, this is then used to update the database.

Below are listed a few of the latest search servers on the web,

- *The Wanderer*: Developed using perl. The Wanderer travels the web searching for web sites. It does a breadth-first search of the web, looking for and including in the queue all the URLs contained in every document it encounters. These are all listed by IP address, host sites, or country.

- *CUSI*: This is forms based interface to many searchable indexes, as well as databases including catalogues, phone books, dictionaries, and technical documents. It's relatively easy to use, and users find it helpful to have so many of these searches available in one document.

- *ALIWEB*: ALIWEB (Archie like Indexing for the Web) is a distributed indexing system. Because it is modelled after Archie, a program that maintains a database of software programs listed in public archives on the Net. ALIWEB's index database is generated from descriptions, stored as files on the servers that contain the documents being indexed. This means that as long as publishers keep their own local descriptions up to date, ALIWEB has the potential to be of high quality. However, because it doesn't index every document on the web, it's not a good place to look for instances of specific words or phrases.

- *CUI W3 Catalog*: The CUI searcher is a database of indexes to several popular lists and resource guides.

- *WebCrawler*: The WebCrawler is a robot that keeps an index of the contents of all the documents it comes across.

- *EINet Galaxy*: This is a combination hierarchical subject catalog and searchable database. This server, unlike others, grabs all the documents it can find (not just HTML) and indexes them in a WAIS (Wide Area Information Service) server. Thus users get as many hits as they ask for, after it completes a search of the web.

- *World Wide Web Worm*: The Worm is a database of Web document titles, URLs, and cross-references. It's power lies in the detailed level of searching that can be done, in addition to searching for words or phrases in a document title, users can search for specific URLs or patterns within URLs. The Worm is run weekly on selected files known to change regularly. In these runs it ignores any files it has already read in a previous run.

In chapter 3, the latest trends in Internet network technology are discussed.

Chapter 3

World Wide Web Network Technology

Since it's inception 20-30 years ago the backbone of the Internet has undergone many changes, as the number of computers or nodes and users burgeon, new and more efficient networks have replaced congested and slow ones. About 20 years ago NSF's network connected super-computer centers using 56 kilobits/second telephone lines, till recently in 1991 the backbone of the network was a 1.5 megabits/second, T1 circuit (T1 is a Standard for digital transmission over phone lines. Capacity of the Internet, T1=1.5 megabits/second). Today the Internet uses a backbone that operates at 45 megabits/ second, or a T3 circuit, however even this capacity is proving to be insufficient for efficient multimedia and large data transfers, e.g., video conferencing. In what follows the latest trends and technologies in Internet, network technology will be described, preceded by an explanation of some network basics.

3.1 Network Basics

Outlined below are some of the common terms and concepts one will come across when using the Internet.

3.1.1 Domains

An address (e.g., email address) on the Internet is of the form username@location.domain, where username identifies the user, and location the institution where the computer resides. Domain identifies the type of the institution, i.e., commercial (com), educational (edu), government (gov), military (mil), gateways and other administrative hosts for a network (net), private organizations (org). Each country also has it's own top-level domain, for example Australia (au) or Canada (ca).

3.1.2 Internet Numbers

Every single machine on the Internet has a unique address, called it's Internet number or IP address. It's actually a 32-bit number, but is most commonly represented as four numbers joined by periods ('.'), like 251.31.254.130. This is sometimes called a dotted quad; there are literally thousands of different possible dotted quads. The ARPAnet originally only had the capacity to have up to 256 addresses. In the early eighties, it became clear that things would fast outgrow such a small limit, thus the 32-bit addressing method was born, freeing thousands of host numbers. Each piece of an Internet address (like 251) is called an octet, representing one of four sets of eight bits. The first two or three pieces (e.g., 251.31.254) represent the network that a system is on, called it's subnet. To apply for IP names and domains, an application must be filed with the Network Information Center (NIC).

3.2 vBNS (Very-High-Speed Backbone Network Service)

On April 30, 1995, the federal government (represented by NSF) stopped funding the Internet, as the NSF shifted it's funding to a new experimental network called the vBNS. The vBNS will exist primarily for research rather than commercial operations. Research on vBNS will focus on broadband, internetworked technologies and services. The core of the project will focus on improving speed and scaling of the Internet and it's underlying technologies. It will provide a 155 megabits/second test-bed for new network applications, and in time, the test-bed will accommodate speeds of 622 megabits/second. The NSF anticipates that the vBNS will upgrade to 2.5 gigabits/second in 1998. The vBNS will have network-access points like those in current backbones and will connect five supercomputer centers across the country. The vBNS will test high-speed router and switching technologies, such as ATM (asynchronous transfer mode) and frame relay (Frame relay is a transmission standard for sending data over public or private leased phone lines.). Other technologies expected to be tested include packet flows, a technique that allows packets of data to be sent from one source to multiple destinations (traditional packet-switched networks are designed for two connection points a sender and a receiver).

Packet flow technology may be critical for multi-casting of multimedia data across networks. Like a TV station broadcasting a program to millions of homes, multicasting broadcasts multimedia programs to hundreds of networked computers. The vBNS will also be a test-bed for IPng (Internet Protocol next generation), or IP version 6, the updated internetworking protocol. IP6 will offer expanded addressing and simplified packet routing, message handling, and improved security features.

3.3 IP Next Generation (IPv6 or IPng)

IPng is a new version of the Internet Protocol, designed as a successor to IP version 4. IPng was designed to take an evolutionary step from IPv4. It was not a design to take a radical step away from IPv4, functions which work in IPv4 remain. The changes fall primarily into the following categories,

Expanded routing and addressing capabilities

IPng increases the IP address size from 32 bits to 128 bits, to support more levels of addressing hierarchy and a much greater number of addressable nodes, and simpler autoconfiguration of addresses. The scalability of multicast routing is improved by adding a scope field to multicast addresses. A new type of address called an "anycast address" is defined, to identify sets of nodes where a packet sent to an anycast address is delivered to one of the nodes. The use of anycast addresses in the IPng source route allows nodes to control the path in which their traffic flows.

Header format simplification

Some IPv4 header fields have been dropped or made optional, to reduce the common-case processing cost of packet handling and to keep the bandwidth cost of the IPng header as low as possible despite the increased size of the addresses. Even though the IPng addresses are four times longer than the IPv4 addresses, the IPng header is only twice the size of the IPv4 header.

Improved support for options

Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing new options in the future.

Quality-of-service capabilities

A new capability is added to enable the labeling of packets belonging to particular traffic flows for which the sender requests special handling, such as non-default quality of service or real-time service.

Authentication and privacy capabilities

IPng includes the definition of extensions which provide support for authentication, data integrity, and confidentiality. This is included as a basic element of IPng and will be included in all implementations.

The IPng protocol consists of two parts, the basic IPng header and IPng extension headers.

3.4 Internetworking with ATM-based Switched Virtual Networks

In the past, internetworking was realized by using bridges and routers. In fact the lack of intelligence in bridges and their weakness for controlling broad-cast traffic resulted in the creation of routers. With the growth of networks, multiprotocol routers were born to interconnect LAN segments. They can control access permissions, broadcast traffic, and traffic flow between segments. Routers added more security and stability to networks, but now they are too slow to handle the heavy load of router-based backbones. Bandwidth bottle-neck, the high and increasing cost of network operation, and increasing complexity of network management of router-based networks are also attributed to slow routers. Traditional approaches to this problem were all based on shared media architecture. Switching technology described below is a more likely solution to this problem.

Improvements in hubs, bridges, and routers have increased the complexity of internetworking in a way that adding, changing, and moving of the endstations in the network results in unacceptable operation costs. Network growth and new applications require a change in paradigm. Switch-based internetworking provides high performance and less complex networks that have lower operation costs. The switch based solution includes two approaches, LAN switching and ATM switching. LAN switching introduces the switching of the existing packet-based LAN technologies (such as Ethernet, Token Ring, and FDDI). The ATM is a cell-based switching technology based on a connection oriented approach. The combination of switching, traditional hubs, and router technologies is the new approach to internetworking.

Switching is the key element of internetworking and routing is used only when necessary. Internetworking provides connectivity over LAN and WAN. In traditional LANs, the forwarding is based on MAC layer or Network layer. For Virtual LANs, the boundaries of virtual LANs define the routing requirements. In the case of ATM networks, the routing decision is made for the first cell in a flow, and a connection is established, then the rest of the cells in the flow are sent via simple switches that forward the cells based on their virtual connection identifier (VCI) and virtual path identifier (VPI).

Switching brings speed to the backbone and virtual networking removes the bottleneck of physical LANs.

Asynchronous Transfer Mode (ATM) is a scaleable solution that offers high speed, high performance, high availability and simplified administration. ATM provides dedicated bandwidth, different classes of service that can support multimedia applications, bandwidth or time sensitive applications, and high traffic client/server configurations. This technology can solve the problem of congested router-based backbones. The endstations on different LAN segments are connected via high speed switches without intermediate hops.

Virtual networks can be realized in different ways. Virtual LANs are a layer 2 switched virtual network. The ATM LAN emulation can be realized by bridging LAN segments. It can also be realized using multiprotocol routing, although the corresponding standards are still to be defined and finalized. However, many implementations that realize routing based on proprietary solutions are already in place. When the number of LAN segments increases, ATM can be used as a backbone. High performance servers can be connected to the ATM backbone, and WAN connection can be also realized via ATM. This is a complete solution that provides the network with high capacity, low and predictable delays, and easy administration.

3.5 Intelligent Caching for World Wide Web Objects

Statistics collected on the NSFNET backbone were indicating that web traffic would exceed that of any other protocol around May 1995. In March of this year, overall traffic fell off as the backbone was being dismantled. Nonetheless, web byte counts have caught up with FTP and together they account for half of all traffic on the NSFNET. As the web continues it's exponential growth, many sites and network links are experiencing severe overload. Host machines at popular sites become overloaded because a large percentage of clients always request documents from the source. Another cause of load problems is less-than-optimal server software designs.

A partial solution to network and server overload is the use of network caches. Studies, simulations, and real-world experience have shown that web caches can significantly reduce wide area network bandwidth and improve latency. The hypertext transfer protocol (HTTP) functions much like those which came before it (i.e., FTP, Gopher). The user runs a web client, also known as a browser. Popular clients are NCSA Mosaic, Netscape and Lynx (described in chapter 2, pg. 19). Organizations publish information by running a HTTP server. Clients make a TCP connection to servers and request a named object. Uniform Resource Locators (URL's) are the names which are used to fully identify objects throughout the Internet.

If the requested object exists it is delivered to the client and the connection is torn down. Requests need not always be made directly between the client and server; a proxy can be used in between them. The client makes it's request to the proxy. The proxy forwards the request to the server and then relays the resulting data back to the client. Since a proxy is relaying requests and responses between clients and servers, it makes sense for the proxy to cache some of the data it has handled. A subsequent request for cached data can be given directly to the cache without connecting to the server. This results in a savings of wide area network bandwidth and usually an improvement in latency to the user. In general it is possible to cache objects which are retrieved with the HTTP, FTP and Gopher protocols.

Of these, FTP is the most difficult because the protocol is more complicated and oriented towards interactive operation. Keeping a set of cached objects up-to-date is a difficult problem. Servers are stateless so there is no mechanism for them to push updates to clients. Instead clients must occasionally poll the servers for changed objects.

3.5.1 Trade-offs of Caching

The three big advantages of caching web objects are reducing i) latency, ii) network bandwidth, and iii) server load. Another positive side-effect is the opportunity to analyze an organization's usage patterns. On the down side, a proxy introduces a single point of failure. It also requires additional resources (CPU, disk, personnel) and oversight. There is always the chance that a user will be sent stale data from the cache. A disadvantage for servers is that they do not log cache hits.

Caching is effective if the following conditions apply, first, retrieving from the cache must be fast, or faster, than from the original source. This implies that cache placement in the network is important. It should be near an organization's connection point to the Internet. Second, on average an object should be requested again while it exists in the cache. If an object changes more frequently than it is accessed, caching would be pointless. Cached objects should have long lifetimes relative to their frequency of access. The HTTP server developed by CERN can also function as a proxy and a cache. Indeed, it is the original web proxy and has set a number of de-facto standards for proxying and caching on the web.

3.6 Improving Network Performance for Multimedia

In the following paragraphs the advances in multimedia network technology are discussed.

3.6.1 Multicast Backbone for Multimedia

The MBone (multicast backbone) is intended for distributing multimedia programs especially video and audio presentations to multiple Internet sites, using multicast packet technology. The MBone operates at relatively low bandwidth, ranging from 56 Kbps to 1.5 Mbps. To receive multicasts from a multimedia source, the client system must support IP multicasting. The service provider or gateway server must also use a multicast router to send multimedia signals (e.g., video conferencing signals) to Mbone capable endpoints by creating tunnels. These tunnels use the routed multicast-routing daemon to channel through nonmulticast routers to the intended endpoints. Because conventional network routers can't cope with multi-cast packets, broadcasters encapsulate the packets using a conventional IP header so that the routers see the packets as normal. Once on the other side of the router, the header is stripped off, and the packets return to multicast form. In the future, some commercial routers will incorporate both conventional and Mbone capabilities.

3.6.2 VuNet - A Network Architecture for a Distributed Multimedia System

The VuNet is a gigabit-per-second local area ATM network built in the context of the ViewStation distributed multimedia project. The ViewStation system takes a software-intensive approach to multimedia where real-time data is processed at the application level and not in specialized hardware. The VuNet hardware has been specifically designed to support this software model. In particular, the network interface has been designed for ease of use. A key feature of the VuNet is that the more sophisticated functions typically found in ATM networks have been pushed out to the edges of the system to become the responsibility of the clients in exchange for simplicity and efficiency. The design of distributed multimedia systems has evolved over the past five years. Early systems were characterized by an analog audiovisual system and a computer system operating in parallel. The analog system handled the communication, presentation and recording of the multimedia information and was controlled by the computer system. With advances in digital IC technology, systems have appeared in which the handling of multimedia information has moved from the analog to the digital domain. The move to digital makes it possible to handle multimedia and traditional computer system information using common processor and communication resources. However, since, in general, the current generation of workstations and networks lack the performance levels to handle video, little sharing of resources occurs.

Multimedia information is again handled in dedicated hardware over a parallel system. The general practice in both the analog and digital systems is for the applications to never touch the multimedia data but to merely control dedicated hardware which processes it. Researchers have used the phrase 'traffic cop' to describe the role of these applications. Since both types of systems use specialized hardware to handle multimedia information, their functionality is restricted to the capabilities of the hardware. Current implementations of these systems have focused on teleconferencing type applications and are limited to those.

According to the developers of the VuNet architecture, the next step in the evolution of multimedia systems involves allowing multimedia information to reach the application. This means channeling the multimedia information to the workstation processor rather than bypassing it with specialized hardware. This will allow a greater variety of multimedia applications that go beyond teleconferencing. The main argument against having the application touch the multimedia information is that workstations cannot handle video. However, this is a short term problem since the software intensive approach to multimedia allows easy migration to faster workstation platforms as they appear.

The main purpose of the ViewStation project is to build a local-area distributed multimedia system. The project's hardware substrate is known as the VuNet. The VuNet is comprised of network components and host interfaces to workstations. The VuNet interconnects workstations and custom network-based video processing resources such as a video capture board and an image processing system.

Overall, the ViewStation project focuses on getting real-time data such as voice and video from the network all the way to the application. From the software end, a socket interface is provided for the video data, and all video data is treated as any other network data. Since the ViewStation takes a software-intensive approach to multimedia, the VuNet and custom video hardware were designed to provide efficient support for software driven handling of multimedia streams. In addition, the hardware has been adapted to the needs of the ViewStation programming model. The following issues motivated the design of the hardware infrastructure,

- *Transparent communication*: Network communication should be as transparent as possible to the application. To this end, a high speed ATM network meets the need for low latency network communication. Gigabit-per-second speeds were targeted because of the bandwidth needs of video applications and the technology for such speeds was ripe. Another important requirements is an easy to use client interface so that workstation network communication is efficient.

- *Bursty traffic*: Applications based on the ViewStation programming model process data in bursts. Therefore, sources of video and audio information must be able to provide their data in bursts and the VuNet and it's hosts must support bursty traffic.

- *Graceful degradation*: A key aspect of the ViewStation programming environment is the support for graceful degradation of audio and video processing as workstation resources become scarce. One such example of this is a temporary decrease in frame rate in a video viewer application as a disk read occurs. This requirement is much more of an issue for video due to the data rates involved. Therefore, VuNet based video sources must provide a mechanism for throttling their frame rate and the VuNet must allow for efficient cross-network control of video sources.

The VuNet has been deployed in project member's offices and is currently interconnected at 500 Megabits/second. The switches currently run at 700 Megabits/second but have been tested internally at 1.5 Gigabits/second. The links are also rated to run at 1.0 Gigabits/second. The current bottleneck in the system is the programmed I/O host interface, which provides an 80 Megabits/second throughput, with an actual useful throughput of 20 Megabits/second.

A suite of applications have been installed and the software toolkit enables simple programming of intelligent video applications. Development of these applications continues with an emphasis on applications which can make decisions based on the content of the video data. The VuNet provides a network ideally suited for intelligent video applications. The ViewStation project, as a whole, provides scalability of video in software, and provides a high speed ATM desk area network where video is treated like other network data.

In the next chapter advancements in on-line multimedia technologies are described.

Chapter 4

Multimedia on the World Wide Web

Through the web users can conduct point-to-point video conferencing, publish on-line material, replete with live video and sound, watch movies or manipulate 3D graphics. Outlined below are some of the latest tools and trends in multimedia.

4.1 VuSystem

The VuSystem is a programming system for the dynamic manipulation of temporarily sensitive data. It is designed to run on any general purpose UNIX workstation running the X Window System, requiring no special real-time facilities. The VuSystem includes modules to read and write a variety of media file formats. Video sequences can be captured and displayed using the X window system. Audio can be captured and displayed with the AudioFile application.

AudioFile (AF) is a device-independent network-transparent audio server developed at Digital Equipment Corporation's Cambridge Research Laboratory. With it, multiple audio applications can run simultaneously, sharing access to the actual audio hardware. Network transparency means that application programs can run on machines scattered throughout the network. Because AF permits applications to be device-independent, applications need not be rewritten to work with new audio hardware. AF does for sound what the X Window System does for text and graphics.

4.1.1 VuSystem Applications

Outlined below are a few of the VuSystem applications in use at MIT's Laboratory of Computer Science,

- *vvdemo* (Image Processing With Live Video Sources): The vvdemo allows one to experiment with various image processing filters, after choosing a video source users are given a range of filters to select from.

- *Office monitor*: The office monitor records a small video fragment of a visitor that drops by an office while the occupant is away. It works by taking continuous video from a stationary camera in the office, and only recording video whenever motion above some threshold is detected. Upon return to the office, the occupant uses a video browser to view videos of those who have dropped by while he was away.

- *Whiteboard recorder*: It keeps a history of changes to an office white board as it's owner adds to it and erases from it. It works by taking continuous video from a stationary camera pointed at the white board. The program filters out persons moving in front of the white board, and saves a minimum set of whiteboard history images by analyzing changes to the board and by following some simple rules.

4.1.2 VuSystem Application Shell

VuSystem applications are all written in a shell program that interprets a version of the Tool Command Language (Tcl). The language has been richly extended to support object-oriented programming. All event-driven code, including all user-interface code, is written as Tcl scripts. A Tcl interface to the Xt intrinsics and the Athena widget set has been implemented for the graphical user-interface code. Media processing modules are implemented as C++ classes and are linked into the shell. Simple applications that use the default set of media processing modules are written as Tcl scripts. More complicated applications require linking in additional modules to the default set.

4.1.3 PAVES - Extension to VuSystem

PAVES, is a direct manipulation system that combines aspects of visualization and multimedia systems to form an interactive video programming environment. PAVES extends the VuSystem media processing toolkit with dataflow-style views for controlling computation-intensive media applications as they run. Thus video may be manipulated graphically and interactively - both program and data are visual and live. PAVES is also novel in it's approach to extensibility. Users may freely combine graphical and underlying VuSystem textual programming methods to restructure and reuse applications. This cooperative programming is available across sessions as well as within them. It is implemented with an object-oriented programming foundation that automatically translates between multiple program representations and maintains them in synchrony.

4.2 The Vidboard - A Video Capture and Processing Peripheral for a Distributed Multimedia System

Asynchronous Transfer Mode (ATM) is a communication paradigm that offers support for seamless broadband communications across heterogeneous networking environments, from wide-area to the desk-area. ATM has many properties which make it well suited to the transport of real-time (audio and video) information. In turn many research groups are designing distributed multimedia systems centered around ATM networks. With the use of ATM, a new class of devices is possible in which devices act as shared network resources and communicate with each other through ATM-based protocols. In terms of multimedia, examples of such devices are video and audio capture boards, frame stores and video servers. Described below is one such device called the Vidboard.

The Vidboard was designed in the context of the ViewStation project (briefly described in chapter 3). The ViewStation system is an all-digital distributed multimedia system centered around a gigabit/second ATM network. The Vidboard was developed as a capture interface for a ViewStation environment. Full-motion video is captured from an analog television source and transmitted to other devices within the system. A typical ViewStation system interconnects general purpose workstations and specialized video resources, such as Vidboards, image processing systems, frame stores and video servers. Communication within the system is supported by a gigabit/second ATM network called the VuNet.

The VuNet consists of switches interconnected by fiber links. Clients connect to the network through a switch port. As described earlier, the software intensive nature of the VidStation environment places constraints on video sources. An application must have a mechanism for adapting the frame rate to system resources. The design of the Vidboard was motivated by the application friendly video source model. In this model, the application sends the Vidboard control information, which is processed by an intelligent video agent, and the Vidboard returns a video stream. Through this closed-loop technique, the workstation regulates the characteristics, and in particular the frame rate, of the video stream. This model led to the following design objectives for the Vidboard,

- *Temporal decoupling*: The ability to process video at a rate which is disjoint from that imposed by television video. The Vidboard serves as an interface between the real-time world of television and the virtual-time ViewStation environment.

- *Uncompressed video*: The Vidboard generates digital video in raw form. The reasons for using uncompressed video are twofold, i) network bandwidth is not expensive in the ViewStation system and ii) video data is handled at the application level. The current generation of workstations does not have the processing power to handle the high data rates of uncompressed video, limiting the frame rate of video streams. However, the ViewStation environment can be easily ported to higher performance workstations as they appear, alleviating this problem. This argument does not mean that compression does not play a role in the handling of digital video. It simply states that compression is not appropriate for shipping video from the Vidboard to an application running on a workstation. Areas where the group has found compression to be appropriate are storage and communication across low bandwidth networks attached to the ViewStation.

- *Distributed control*: The ability to receive and execute commands from other devices within the system.

- *Network transport options*: The ability to tailor the video stream to the needs of the destination device on a network transport level. Two factors are involved, i) the packing of video data into cells and ii) the generation of video streams having different traffic characteristics.

The Vidboard is based on a front-end frame-memory processor architecture. The architecture is centered around a Texas Instruments Digital Signal Processor (DSP). In turn, the Vidboard is actually a system having both hardware and software components. Video is captured by the front end and the resultant pixel information is stored in the frame memory. The organization of the pixel information within the frame memory is controlled by the format convert circuitry. The pixels are then read by the DSP, packed into ATM cells and sent to the network. The DSP is also responsible for receiving and interpreting commands from other devices within the ViewStation system.

4.3 Vex (A Programming System for Extracting Content Information from Unstructured Media)

While computers today are very capable of examining and transforming text, they are largely incapable of penetrating media. Media remains an opaque data type which may be captured, stored, retrieved and presented. However, it can not be examined, searched, transformed, and digested in such a way that the computer may perform independent actions based on the media's content. Vex, has been proposed as a possible tool to overcome this limitation.

Vex, is a general purpose search and analysis tool which enables it's user to specify the characteristics of the video that he or she is interested in and to perform actions based upon whether the video meets the specification. Vex poses the problem of understanding media in terms of recognizing events in media streams. Vex approaches this problem by adapting computer science tools for matching patterns and parsing text. Thus, Vex is a media equivalent to textual analysis tools such as grep or awk. It may be used as an interactive tool whereby users use their knowledge and intelligence to describe the imagery in which they are interested, while Vex uses it's ability to match patterns and process images to find video clips which meet the specification. For example, a user might be interested in searching an archive of news footage for highlights to Minnesota Twins games. Vex has no knowledge of baseball or even of the format of a news broadcast.

However, the user may, for instance, know that each video highlight is typically preceded by a shot of an anchor person and followed by a scoreboard graphic. Thus, if the Vex image processing library contains modules which could recognize these events (perhaps with the additional assistance of user supplied templates) then Vex can perform the search.

Vex represents an approach to content understanding of video which differs from much of today's research in computer vision. Using the analogy to text, one may argue that while considerable progress has been made in the area of natural language understanding, most users use much simpler search tools, such as grep, to find and manipulate text. For example, if one is searching through one's email to find an old message about a seminar, one is likely to search for words related to the topic of the seminar, the speaker's name, or even the date to find the message for which he or she is looking.

4.4 Content-based Indexing of Captioned Video on the ViewStation

Many broadcast television programs are close-captioned to provide a transcription of their audio content for the hearing impaired. Described below is a system that captures and processes the closed-captioned information to support automatic content-based indexing of these programs. The captioning system has been integrated as a component of the ViewStation, a desk-area network and programming system that is capable of transferring and displaying multimedia information. In addition to the spoken content of television programs, captions contain data regarding changes of speaker and background noises. A caption parser module that breaks captions into their distinct components has been developed. This module can be used to parse the captions associated with any television program. Caption cues are often specific to the television program and vary depending on the information that the application is designed to extract.

For each specific television program, a subclass of the caption parser can be developed that analyzes the cues in the captions, using general routines provided by the caption parser. When dealing with large amounts of video, it is desirable to enable access to smaller segments. The problem with segmenting video lies in determining what the segments should represent and where the divisions between segments should occur.

The goal is to find a method whereby a workstation could process a digitized video sequence and segment sizable video sequences into smaller meaningful bits that can be replayed and examined individually. One approach to the problem of segmentation is to perform it manually. This becomes infeasible as the amount of information grows and the needs of the community broaden. The idea of segmenting large chunks of information into smaller segments becomes much more interesting when it is done automatically. In order to perform this segmentation, regularities that are known to be present within the data can be exploited. These regularities can be simple things, such as a return to the shot of the news anchor before the introduction to the next story. Often the regularities are cues that can be used to indicate the division points between segments. Methods of automatically segmenting large amounts of data by finding these cues and defining the points at which the data streams can be broken into smaller blocks can be constructed.

The motivation behind the caption processing system is primarily to provide a means of determining the content of a television program. Since a number of television programs are captioned (including news and prime-time shows), a text translation of the audio component of the program can provide a significant amount of information.

4.5 MICE project (Multimedia conferencing on the World Wide Web)

The MICE project has been piloting multimedia conferencing (audio, video and shared workspace) over the Internet in Europe, with links to the US, over the last two years. Listed below are the findings of the project team,

Video tools

The group has evaluated the following tools for video conferencing, *nv* from Xerox, *IVS* from INRIA and *VIC* from Lawrence Berkeley Laboratories (LBL). As workstations have become more powerful, and the video cards on the workstations have improved, these tools have improved in terms of quality and frame rates that are attainable. In addition some work has been done to make video more tolerant to packet loss.

Audio tools

The most widely used audio tool for Mbone conferencing is *vat*. It has been constantly improved over the past two years, however audio quality is still a problem in many conferences. Sometimes, bad audio is due to peripherals (microphones and speaker boxes), whose quality varies considerably. The conferencing environment (echo and background noise) can cause more problems than newcomers to conferencing realize, particularly in conference rooms. The most persistent audio problem, however, is due to packet loss as a result of network congestion. Since this problem requires high-level concerted action, it is unlikely to be resolved in the near future.

Shared Workspaces

Tools for shared drawing have been developed by SICS (*Multidraw*) and INRIA (*Mscrawl*). These tools are not ready for use yet. The most commonly used joint working tool is still *wb* also from LBL, however it is not optimal for all shared workspace tasks. The University of Oslo has been developing an electronic whiteboard to be used in distance learning. This is a combination of hardware and software designed specifically for distance teaching in a classroom environment. A network text editor called *nt* that allows scaleable shared editing of the same document between Mbone sites has been developed by the MICE group. The systems use many different paradigms for the shared work-spaces, some of these scale as one moves to tens or hundreds of sites; some do not scale to more than a handful of sites. In contrast entirely distributed applications such as *wb* and *nt*, which use multicast and provide their own carefully designed protocols for dealing with inconsistency and retransmission can scale to hundreds of simultaneous users.

4.5.1 Multimedia Server

The MICE group has developed a multimedia server, the main objectives of this server are the following, (i) To record multicast conference data whose source may be any of the conferees, (ii) To playback recorded material either directly to one user or into another multicast conference, (iii) To allow users to create their own edits of recorded material, and create their own material for playback, (iv) To supply a large repository of archive space which is accessible to authorized network users who wish to record or playback multimedia data, (v) To allow synchronization between streams from each source.

In order to achieve the above, a system called the Video Conference Recorder (VCR) is under development. VCR has been designed as a client/server system, in which the server, (i) acts as a single point of contact for recording and playback, thus avoiding the use of multiple tools for different media, each of which may have a different interface, and (ii) has access to large amounts of disc space, which saves the user fro searching for enough space for each recording as is common at present. For this purpose UCL (University College London) has access to a large (180GB) magneto-optical jukebox as part of their distance learning initiative.

On the client side, there are currently three interfaces, one for recording, one for playback, and one for editing. The record interface will let the user start a recording of the audio and/or the video for one or many of the source streams, and to add a text description of the recording. The playback mechanism will allow the user to peruse all the on-line archives, and then to select one or more streams for playback. The editing client will provide sophisticated facilities to edit, annotate, and retrieve sub-items inside a recorded stream.

One of the main design concepts of VCR is the index. These indexes allow access to the source streams in a multitude of ways. The primary index used in VCR is created when a stream is originally recorded. For each source of a media, the incoming data will be saved together with an entry in this index. Each index entry contains a reference to the data, the time it arrived at the recorder, and a reference to some meta-data, which will initially be empty. At the end of a recording, each source will have a stream of data and a stream of index entries. It is the editing client which allows the user to manipulate these indexes in order to gain the flexibility required. For example, the user can add a text annotation to any part of the recording, and VCR will attach this to the source stream in the relevant place in the index by updating the meta-data field of an index entry.

Unlike some other systems designed for the purpose of multimedia recording, VCR provides a unified interface to the users so that recordings of different media can be shared together; moreover, because it is available over the network to authorized users and because a large store is guaranteed to be available for the purpose, remote clients can rely on it. Finally, the system is being integrated with the World Wide Web, so that web tools can be used for the retrieval and playback process.

4.6 Other Multimedia Tools on the World Wide Web

Outlined below are some of the other tools that were encountered in the research of multimedia technologies on the web.

4.6.1 CoReview (Mosaic + XTV)

CoReview is an interactive document and data retrieval tool. It has been developed to provide a seamless environment for groups and individuals, distributed across the Internet, that need to interact on the progression of a project. It can also assist individuals to put together a document in a collaborative manner. CoReview is based on the strengths of the World Wide Web server, Mosaic, and XTV, an X-window teleconferencing system. While Mosaic will be used to manage the project documents and review annotation files involved in proposals and their evaluation, XTV will aid in real-time remote collaboration among a group of users. CoReview incorporates the XTV features into a user friendly graphical interface and enables Mosaic to be shared by multiple networked participants. The system architecture embeds the concept of a chair that manages the shared resources. CoReview allows for easy creation of a pool of reviewers or proposal writers and automates the process of creating the necessary infrastructure daemons and directories at the needed sites.

The CoReview system consists of the CoReview utility and a set of daemons that are running at the sites of all reviewers. The CoReview daemons are used to setup a document for review by generating the appropriate directory structures and the links required by the review process. These daemons also facilitate the synchronous review of documents between the chair and a set of reviewers distributed across the Internet. The CoReview system provides the necessary tools to review the documents such as calls to Mosaic, an editor to generate annotations, the join mechanism, and a host of other features needed by the reviewers. The HTTP server at the chair's site is critical since most of the documents related to the review/edit process are located here. The individual annotations to the document as well as any other files created by a group member reside at the creator's site. Hence, a HTTP server is run at each of the member sites.

4.6.2 Hyper-G (Serving Information to the Web with Hyper-G)

The provision and maintenance of truly large-scale information resources on the World Wide Web necessitates server architectures offering substantially more functionality than simply serving HTML files from the local file system and processing CGI requests. Hyper-G is a large-scale, multi-protocol, distributed, hypermedia information system which uses an object-oriented database layer to provide information attribute and content search.

Hyper-G combines the intuitiveness of top-down hierarchical navigation with the immediacy of associative hyperlinks and the power of focused attribute and content searches. The basis for these three tightly-coupled navigational facilities is the Hyper-G data model. Documents may be grouped into aggregate collections, which may themselves belong to other collections and which may span multiple Hyper-G servers, providing a unified view of distributed resources. A special kind of collection called a cluster is used to form multimedia and/or multilingual aggregates (the appropriate language version of a document to be displayed is selected according to the user's language preference setting). Documents and collections may belong to multiple parent collections, opening up the possibility of providing multiple views of the same information (for example, incoming mail sorted by date, author, and subject). Collections are typically provided with an introductory text, the collection head, which is displayed automatically when a collection is accessed and often contains links to specific parts of the collection.

Hyperlinks in Hyper-G connect a source anchor within one document to either a destination anchor within another document, an entire document, or a collection. Links are not stored within documents but in a separate link database, hence links are not restricted to text documents, they can be followed backwards, they are updated and deleted automatically when their destination moves or is deleted (no "dangling links") and they are easy to visualize graphically.

Scaleable performance is achieved using a document naming scheme that allows replication and caching with weak consistency. Integrity of the database across server boundaries is maintained using a scaleable flood algorithm. Hyper-G has fully integrated search facilities: every document and collection is automatically indexed upon insertion into the database, no extra indexing steps are required. Both attribute and full text are supported, including Boolean combinations and term truncation. Searches may be focused, restricted in scope to particular sets of collections which may span multi-servers. Both anonymous and identified users are supported, with access rights assignable on a per document or per collection basis to user groups or individual users. Identified users have "home collections" within which to organize personal documents and keep pointers to resources.

As described earlier Hyper-G is a multi-protocol system. When accessed by a Gopher client, the Hyper-G server maps the collection hierarchy into a Gopher menu tree (hyperlinks cannot be represented in Gopher). A synthetic search item is generated at the foot of each Gopher menu to allow searching the corresponding collection. When accessed by a W3 client, each level of the collection hierarchy is converted to an HTML document containing a menu of links to it's members. Hyper-G text documents are transformed on-the-fly into HTML documents. Additional Hyper-G functionality such as user identification, language preferences selection, and searching is implemented via

HTML forms which are accessible at any time. The Hyper-G server is able to store pointers to remote objects on Gopher and W3 servers. This allows the incorporation of information on remote non-Hyper-G servers (almost) seamlessly. Interoperability with WAIS and FTP servers is planned. The Hyper-G server is currently available form most common UNIX platforms.

4.6.3 CAVE (Cave Automatic Virtual Environment)

Mosaic gives access to hypertext documents on the Internet, most of which contain primarily text and in-line images. However, some documents contain high-resolution images, audio clips, animations and video sequences. Mosaic presents these items by invoking an external viewer like xv, ghostview or mpeg (described in section 2.4.3). Mosaic serves as an interface to a viewer by locating, retrieving, storing the data and calling the viewer with the appropriate arguments. Viewers have specific control functionality that allow the user to modify or change viewing parameters of an image, animation or sound. For example, the user may be able to change color, size or crop an image, replay an animation at different speeds or view the animation backwards. Recently, Mosaic added forms support which allows the Mosaic user to input parameters and receive back information reflecting their choices. For example, a user may request an image of the United States weather system, designed to his own specifications. CAVEview expands this kind of user driven data and information gathering system to three dimensional computer graphics and virtual reality by presenting interactive CAVE applications within Mosaic HTML documents.

The CAVE is a virtual reality environment designed and implemented at the Electronic Visualization Laboratory at the University of Illinois at Chicago. The CAVE, is a surround screen, surround sound, projection based virtual reality environment system. The actual environment is a 10*10*10 foot cube, where images are rear-projected in stereo on 2 or 3 walls, and down-projected onto the floor. The viewer explores a virtual world by moving around inside the cube and grabbing objects with the 'wand', the CAVE input device. In the CAVE, high-resolution stereoscopic images are generated by a multi-processor Silicon Graphics workstation or multiple workstations and rear-projected onto several walls and front projected onto the floor. Several people can be in the CAVE simultaneously all wearing stereo glasses. One person is tracked and as they move in the CAVE, the correct perspective and stereo projection of the environment are updated to match their viewpoint. The rest of the projections are passive viewers, as though watching a 3D movie.

To aid in the development of the CAVE application, the *CAVE Simulator* has been developed. The *CAVE Simulator* emulates the CAVE by providing control of navigation, viewpoint and wand input. Recently, the CAVE has grown in popularity. Many research scientists, application programmers, and artists have begun utilizing the CAVE and the *CAVE Simulator.* As interest in the CAVE grew, there was a need for a presentation tool for CAVE applications. This tool needed to be geared toward the novice user, the individual who has little or no experience with the CAVE environment.

4.6.4 HUSH (Hyper Utility Shell)

HUSH (Hyper Utility Shell) is a C++ library that can be used to build applications with a graphical user interface. It contains classes that provide convenient yet flexible access to the functionality offered by the Tcl/Tk toolkit and it's extensions. Tcl is an interpreted script language, Tk is a window and graphics toolkit based on X11. HUSH provides a C++ class interface on top of Tcl/Tk. It provides flexible access to scripting languages. Not only Tcl, but also other script languages such as Python are supported. New widgets, written in C++, can be accessed from the script language. The HUSH library is intended to support the needs of both novices and experienced (window) programmers. Although HUSH can be used to write all sorts of programs with a graphical user interface, it mainly focuses on hypermedia applications (i.e., combination of hypertext and multimedia). Currently HUSH provides several hypermedia extensions, including a MPEG widget. MPEG is a widely used video standard. HUSH also includes a C++ API and a Tcl interface for *Csound*, a software sound synthesis package developed at MIT's Media Lab. Another extension is the web widget, which gives access to the World Wide Web.

4.6.5 Visualization through the World Wide Web using Geomview, Cyberview, W3Kit and WebOOGL

The following description is based on work done at the Geometry Center at the University of Minnesota.

Much of the center's effort is built around *Geomview*, the center's extensible mathematical visualization package. It is a public domain 3D viewer which runs on SGI, NeXTStep, and standard X workstations and allows full interactive control of the motion and appearance of multiple objects and cameras. Among it's features are full support for non-Euclidean and higher dimensional spaces. At the most basic level of web integration, *Geomview* can be run as a local external viewer for 3D scenes just as image viewers are used for 2D images.

Cyberview is a 3D viewer that runs through the web. *Cyberview* is essentially a web front-end for some of the capabilities of *Geomview*. For instance, clicking on an image in the *Cyberview* fill-out form sends a request to a specialized version for *Geomview* running on one of the Geometry Center computers, to compute a new image, a rotated view of the 3D object, which the group's web server then transmits back to the user as a part of a new HTML document created on-the-fly. That new document is itself a *Cyberview* fill-out form, so users can repeat this process. Clicking on an image in a web browser quasi-interactively positions the 3D object in the same way that using the mouse in an interactive 3D viewer on a user's local machine would change the user's point of view. *Cyberview* is one of a number of 2D and 3D applications in the Geometry Center's Interactive Gallery which are built on top of W3Kit toolkit (developed at the University of Minnesota). While their interactivity is limited by the response time of the web, these applications address the problems of portability and distribution in a novel way by taking advantage of widely available web browsers.

The *WebOOGL* is an experiment in 3D distributed hypermedia that implements bi-directional communication between *Geomview* and Mosaic. *WebOOGL* is an implementation of URL hyperlinks in 3D scenes in the spirit of VRML(Virtual Reality Modeling Language). Clicking on a *WebOOGL* hyperlinked object in *Geomview* can trigger a number of actions: for example, teleporting to a new *WebOOGL* world in *Geomview*, opening up a new HTML page in Mosaic, or downloading an image file. *WebOOGL* provides a general purpose tool for closely integrating 3D visualization with the web, and has been said to visualize webspace itself. An obvious future direction is to embed the webspace visualization in hyperbolic space, where there is an exponential "amount of room" (a given distance from a point), rather than the quadratic "amount of room" at the group's disposal in Euclidean space.

4.6.6 A MOO-Based Collaborative Hypermedia System

This is a hypermedia system intended to support groups of writers and scholars in writing and publishing hypertext articles by collaborating on a network. The current system supports the importation of individually-developed Storyspace hypertext documents into a MUD-based (Multi User Domain) collaborative workspace for integration and expansion, and allows for the immediate publication of these dynamically generated multimedia documents onto the World Wide Web. In addition, a forms-based writing and linking interface to the text is available, so that writers can write using either the MUD-based or the forms-based authoring tools.

MUDs, or Multi User Domains, evolved out of multi-player Adventure-style games in the early 80s. These began as hack-and-slash style games, but some of the MUDs began to evolve into more social areas, somewhat reminiscent of chat lines Although many of the earlier systems relied on hard-coded behaviors, MUD systems began to incorporate internal scripting languages. One particularly flexible server, MOO (MUD Object-Oriented), is now being widely used by the research community to support collaborative work, due to the ease of modifying the environment to support scholarship and sharing of information.

4.6.7 Multicast Trees for Data Distribution over Global Networks

Recently, high bandwidth applications, such as video conferencing, have penetrated the internet with requirements to multicast the same information to multiple destinations. The implications of introducing such technology are just now being realized and yet, at the same time developers are trying to exploit this new functionality and extend it to the global internet. Example of applications that use multicast are:

- Distributed database updates
- Video Conferencing over packet networks
- Wide area distributed computations
- Distance learning
- Distributed interactive simulation

The standard method for an internetwork to provide a multicast service is to construct a multicast delivery tree upon which information is routed amongst a known group of members. The purpose behind using a delivery tree is to minimize the number of copies of a single packet that the network must transmit, especially with the high volume of data originating from video sources. It is imperative that whatever method is used to accomplish this data distribution, it must be done efficiently from both the network and the end user's point of view.

There are two general approaches to multicast tree construction. The first uses source specific trees (SST) which involves finding the shortest path through the network from each sender to each receiver. The use of multiple source specific trees provides the minimum delay from each sender to every receiver but incurs high overhead. Center specific trees (CST), or shared trees, use a distribution center in their construction of a single multicast tree. The use of a single minimum weight tree to connect a set of nodes in a graph is the well known Steiner Tree Problem. A center specific tree is a low overhead method that sacrifices minimal end-to-end delay.

The current draft-standards for network level multicast being considered by the Internet Engineering Task Force (IETF) are Protocol Independent Multicast (PIM) and Core Based Trees (CBT). In the case of these two proposed approaches, both PIM and CBT use centers to accomplish CST construction. Neither PIM nor CBT specify an algorithmic method to select the location for the center which, can affect the quality of the resulting tree. Presently, the draft standards choose to have these locations selected administratively. The desirable features for multicast data distribution service in a global network are:

- Scalability (with respect to the number and distribution of members)
- Low end-to-end delay (includes transmission and routing delays)
- The ability to chose between source specific and center specific trees (or a combination thereof)
- Group specific centers for center specific trees
- Use of existing unicast routing tables

Chapter 5 delves into techniques that are being developed to make the web more secure.

Chapter 5

Security on the World Wide Web

Communication on the Internet is constantly under the threat of attack from such sources as eavesdropping, unauthorized access and message modifications. Several tools, methodologies and protocols have been developed over the years to allay the fears of such attacks, described herein are some of these.

5.1 Secure TCP (providing security functions in TCP layer)

On the Internet, some application protocols handle important or sensitive data. However, because of lack of security functions in the current TCP/IP protocols, they are under threat of attack. Some of the mechanisms to guard the Internet are Privacy Enhanced Mail (PEM), Privacy Enhanced Telnet (PET), Kerberos authentication service, *swIPe* and *IP/Secure*. In order to reduce overhead for adding security functions to current existing application protocols and systems, various kinds of security functions for host/ host communication are performed by IP level security such as swIPe or IP/Secure. However, for process/process communication, enhancement of each application protocol has been the only way to add security functions. Therefore, because of the overhead involved in enhancing application protocols, only a few applications provide security functions such as data encryption or user authentication.

In order to reduce this overhead and accelerate the integration of security functions to currently available network applications, security functions should be provided in the transport layer. To this end a new mechanism called "*Secure TCP*", a security enhanced TCP, is under development, to guard communication against attacks. Since TCP is widely used for various applications, the Secure TCP can be applied to many currently existing network applications. Secure TCP is a common platform for network applications to achieve secure process/process communication. Furthermore it provides a solution for the smooth migration to secure services for existing applications. Secure TCP is interoperable with classical TCP since it uses a natural extension of TCP's three-way handshake to exchange session keys for TCP segment encryption. This feature is necessary for it to be applied to the existing network environment

5.1.1 Security problems with TCP

TCP provides fundamental communication service and is used by many application protocols. Therefore, TCP has become one of the most popular protocols. But TCP has several security problems as illustrated below,

- TCP cannot guard a segment from message modification attacks. TCP's checksum is used to identify the modification of a segment. However, since this field is not protected against message modification, it is possible to modify any TCP segment. Moreover, there is no way for peer entities to diagnose message modification.

- TCP cannot keep segment data secure from message eavesdropping attacks, TCP transports stream data used in the application layer. Since TCP does not provide any data encryption functions, anyone can gain access to valuable information.

- TCP cannot protect connections from unauthorized access, TCP certifies a peer entity by a source IP address and a port number. However, it is possible to modify the source address and port number.

5.1.2 Goals of the Secure TCP project

The goals of this project are, (i) Provide data integrity for the TCP segment. The data integrity is a security service where right entities can identify data modified by unauthorized entities. (ii) Provide data confidentiality for TCP segment data. The data confidentiality is a security service that makes data unreadable to unauthorized agents. (iii) Provide a negotiation service for using various security functions in process/process communication. (iv) Keep interoperability with current version of TCP.

Secure TCP by itself does not provide any authentication services. Because several third party services such as Kerberos are now available, Secure TCP relies on these for authentication. In this sense, each peer entity in TCP layer has it's certificate and the Secure TCP uses it for authentication.

5.1.3 Secure TCP Protocols

The essential parts of Secure TCP are the security service negotiation, the key exchange for encryption, and TCP segment encryption. In the following subsections, the design of the security service negotiation and key exchange for encryption are described,

Security Service Type Negotiation

Secure TCP can choose from several kinds of cipher (the unintelligible text that results from encrypting original text) methods used for data integrity and confidentiality, and select various hash algorithms for data integrity. In the negotiation phase Secure TCP peer entities exchange a security service type that indicates both cipher methods and a hash algorithm.

A negotiation procedure is carried out as follows. A sender entity sends a list of types to a peer entity. The list order indicates the priority the sender wants to use. For example in the case the sender sends "421", it wants to use MD5 and DES (Data Encryption Standard) for it's data integrity and DES for it's data confidentiality. The receiver entity chooses one of the methods that the sender had advertised in the list, and sends back a message. If the receiver can't pick any methods in the list, then it can either give up the connection established or proceed through classical TCP. In the case that SYN+ACK segment is sent back to the sender, the sender can select to reset the connection and retry with another listed security method or proceed with classical TCP.

Key Exchange

Another objective of the negotiation phase is key exchange between peer entities. This exchange should be done in a secure manner. In key exchange, a "public key certificate", is introduced in order to protect the key exchange against eavesdropping. After the key exchange, peer entities share session keys for datagram integrity and confidentiality. Before a procedure of the key exchange is carried out, peer entities of communication have a public key certificate that is signed by a reliable certificate authority.

5.2 Network Security Projects

There are a number of security projects underway in the network security area. This section briefly describes some of these projects.

5.2.1 DISNet and BLACKER

A major network security project now underway within the DoD (Department of Defense) is attempting to consolidate the department's various Defense Secure Networks into a single Defense Integrated Secure Network. Known as DISNet, the different component networks are,

- DSNet 1, which processes secret information.
- DSNet 2, which processes top secret information.
- DSNet 3, which processes sensitive compartmented intelligence (SCI).

The network is expected to be available in the next few years. By using the BLACKER system from UNISYS Corporation, the consolidated network will use different encryption algorithms to keep separate the different security levels of information stored and transmitted over the networks. The system will employ key distribution centers for the centralized distribution of electronic cryptographic keys.

5.2.2 SDNS

A major effort under development by NSA (National Security Agency) is the Secure Data Network System (SDNS) Program. This program was established in 1986 as a partnership between NSA and ten major computer and telecommunications companies, including AT&T, Digital Equipment Corporation, and IBM. The goal of SDNS is to develop a security architecture based on the OSI (Open Systems Inteconnectivity) model, as well as standard security protocols for the Government Open System Interconnect Profile (GOSIP). Through NIST, the program has published documents describing proposed standards for controlling access to a network, protecting data in networks, and managing encryption keys. SDNS is expected to use a distributed key generation system in which computers will be able to establish cryptographic keys as needed to communicate with other computers.

5.2.3 Kerberos

Kerberos is an authentication system for open systems and networks. Developed by Project Athena at MIT, Kerberos can be added onto any existing network protocol. Historically used with UNIX-oriented protocols like Sun's Network File System (NFS), Kerberos is expected to be used by the Open Software Foundation's OSF/1 operating system, and by other vendors. Kerberos uses an encryption standard based on the Data Encryption Standard. Each user has a private authentication key. Kerberos works by guarding the data transmitted between machines that communicate over the network. Kerberos uses cryptographic keys known as tickets to protect the security of the messages the user sends to the system-and the messages the system sends back to the user. Kerberos never transmits passwords, even in encrypted form, on the network. Passwords reside only in a highly secure machine called a key server. Kerberos performs authentication both when users log into the system and when they request any type of network service.

5.2.4 Project MAX

Project MAX is a network security research effort established by a number of commercial vendors, including SecureWare, Sun, and Hewlett-Packard. The goal of the project is to develop a trusted network technology that's independent of vendor platform, specific operating system, and network protocol. This will allow information to be exchanged securely between disparate systems and networks. The group has developed a trusted network product called MaxSix (Multilevel Architecture for X for Security Information Exchange). MaxSix consists of a set of enhancements to UNIX and networking interprocess communication.

An important goal of the project is to standardize the labeling of the data that's imported into or exported form a system. In MaxSix trusted networks, the networking software transmits security attributes along with data, and inserts appropriate checks in the network subsystem that enable a local system to make access decisions (based on it's own security policies) about the data imported into the system from the outside world. MaxSix was designed to support network security standards, including DNSIX and TSIG. DNSIX is the Department of Defense Intelligence Information System (DoDIIS) Network for Security Information Exchange, a specification developed by the Defense intelligence Agency (DIA) and Mitre for network level access control and session management. The DNSIX standard is specified as a requirement for many DoD and DIA System High Workstation and Compartmental Mode Workstation procurements. TSIG is the Trusted Systems Interoperability Group, which has developed a specification for message headers that includes attributes, such as sensitivity labels, used for trusted routing of messages in a network. TSIG has proposed a Commercial Internet Protocol Security Option (CIPSO) header that is currently being accessed.

5.2.5 Secure NFS

Sun Microsystems's Network File System (NFS) is a set of file transfer protocols that has become virtually an industry standard for UNIX systems. NFS allows a client system to mount devices on a file server just as if the server were physically connected to the system. File transfers are supported over heterogeneous networks. Sun has recently introduced Secure NFS, which uses two types of encryption-the Data Encryption Standard (a private key system) and a public key encryption system.

5.3 Security for Electronic Commerce

Today there are two basic approaches to secure electronic commerce. The first one focuses on protecting resources by securing individual servers and network sites. This access security is generally addressed by firewalls (gateway computers programmed to block unwanted traffic arriving from the Net) or other means of "perimeter security". The second approach focuses on transaction security. Transaction security addresses unauthorized listening in or eavesdropping on buyer/seller communications; authentication, so both parties are confident they know who they're talking to; message integrity, so the message contents can't be changed or tampered with; and a nonrepudiable record of the transaction in the form of a receipt or signature. One way to achieve these security properties is channel-based security, which secure the channel along which the transaction is taking place. Users assume that any data passing through this channel is secure. Document based security focuses on securing the documents that make up the transaction.

Two emerging standards address channel-based and document based security. The SSL (Secure Sockets Layer) system from Netscape Communications is the leading channel-based technology. The key document based approach is the SHTTP (Secure Hypertext Transport Protocol) system from Enterprise Integration Technologies, which is the lead sponsor of CommerceNet, a nonprofit consortium. It's members include Apple, Bank of America, and Hewlett-Packard. CommerceNet has created many pilot electronic-commerce projects and is a key proponent of SHTTP.

5.3.1 CyberCash Payment System

One of the critical technologies that must be in place to enable electronic commerce is the ability to move funds between parties. Until now, the only means that have been available have required separate arrangements for payments or very limited use of credit cards. Some have attempted to introduce an integrated electronic mall with facilities for making credit payments within the context of the mall. Others have attempted to define a more general scheme but have not yet been connected to the banking system. The CyberCash system is a separate system which can be used by any user, any merchant, and any bank. Moreover, it is introducing a new form of payment, peer-to-peer, which permits individuals to pay each other without requiring the payee to be a merchant authorized to accept credit cards.

5.3.2 Customer to Merchant Payments

CyberCash (CC) will provide two classes of payment services, customer to merchant and peer to peer. Customer to merchant payments include credit card and debit card. For customer to merchant payments, there are three parties connected to the Internet, the customer, the merchant, and the CyberCash server. A customer will engage in a dialog with the merchant's web server, eventually arriving at a decision to buy some goods or services. All of the dialog related to the sale is under control of the merchant until it's time to make the payment. At that point, the merchant presents the user with the amount to be paid and a transaction identifier, and then turns control over to the CyberCash payment system.

The CyberCash payment system consists of "CC user software" on the user's machine, "CC merchant software" on the merchant's machine, and the CyberCash server (CC server). The user invokes the CC user software, provides his credit card information, and authorizes the sale for a specific amount. Using public key technology, the CC user software encrypts this information under the CC server's key and sends it to the merchant. The merchant thus has a message that it cannot read completely and cannot tamper with but which authorizes the purchase. The merchant adds his identification information and sends it to the CC server. The entire message is digitally signed by the merchant to prevent tampering intransit.

The CC server unwraps the message and creates a standard credit card authorization request. The CC server then forwards the request to the appropriate bank or processing house for authorization and returns the result to the merchant. This is the basic protocol, but there are a number of variations that can be accommodated. As noted, the credit card information is encrypted so the merchant cannot see it. This prevents handling errors such as unauthorized charges originating with the merchant, but some merchants may require the credit card information as part of their record keeping. If the merchant's bank authorizes the merchant to receive the credit card numbers, the CC server will forward the information back to the merchant with the authorization. Not all transactions will take place via the web. Email and other forms of interaction are possible, and the same protocols are usable in those settings as well.

5.3.3 Peer to Peer Payments

CyberCash enables peer to peer payments between CyberCash account holders and from CyberCash account holders to others. CyberCash accounts are established by users on demand. A user who wishes to open a CyberCash account first obtains the

CyberCash software for his machine and then contacts the CyberCash server. As part of this process, the user will generate a public/private key pair. The public key will be known to the CC server as will be the identifier for the user's account. The private key will be known only to the user and will permit the user to sign the instructions he gives to the CyberCash server. Once the user has opened a CyberCash account, he can load his account by sending an instruction to move money out of his checking account. This instruction has to be digitally signed and accompanied by information which authenticates that the user legitimately has access to that checking account. The funds are then moved from the user's checking account to a CyberCash agency account in the same bank or another bank that is cooperating with the CyberCash service. At the same time a mention is made on the CyberCash books that the user has those funds in this CyberCash account.

Whenever the user desires to make a payment to another CyberCash account, the user sends the CyberCash server an instruction to do so. The CyberCash server checks that the funds exist in the account and that the instruction is not a duplicate. It then updates both accounts and issues a signed receipt back to the user. The user may then present the receipt to payee to demonstrate that the payment has been made. The receipt is signed by the CC server and can be checked for validity without interacting with the CC server, The payee may also contact the CC server and check the status of his account. When the user desires to move the money out of his CyberCash account, he sends an instruction to the CC server and the funds are moved out of the CyberCash account into the user's checking account. These funds are then available for any kind of transaction, including disbursement in cash.

5.3.4 CyberCash Security

Security is obviously paramount in this operation. As noted above, public key technology is used, in combination with symmetric algorithms and hash algorithms, to protect the messages in transit and authenticate origins. The public key algorithm is RSA, the symmetric algorithm is DES, and the hash is MD5.

Two other aspects of security are also important. The user must be certain he is using a true copy of the CyberCash code. To provide assurance to the user that he is not using modified code, a third party verification service will be available which provides certificates which convey the correct hash code for the software. Routines which check the hash code will be available from a variety of sources and will be well documented. Protection of the user's private key is also paramount. User software will keep the user's private key encrypted and available only with a local password. Strong local protection of the password does not protect against loss of a private key. A separate facility is needed in case the user loses access to his private key. To solve this problem users will be encouraged to make use of third party data recovery centers that can store encrypted information and make it available only upon successfully completing an interrogatory which had been provided by the user when the data was originally stored.

A discussion of data integrity followed by the specification for a tool that will enable web users to measure data integrity of World Wide Web sites makes up the next chapter.

Chapter 6

Evaluating Data Integrity on the World Wide Web

Databases/web-sites rife with incorrect, untimely and incomplete data can put a damper on effective decision making, to illustrate, a popular financial data provider[1] had "-6" listed as the number of employees for General Motors corporation, several other of the source's columns did not contain any data. Upon further investigation, it was detected that the problem arose during data-feed, early detection of this problem could have resulted in an improvement of the data-feed technique, however by the time the data integrity problems were detected and corrected several man hours had been lost and most of the data was no longer relevant. Problems such as the above call for a formal means to measure data integrity, which is the goal of the Data Integrity measurement process, described here-in.

6.1 Overview of Data Integrity Measurement process

Data integrity measurement is the process of analyzing, comparing and measuring the integrity of disparate databases.

In what follows, a system specification that will aid in the measurement of data integrity is described (refer Figure 6.1). Several modules are involved in the process, they include data integrity source trackers, error checkers, data integrity enhancers and the data integrity software tool. Source-trackers are used to periodically gauge the integrity of data in databases, in performing there function source-trackers compare data integrity in candidate databases to data in an ideal database (i.e., database with an optimal data integrity rating). Error checkers and data integrity enhancers aid in the mediation process

1. Fortune500 database service.

by detecting erroneous data, reducing error where applicable and enhancing data integrity. This discussion will focus on the data integrity tool which works synergistically with the afore-mentioned modules to measure and compare data integrity across databases.

Figure 6.1: Measuring Data Integrity

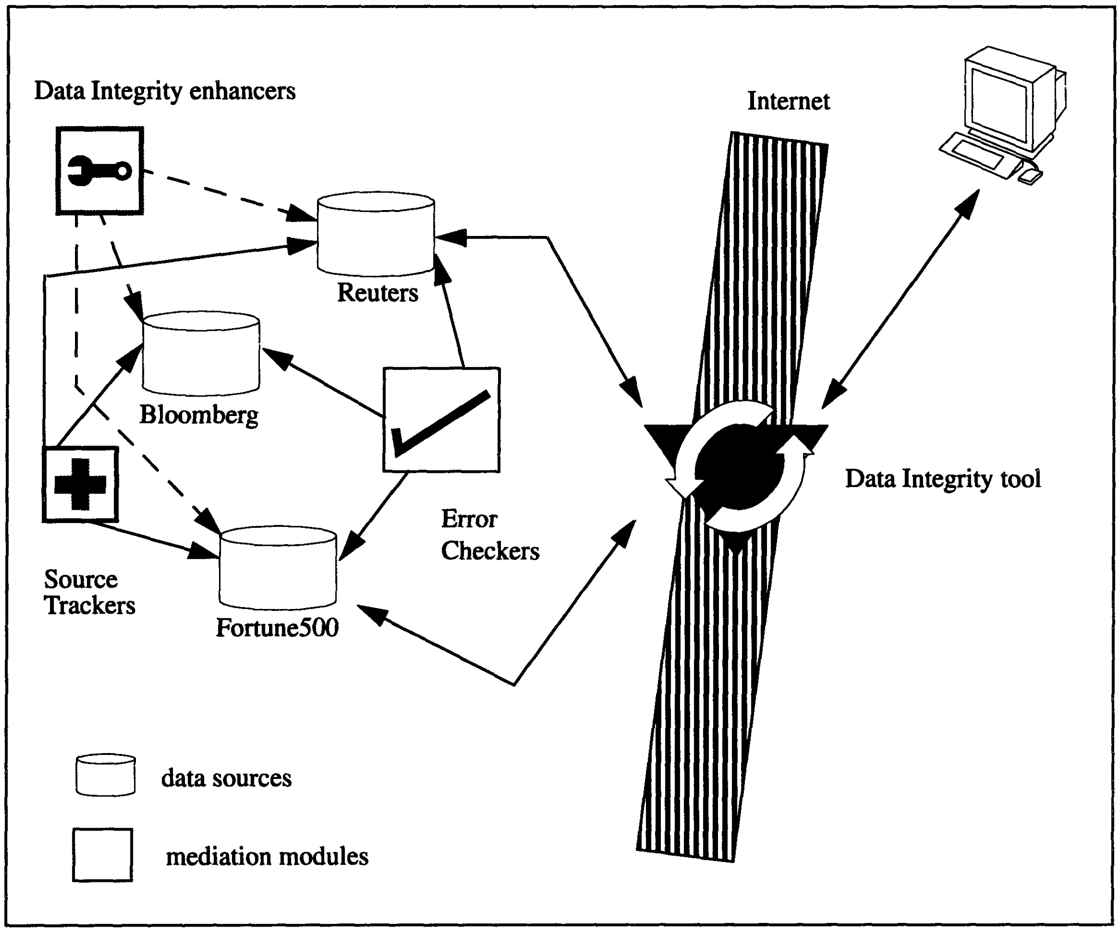

6.2 Data Integrity Tool

The tool is used to make periodic (daily) measurements of integrity of the various databases on the internet, the results of these measurements are stored in system catalogues. When an application[1] or user starts up the tool and specifies a domain, the tool using an in built expert system determines the integrity requirements of the domain, and proceeds to match these requirements with the integrity measures of databases residing in

1. Applications like excel, word.

it's system catalogues. The tool presents the results of these searches to the user or calling application, and extracts data from the selected source.

6.2.1 Data Integrity Rating

In measuring the aggregate data integrity of a web-site/database, several dimensions are taken into consideration, particularly, correctness, timeliness, completeness, credibility and format integrity [Wang, Strong, Guarascio, 1993]. The criteria user for selecting dimensions are application-independence and quantifiability.

- *correctness*, is a measure of the number of non-garbled[1] data values that match data values residing in an ideal web-site/database.
- *timeliness*, is a measure of how current the data is.
- *completeness*, is a measure of the number of non null mandatory values present.
- *credibility*, is a measure of how credible a data source is. The origin and method of data collection, will be used in determining credibility.
- *format integrity*, is a measure of the number of data values that conform to the column data type and data base design restrictions (i.e., referential integrity, normalized schema).

6.2.2 Measuring Integrity Rating

In determining the integrity rating of a web-site/database, each individual dimension is first measured, then the function outlined below is applied to these results

$$rating = \sum_{i=1}^{n} w_i \bullet r_i$$

where,

w, weight associated to each dimension

r, measure of a data integrity dimension of the column

n, is the number of dimensions considered in determining the rating of a column

1. erroneous data, usually attributed to bad data entry.

For the sake of measuring the rating of a database, weights are assigned to dimensions based on their importance. One possible mix of weights:

scale of 0 - 10

w(completeness) = 10, *w(format integrity)* = 10, *w(timeliness)* = 8,

w(credibility) = 9, *w(correctness)* = 10

In this mix timeliness and credibility are assigned lesser weights.

6.2.3 User Interface

The user interface shown in Figures 6.2 and 6.3, indicates how the tool's interface would look if it were geared towards users in the financial domain [Wang, Page, & Kaomea, 1994]. The interface consists of pull down menus and sliding bars. Users specify their domains using the menus (e.g., Figure 6.2 indicates *company analyst*, analyzing *technology stocks*), data requirements are specified using sliding bars, results are presented in a graphical format shown in Figure 6.3. Users can also specify how much they are willing to spend (in dollars) and the data retrieval delay they can tolerate.

Figure 6.2: Menu interface

USER	INDUSTRY	AREA	ITEM
Company Analyst Industry Analyst Portfolio Manager	Automobile Banking **Technology** Utilities	**Company** Industry Macro-economy	**Stock** Bond Option

Figure 6.3: Data integrity requirements

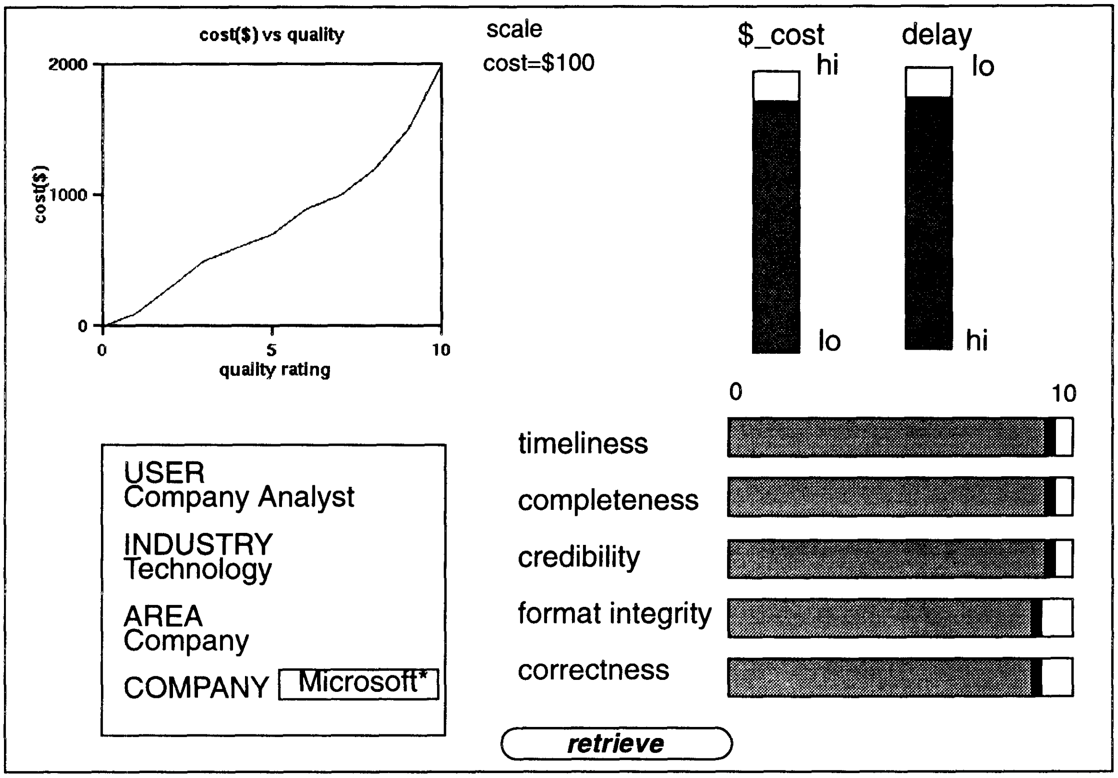

6.2.4 Data flow

Once a domain is specified as above, the tool determines what data integrity requirements are needed, i.e., it determines what trade-offs are tolerable between the different dimensions. It arrives at an aggregate value for required data integrity by assigning appropriate weights to the dimensions; these weights are applied based on their importance to the specified domain. For e.g., to a foreign currency trader timeliness would be very important, while a stock analyst, might be more interested in historic data over a broader area (i.e., completeness is more important than timeliness).

The aggregate value, is then matched against integrity measures of databases, residing in system catalogues. Results of this match produces a list of databases that match the data requirements of the domain. The results of this matching are then presented to the user in a graph, the graph is an increasing function of cost (dollar cost of accessing source) vs. rating. The user can select a specific source either by clicking on the appropriate point in the graph or by clicking on a drop down list. [Wang, Page, Kaomea, 1994]

Figure 6.4: Data flow

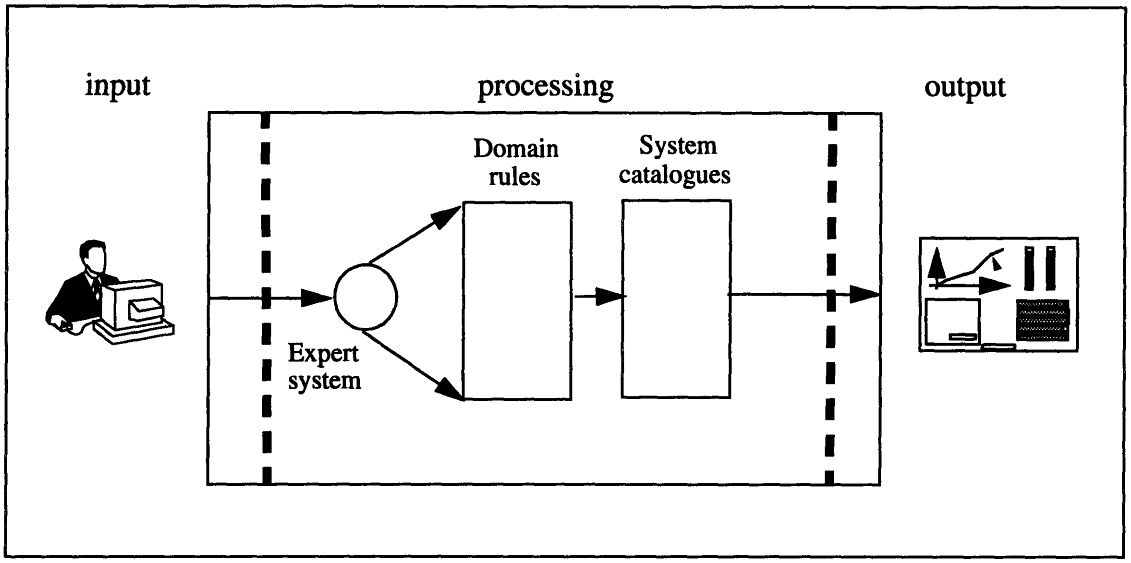

6.2.5 Processing

The tool uses the specified domain (company analyst, analyzing technology stock), to determine integrity requirements. The requirements (high timeliness, high credibility etc.) are determined using an expert system and in built rules. The user may also explicitly specify requirements using the interface, thereby over riding the expert system.

6.2.6 Expert system

An expert system uses the domain knowledge specified by the user (specified in menu interface) to determine appropriate data integrity requirements. The expert system with the aid of it's forward chaining mechanism finds the set of rules that best fit the user's domain. The rules specify the data requirements of the user.

6.2.7 Domain rules

Prior knowledge in the form of rules, will enable retrieval of data from data sources that fit domain requirements. Rules for retrieval will be user and domain specific.

A sample rule may look as follows,

IF domain = *(("speculator")* AND *("foreign currency")*) THEN

dimension_measure =

((**timeliness** = 10) AND (**correctness** = 10) AND

(**completeness** = 10) AND (9 < **credibility** < 10) AND

(9 < **format integrity** < 10))

6.2.8 Calculating data rating

Once the requirements have been determined, an aggregate value of data integrity is determined using the function outlined in section 6.2.2. The integrity measure of databases residing in the system catalogues are also determined by applying the same function, the rating from the tool is then matched against the integrity measure of databases in the system catalogues. A cost vs. rating graph presents the results of the calculation to the user, the graph indicates the ratings of databases that are greater than or equal to the required rating.

6.3 Data Integrity tool: scenario

In the financial world, there is a great need for data of high integrity. In the following scenario, we explain how a stock analyst can use the data integrity tool to assess the value of a company stock (using stock data from financial data providers on the Internet).

A stock analyst working for a brokerage firm on Wall St., is analyzing Microsoft Corp's stock for one of the firm's brokers. Based on his suggestions the broker will either buy or sell Microsoft Corp. stock. The analyst uses technical analysis principles. The analyst has noticed a recent upswing in the purchasing of Microsoft stock, he would like to take a look at price trend data to determine if investors are still Bullish on the stock. He has had some difficulties with errant price data in the past, and would like to ensure that he extracts data from a source with data of the highest integrity, he decides that this is a good opportunity to test the tool and proceeds to do his analysis with it.

6.3.1 Interaction with tool

He installs the tool on his PC and fires it up. He is lead through a series of menus, requesting him to specify his domain. Once the domain is specified the cost-rating screen appears (Figure 6.6). The dimension measures are set by default for the domain specified, as indicated by the sliding bars in Figure 6.3. He then enters the name of the company he is interested in evaluating and clicks on the retrieve button. He notices a graph, he assumes it is representative of how much he will be charged for the data retrieval for a given data integrity.

Figure 6.5: Menu interface

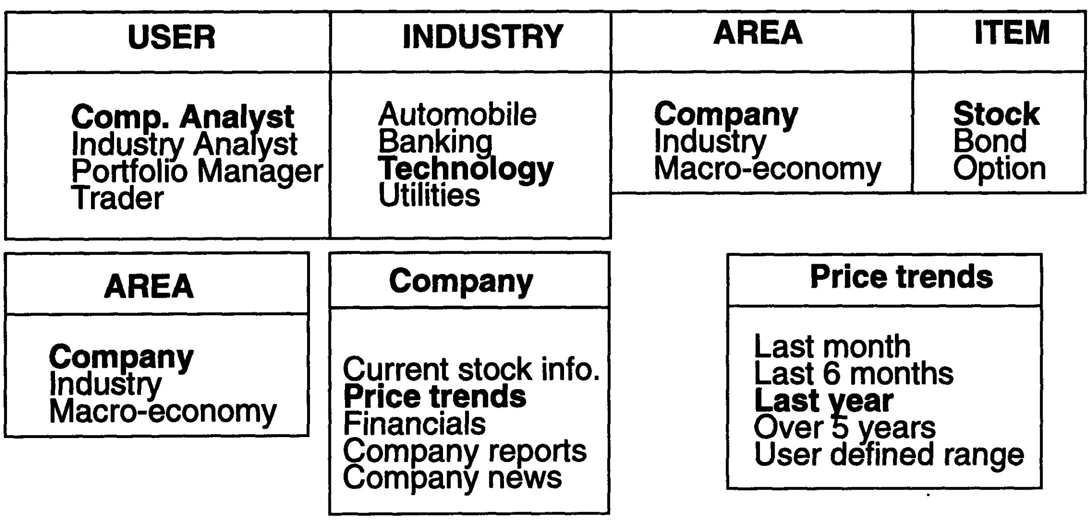

Figure 6.6: user requirements - cost vs. rating graph

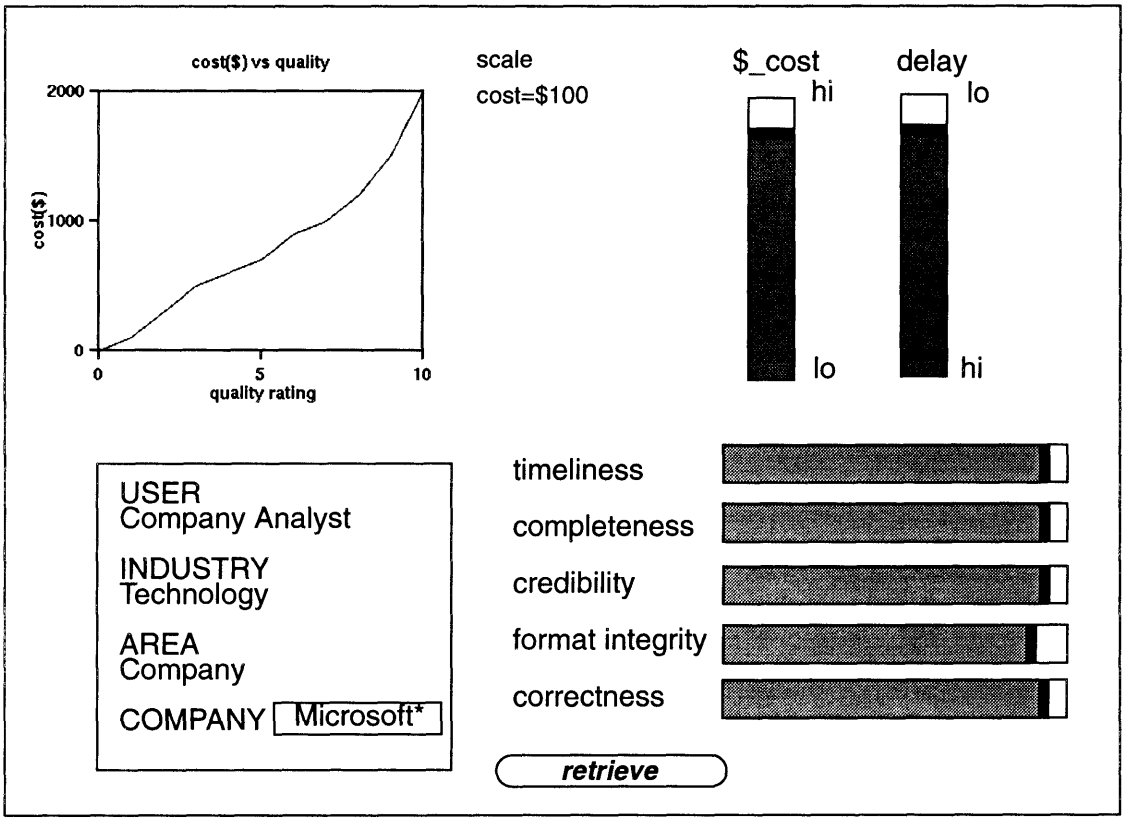

6.3.2 Behind the scenes

Once the application domain is specified as above, the tool makes use of domain specific rules to determine the requirements of the user, in this case the measures in the sliding bars are the requirements for a company analyst. The tool then proceeds to match the requirements against all pertinent data sources. As Figure 6.6, indicates, the tool will look for data from a database that has an optimal or near optimal measure for timeliness, completeness, credibility and correctness.

On this occasion, the tool proceeds to compare data integrity and extract price trend data from, World Wide Web data feed services, like Reuters, Bloomberg, Datastream, Worldscope. The tool having measured and compared the data integrity of the above sources, finds that the 'MIT AI lab' price feed database, had data of highest integrity. It returns this information to the user and also presents the user with the reasons for why it chose 'MIT AI lab' and not the other sources.

The tool extracts the data and returns it in the following chart format:

Figure 6.7: Company stock chart

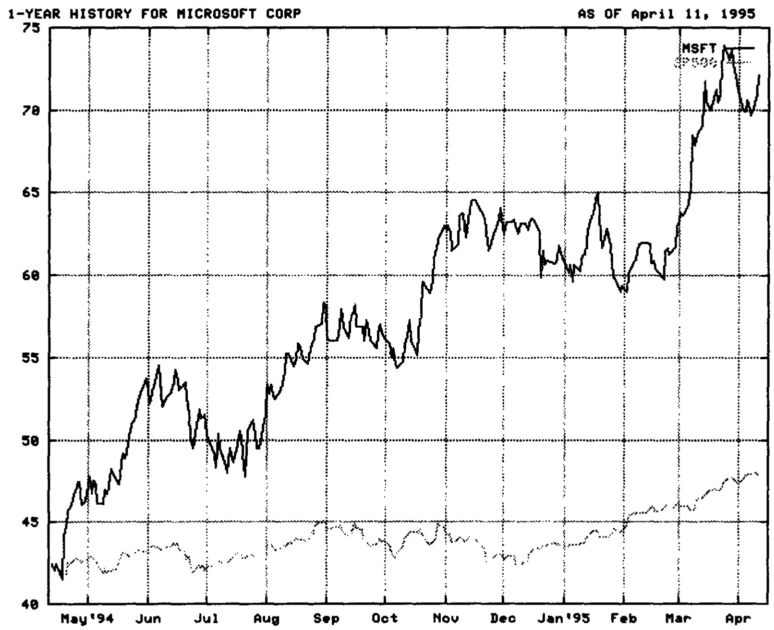

6.3.3 Benefits

i) The tool in choosing a source, compares the data integrity of several sources, ensuring that the user receives data of the highest integrity (given cost user can afford, and application requirements) of all available sources. **ii)** The tool in performing it's analysis and comparison of data integrity will be able to identify and flag data sources providing errant or irregular data, there-by eliminating these sources from the list of sources that data can be retrieved from. **iii)** The tool saves considerable time for the user, (i.e., time spent determining wether data is accurate) also saves the user the dollar cost of having to access several sources to confirm data accuracy. **iv)** The tool also provides reasons and explanations as to why a certain source was selected and others not selected.

Chapter 7

Conclusions and Frontiers of new research for World Wide Web computing

7.1 Conclusions

Outlined below are the salient features of this thesis, with an emphasis on Web technologies that will dominate on-line computing in the near-a-future.

7.1.1 Networking technology

As the influence of technology and communication on our day-to-day lives increases, so does the size of the World Wide Web, from modest beginnings in the '60s when the total number of hosts on-line numbered 15-20 to a staggering 4, 000, 000 and counting today. With this rapid growth has risen the need for networks that can handle huge data loads, in the thesis the technologies that are being built to cater to these needs like the Mbone (Multicasting Backbone, able to handle up to 1.5 Mbps), IPng or Ipv6 (next generation internetworking protocol) with larger addressing room and improved security features, are described.

7.1.2 Multimedia technology

Multimedia technology on the web, has also seen tremendous advances, from text based menuing interfaces in the '80s to sophisticated, video conferencing, on-line publishing and 3-D graphics, today, users of the web are constantly seeking more sophisticated multimedia tools. The latest tools and trends in this area are outlined in the thesis. The VuSystem, a distributed Multimedia system being developed at MIT's Laboratory of Computer Science that promises to revolutionize multimedia technology by taking a software-centric approach is described in detail. Improvements in networking technology to handle sophisticated multimedia programs are presented.

7.1.3 Security technology

The need for Internet security has grown exponentially over the years as well. In recent years the Internet has been under attack from hackers, who, eavesdrop on transactions, modify messages, and gain unauthorized access to sensitive systems, some of the measures being taken to secure networks from such threats are discussed. Techniques and tools that will make on-line commerce secure and safe from tampering are also presented. New protocols like the Secure TCP protocol that are being developed to enhance network security are explained.

7.1.4 Data integrity technology

On-line data integrity is one area of research that has not seen much change or improvement, as the number of World Wide Web sources burgeon the need to be able to evaluate the integrity of these sources before accessing them has increased. Evaluation of data source integrity can save web users a lot of time and effort, by restricting their data source access to only those that satisfy their data integrity requirements users can optimize time spent on-line. The specification for a tool that will enable the measuring of data integrity for on-line data sources is presented.

7.2 Where the Web is headed

Today there are 4, 000, 000 million nodes or servers and several million users on the Internet, with the number of nodes and users growing exponentially. At this rate, very soon current networking technologies will not be able to handle the load, hence work has begun on building faster, broader networks. Many commercial services already provide networks that are faster than the Internet. NSF is funding the development of vBNS (Very-

High-Speed Backbone Network Service), that focus on broadband, internetworked technologies and services. The core of the project will be to improve the speed and scaling of the internet and it's underlying technologies, developers promise 155Mbps service, by contrast the internet backbone presently operates at 45 Mbps. Efforts are also underway to improve the current Internet Protocol(IPv4), that can only handle a limited number of addresses and is highly susceptible to security breaches, the proposed protocol is IP6: it will offer expanded addressing, simplified packet routing, message handling, and improved security features to be implemented in 1996.

There has also been a tremendous growth and increasing sophistication in multimedia technologies, however in this area too, the current network capacity is unable to fully exploit the benefits and power of these new technologies. One area of research that promises some solace is the MBone project (Multicast Backbone, currently available to restricted groups). It is used to allocate network resources based on the packet size and packets-per-second speed required by users multimedia content, and to multicast video and audio presentations to some Internet sites. It will soon be able to provide moderate video conferencing facilities to any individual or business with an internet connection, it currently operates at speeds between 56kbps and 1.5mbps.

Efforts are underway to enhance the security of the web as well, in recent years many a sensitive system has been compromised, however newer and more effective security protocols are under development, one among them is the SecureTCP project. Improved security takes on greater importance as merchants and corporations proceed to sell their wares on-line, different standards and means for performing transactions have been proposed, very soon terms such as Ecash (Electronic cash) will become common-place, maybe someday even replacing traditional forms of commerce.

Chapter 8

Glossary

ANS: Advanced Network and Services, responsible for operating NSFNET, done by MCI, IBM, Merit network. Contract ended in April 1995. Now backbone of internet is operated by MCI, AOL, ANS, Sprint.

ATM: Asynchronous Transfer Mode, a high speed digital switching and transmission technology that allows voice, video, and data signals to be sent over a single line at speeds ranging from 25 million to 1 billion bits per second (bps). An analog phone line transmits at about 2 million bps.

Archie: A system for locating files that are publicly available by anonymous FTP.

ARPAnet: An experimental network established in the 70's where the theories and software on which the Internet is based were tested. No longer in existence.

Bridge: A device that links two local-area networks together so they can share data.

Broadband: Any transmission at a speed higher than 2 million bps. A broad-band network can carry voice, data and video signals simultaneously.

Brouter: A cross between a bridge and a router.

CERN: European high energy physics laboratory in Geneva, Switzerland. Hypertext technologies developed at CERN to allow physicists to share information provided the basis for the World Wide Web.

CGI: Common Gateway Interface. This is the specification for how an HTTP server should communicate with server gateway programs. Most servers support the CGI specification.

DES: Data Encryption Standard, A private key encryption algorithm adopted as the federal standard for the protection of sensitive unclassified information and used extensively for the protection of commercial data as well.

Electronic data interchange (EDI): A series of standards that allows computer-to-computer exchange of business documents, such as invoices, shipping orders, payments, etc., between different companies.

Ethernet: A set of local-area network standards that allows networking products from different vendors to communicate. Introduced some 20 years ago, it is the most widely used LAN technology.

FDDI: A cabling system, based on the token-passing stndars used in IEEE 802.5, but with a few adaptations to suit fibre and to build a highly reliable dual ring mechanism.

Firewalls: Used to secure servers and networks on the internet. (perimeter security)

Frame relay: A transmission standard for sending data over public or private leased phone lines. Data is broken down and placed in frames, each the same size, for relaying.

Gopher: A menu-based system for exploring Internet resources.

HTML: Hyper text markup language, language in which World Wide Web documents are written.

HTML 3.0: Latest version of HTML. Look in *http://www.w3.org/hypertext/WWW/T* (for explanation).

HTTP: Hyper Text Transfer Protocol

IETF: Internet Engineering Task Force largely governs the Internet on an ad hoc basis, IETF, consisted of govt and academia, now commercial members also part of it.

Integrated services digital network (ISDN): Offered by local phone companies, this protocol turns a standard copper phone line into a high-speed digital link that can send voice and data simultaneously. With an ISDN line at home, users can pretend they're in the office using the LAN.

Internet service providers(commercial): BBN Planet, internetMCI, IBM Global Network.

IP 6: Updated internetworking protocol (to replace IP4), it will offer expanded addressing, simplified packet routing, improved message handling and security. To be implemented in 1996.

ISO: The International Organization for Standardization; an organization that has defined a different set of network protocols, called the ISO/OSI protocols. In theory, the ISO/OSI protocols will eventually replace the Internet Protocol. When and if this will actually happen is a hotly debated topic.

LAN: Local Area Network, usually local to a well-defined area. It is a small scale network.

MAC layer: All stations on a LAN are identified by a special address. This is known as the MAC-layer address or the physical-layer address. This address is used to physically identify a station on the LAN. It is called the MAC address, for it is at this sub layer (the Media Access Control, or MAC) that addressing is defined.

MBone: Multicast Backbone, used to allocate network resources based on the packet size and packets-per-second speed required by user's multimedia content. Multicast distributes video and audio presentations to some Internet sites. Will be able to provide moderate video conferencing facilities to any individual or business with internet connection. Currently operates at 56kbps - 1.5mbps. For more information refer: http://www.eit.com/techinfo/mbone/mbone.html

MPEG: Motion Picture Experts Group, is an acronym for a common video file compression method.

NCSA: National Center for Supercomputing Applications. The NCSA is situated at the Urbana-Champaign campus of the University of Illinois. The NCSA software developments team developed the programs Mosaic and the NCSA HTTPD server.

NSF: National Science Foundation. Responsible for funding and developing major parts of the Internet.

NSFNET: The National Science Foundation Network; the NSFNET is not the Internet. It's just one of the networks that make up the Internet.

PERL: Practical Extraction and Reporting Language, refer to *ftp://ftp.cis.ufl.edu* for information.

Private Branch Exchange (PBX): An automatic telephone switch for an internal phone system. It replaces the old-fashioned office switchboard and manages the voice-mail system.

Router: A device that links two LANs of different standards together, also referred to as a sophisticated bridge.

RSA: Is a very powerful public key algorithm that has resisted efforts at penetration. Named for it's developers, Rivest, Shamir and Adelman, who invented the algorithm at MIT in 1978, the algorithm uses two keys: a private key and a public key.

SGML: Standard General Markup Language. This is a standard for describing markup languages. HTML is defined as an instance of SGML.

SHTTP: Secure Hyper Text Transfer Protocol, provides security for documents used between two parties.

SONET: Synchronous Optical Network Technology. A high speed transmission arch, meant to exploit the huge bandwidth available on fibre-optic networks.

SSL: Secure Socket Layer, provides a secure communication channel between two parties on the internet. presently most popular method.

T1, T3: Standard for digital transmission over phone lines. Capacity of the internet is as follows,

T1=1.5 Mbps,

T3=45 Mbps

TCP/IP: Transmission Control Protocol/Internet Protocol. TCP/IP is the basic communication protocol that is the foundation of the Internet. All the other protocols, such as HTTP, FTP and Gopher, are built on top of TCP/IP.

UCL: University College London, the centre for the MICE project in multimedia conferencing.

URL: Uniform Resource Locator, it is the scheme used to address Internet resources on the World Wide Web.

vBNS: Very-High-Speed Backbone Network Service, focuses on broadband, internetworked technologies and services. The core of the project will be to improve the speed and scaling of the internet and it's underlying technologies. Provides 155Mbps service, by contrast the internet backbone presently operates at 45 Mbps.

Veronica: A service, very similar to Archie, that's built into Gopher. Just as Archie allows users to search all FTP sites for files, Veronica allows users to search all Gopher sites for menu items (files, directories, and other resources).

VideoVu: Provided by Future Communication Systems, allows users to send and receive audio signals over the Internet using standard modem technology, and a digital camera. Functions at 15fps.

WAIS: Wide Area Information Service; a very powerful system for looking up information in databases (or libraries) across the Internet.

WAN: Wide Area Network. Connects computers across borders or countries (large scale network).

Xerox PARC: Palo Alto Research Center.

Chapter 9

Literature Survey

In this chapter summaries of certain key publications that contributed to the thesis are presented.

9.1 Network technology

1. Robert M. Hinden, "IP Next Generation Overview"

This paper presents an overview of the Next Generation Internet Protocol (IPng). IPng was recommended by the IPng Area Directors of the Internet Engineering Task Force at the Toronto IETF meeting on July 25, 1994, and documented in RFC 1752 titled, "The Recommendation for the IP Next Generation Protocol". The recommendation was approved by the Internet Engineering Steering Group on November 17, 1994 and made a Proposed Standard.

IPng is a new version of the Internet Protocol, designed as a successor to IP version 4. IPng is assigned IP version number 6 and is formerly called IPv6. IPng was designed to take an evolutionary step from IPv4. Functions which work in IPv4 were kept in IPng. Functions which didn't work were removed. The changes from IPv4 to IPng fall primarily into the following categories:

- Expanded Routing and Addressing Capabilities
- Header Format Simplification
- Improved Support for Options
- Quality-of-service Capabilities
- Authentication and Privacy Capabilities

2. Joel F. Adam, Henry H. Houh, David L. Tennenhouse, "Experience with the VuNet: A Network Architecture for a Distributed Multimedia System"

The VuNet is a gigabit-per-second desk/local-area ATM network built in the context of the ViewStation distributed multimedia project proceeding at the Laboratory of Computer Science at MIT. The ViewStation system takes a software-intensive approach to multimedia where real-time data is processed at the application level and not in specialized hardware. The VuNet hardware has been specifically designed to support this software model. In particular, the network interface has been designed for ease of use. A key feature of the VuNet is that the more sophisticated functions typically found in ATM networks have been pushed out to the edges of the system to become the responsibility of the clients in exchange for simplicity and efficiency. This paper describes the VuNet hardware and presents the results of experiments in which video is transported across the system.

The following issues motivated the design of the hardware infrastructure, transparent network communication, bursty traffic (sources of video and audio must be able to provide their data in bursts, since applications based on the ViewStation process data in bursts) and graceful degradation of audio and video processing as workstations become scarce.

3. Ajit S. Thyagarajan, Stephen L. Casner, Stephen E. Deering, "Making the MBone Real"

IP multicasting is an extension of the Internet Protocol that efficiently delivers a single datagram to multiple hosts instead of a single host. It's benefits for applications such as live audio and video conferencing among Internet sites around the world have been clearly demonstrated over the past three years. This paper discusses some of the enhancements that are underway to make the MBone and IP multicasting real: scaleable routing, congestion management, and diagnostic tools.

4. Kamran Ghane, Neda Communications, Inc., "Internetworking with ATM-based Switched Virtual Networks"

Switched fabrics are going to play an important role in the future of internetworking. At present the major architecture for internetworking is router-based. Bridges are also important players in internetworking but considering the variety of protocols and lack of intelligence in bridges, routers are the major players. The problem with the existing router based internetworking architecture is limitation on speed. New applications, specifically multimedia and real-time applications demand more and more bandwidth and speed.

Virtual switch-based networks are one of the proposed solutions, since the latency in switch fabrics is much less than the latency of fast routers. The simple form of such switch fabrics is switched LAN. Switched LANs are basically layer 2 bridging. But the new generation of virtual networks is mostly based on ATM. ATM-based switched virtual LANs have different formats. ATM LAN emulation, ATM virtual routers, and ATM edge routers are different approaches that try to combine routing and switching.

ATM LAN emulation started with simply bridging LANs and now has started to support multi-protocol routing too. Edge routers and virtual routers are based on switching, routing, and virtual subnets. Relational networks combine multilayer relational switches with traditional routers. In this case the relational switches don't do the routing. Each of these different approaches to switched virtual networks has it's own pros and cons. Lots of standards have to be defined in this field. The ATM-based switched virtual network is used in all levels of internetworking, from small workgroups to large WANs.

5. Duane Wessels, "Intelligent Caching for World-Wide Web Objects"

The continued increase in demand for information services on the Internet is showing signs of strain. While the Internet is a highly distributed system, individual data objects most often have only a single source. Host computers and Network links can easily become overloaded when a large number of users access very popular data.

Proxy-caching is currently a popular way to reduce network bandwidth, server load and to improve response time to the user. The original caching proxy, form CERN, is probably still the most widely used. This paper describes software developed by the author that investigates some alternative techniques for caching World-Wide Web objects. This software complements traditional proxy-caching by allowing servers to explicitly grant or deny permission to cache an object, and with support for server-initiated callback invalidation of changed objects.

9.2 Multimedia Technology

1. Christopher J. Lindblad, David J. Wetherall, William F. Stasior, Joel F. Adam, Henry H. Houh, Mike Ismert, David R. Bacher, Brent M. Phillips and David L. Tennenhouse, "ViewStation Applications: Intelligent Video Processing Over a Broadband Local Area Network"

This paper describes applications built on the ViewStation, a distributed multimedia system based on UNIX workstations and a gigabit per second local area network. A key tenet of the ViewStation project is the delivery of media data not just to the desktop but all the way to the application program.The project group has written applications that directly process live video to provide more responsive human-computer interaction.The project group has also developed applications to explore the potential of media processing to support content-based retrieval of pre-recorded television broadcasts. These applications perform intelligent processing on video, as well as straightforward presentation. They demonstrate the utility of network-based multimedia systems that deliver audio and video data all the way to the application. The network requirements of the applications are a combination of bursty transfers and periodic packet-trains. Some applications described here are, the Office Monitor, and the Whiteboard Recorder.

2. Tamara Munzer, Paul Burchard, and Ed Chi, The Geometry Center, University of Minnesota, "Visualization through the World Wide Web with Geomview, Cyberview, W3Kit, and WebOOGL"

The Geometry Center has integrated it's visualization software into the web at many levels. Much of their effort is built around Geomview, the Center's extensible mathematical visualization package. It is a public domain 3D viewer which runs on SGI, NeXTStep, and standard X workstations and allows full interactive control of the motion and appearance of multiple objects and cameras. Among it's features are full support for non-Euclidean and higher dimensional spaces. At the most basic level of Web integration, Geomview can be run as a local external viewer for 3D scenes just as image viewers are used for 2D images.

3. Matthijs van Doorn, Anton Eli ns, "Integrating applications and the World Wide Web"

This paper describes how to integrate the World Wide Web (WWW) with applications. By means of the web widget, which is part of Hush (Hyper Utility Shell, is a C++ library that can be used to build applications with a graphical user interface), the WWW is made available to Tcl/Tk/Hush programmers. Apart form using WWW as part of an application, it also allows one to embed scripts into a web page. This results in a mutual integration of applications and the WWW.

Both forms of integration will be described. Some new possibilities of embedded scripts such as in-line MPEG, interactive games and navigation facilities will be discussed.

4. Joanna Mason, Marek Czernuszenko, Dana Plepys, Thomas A. DeFanti, "CAVEview, An Interactive tool for exploring virtual reality applications over the Internet."

CAVEview is an interactive tool for exploring virtual reality applications over the Internet, via Mosaic. Mosaic currently uses external viewers to present hypertext documents with images, text, sounds and animation. With CAVEview, virtual reality applications can be inserted into hypertext documents. CAVEview operates like other external viewers by bringing data files over the network; however, the transferred data is an application object file. The object file fully describes the environment, thus allowing the user to explore the scene for an unlimited amount of time.

Accessing applications program files lets the user have control over the story. This is in sharp contrast to being able to upload and view animations where the exploration of the scene is restricted to a fixed interval of a pre-recorded sequence. To attain a variation of that sequence, the animator would have to add more frames increasing the size of the data file that must be transferred. CAVEview adds stereo realism, interactivity and sound to viewing 3D objects in Mosaic without the cost of large animation files.

5. Keith Andrews, Frank Kappe, and Hermann Maurer, "Serving Information to the Web with Hyper-G"

The provision and maintenance of truly large-scale information resources on the World-Wide Web necessitates server architectures offering substantially more functionality than simply serving HTML files from the local file system and processing CGI requests.

This paper describes Hyper-G, a large-scale, multi-protocol, distributed, hypermedia information system which uses an object-oriented database layer to provide information structuring and link maintenance facilities in addition to fully integrated attribute and content search, a hierarchical access control scheme, support for multiple languages, interactive link editing, and point-and-click document insertion.

6. Tom Meyer, David Blair, and Suzanne Hader, "A MOO-Based Collaborative Hypermedia System for WWW"

This paper describes the development of a networked collaborative hypermedia system intended to support groups of writers and scholars in writing and publishing hypertext fiction and criticism. The current system supports the importation of individually-developed Storyspace hypertext documents into a MUD-based collaborative workspace for integration and expansion and allows for the immediate publication of these dynamically-generated multimedia documents on the World-Wide Web. In addition, the developers provided forms-based writing and linking interface to the text, so that writers can write using either the MUD-based or the forms-based authoring tools.

Since the HTML is generated dynamically from the underlying database, they have added the capabilities to: allow for user negotiation of content and bandwidth, and provide the bandwidth-intensive media from distributed mirror sites. This system is being used by dozens of students, writers, and theorists around the world to support such projects as: hypertext writing and theory classes.

7. K. J. Maly, H. Abdel-Wahab, R. Mukkamala, A. Gupta, A. Prabhu, "Mosaic + XTV = CoReview"

CoReview is an interactive document and data retrieval tool, it has been developed to provide a seamless environment to groups and individuals, distributed across the Internet, that want to interact on the progression of a project. It can also assist individuals to interactively put together a document in a collaborative manner. CoReview is based on the strengths of the World-Wide Web server, Mosaic, and XTV (an X-window teleconferencing system). While Mosaic will be used to manage the project documents and reviewer annotation files involved in proposals and their evaluation, XTV will aid in real-time remote collaboration among a group of users. CoReview incorporates the XTV features into a user friendly graphical interface and enables Mosaic to be shared by multiple networked participants for each document to be managed. The CoReview chair manages the shared resources. CoReview allows for easy creation of a pool of reviewers or proposal writers and automates the process of creating the necessary infrastructure daemons and directories at the needed sites.

8. Peter T. Kirstein, Mark Handley, Angela Sasse, Stuart Clayman, "Recent Activities in the MICE Conferencing Project"

The MICE project has been piloting multimedia conferencing (audio, video and shared workspace) over the Internet in Europe, with links to he US, over the last two years. In this paper, the group summarizes their experience and discuss recent developments in several areas: tool improvements, their use over ISDN and ATM, security, management and control, multimedia servers, new application areas and a recently created network of support centers.

9. Robert Voigt, Robert Barton, Shridhar Shukla, "A Tool for Configuring Multicast Data Distribution over Global Networks"

This paper addresses the challenge of constructing center specific multicast trees in a global network. The described approach is an interaction specific algorithm, for data distribution center location using asymmetric traffic loads on network links. A fast and scaleable algorithm for locating distribution centers based on the network load and participant's locations is given. This algorithm is presented as part of a tool which interactively allows a user to build networks, specify the group of multicast participants, select manually or generate algorithmically a center for the multicast tree and evaluate the tree quality. Simulation results on various topologies are presented showing that, with the above center location mechanism, algorithmically located center-specific trees yield lower tree cost and delay than center-specific trees with randomly selected centers.

10. William F. Stasior, "Vex: A Programming System for Extracting Content Information from Unstructured Media"

This paper describes Vex: a programming system for extracting content information from unstructured media. Vex poses the problem of understanding media in terms of recognizing the events in media streams. Vex approaches this problem by adapting computer science tools for matching patterns and parsing text. Vex allows the user to write an application directly in terms of patterns and actions. Vex transforms such a specification into a program that examines a stream of video, recognizes when a specified pattern of imagery has appeared, and performs the appropriate action.

Vex addresses important issues in both multimedia computing and computer vision. As a multimedia tool, Vex is a language for penetrating opaque media. In computer vision, Vex represents an approach to building general purpose tools.

11. C. J. Lindblad, "Using Tcl to Control a Computer-Participative Multimedia Programming Environment"

This paper describes how the VuSystem, a programming environment for the development of computer-participative multimedia applications, is controlled through Tcl scripts. In the VuSystem, networks of in-band media-processing modules are created and controlled by interpreted out-of-band Tcl scripts through object commands and callbacks. Tcl's extensibility, simple type system, efficient interface to C, and introspective capabilities are used by the VuSystem to produce a highly dynamic and capable media processing system.

12. David R. Bacher, Christopher J. Lindblad, "Context-based Indexing of Captioned Video on the ViewStation"

MIT's Telemedia Networks and Systems Group at the Laboratory of Computer Science have designed and constructed a mechanism for using caption text from broadcast television programs to analyze their content. In this paper, they describe the method by which captions are captured and translated from the raw video signal into text on the ViewStation. They also describe a tool called the Caption Parser, which analyzes the text and extracts information about the content of broadcast television programs. A specialized version of the parser is used by their Joke Markup program to index captioned monologues of late-night talk shows. The output of the markup program can be viewed through the Joke Browser, enabling selection and replay of jokes based on their content.

13. Joel F. Adam, David L. Tennenhouse, "The Vidboard: A Video Capture and Processing Peripheral for a Distributed Multimedia System".

This paper describes a stand-alone network-based video capture and processing peripheral (the Vidboard) for a distributed multimedia system centered around a gigabit/second Asynchronous Transfer Mode (ATM) network. The Vidboard captures video from an analog NTSC television source and transmits it to devices within the system. Devices control the Vidboard through a set of ATM protocols. Whereas capture boards typically generate video streams having fixed frame rate characteristics, the Vidboard is capable of decoupling video from the real-time constraints of the television world. This allows easier integration of video into the software environment of computer systems. The Vidboard is based on a front-end frame memory processor architecture that is also capable of generating full-motion video streams having a range of presentation (picture size, color space) and network (traffic, pixel packing) characteristics. A series of experiments are presented in which video is transmitted to a workstation for display. Frame rate performance and a remote video source control model are described.

14. D. Wetherall, "An Interactive Programming System for Media"

This paper presents PAVES, a direct manipulation system that combines aspects of visualization and multimedia systems to form an interactive video programming environment. PAVES extends the VuSystem media processing toolkit with dataflow-style views for controlling compute-intensive media applications as they run. Thus video may be manipulated graphically and interactively-both program and data are visual and live. PAVES is also novel in it's approach to extensibility. Users may freely combine graphical and underlying VuSystem textual programming methods to restructure and reuse applications. This cooperative programming is available across sessions as well as within them. It is implemented with an object-oriented programming foundation that automatically translates between multiple program representations and maintains them in synchrony.

9.3 Security Technology

1. Nicholas Baran, "Internet and Beyond - Security Truce?", BYTE magazine - July 1995.

Inadequate security is the biggest challenge to making the Internet a commercial marketplace. Because incidents of security breaches on the Internet are legion, many businesses and consumers are understandably uneasy about performing financial or other confidential transactions on-line. The two leading camps in security are i) SSL (Secure Sockets Layer) secures the transaction channel, ii) SHTTP (Secure Hypertext Transfer Protocol) secures documents that compromise the transaction.

Several articles from INET'95 were also researched and used in the thesis, many useful articles were found while conducting research on the Web, references to these are provided in chapter 10.

Chapter 10

References

10.1 General/World Wide Web overview

[1] Brendan P. Kehoe, (1992) "Zen and the Art of the Internet", 1st edition, World Wide Web resource.

[2] Bruce Sterling, "History of the Internet", World Wide Web resource.

[3] Jean Armour Polly, (1992) "Surfing the INTERNET: an Introduction Version 2.0.2", World Wide Web resource.

[4] Mitch Kapor, Steve Cisler, Adam Gaffin, (1991) "Welcome to the Big Dummy's Guide to the Internet", World Wide Web resource.

10.2 Publications on Browsing the World Wide Web

[5] Ari Luotonen, Tim Berners-Lee, (1994) "CERN httpd Reference Manual - A Guide to A World Wide Web HyperText Daemon", CERN.

[6] Dale Dougherty, Richard Koman, & Paula Ferguson, (1994) "The MOSAIC HANDBOOK", O'Reilly & Associates, Inc., 1st edition.

[7] Dave Raggett, "HyperText Markup Language Specification Version 3.0", Internet-Draft, Internet Engineering Task Force (IETF).

[8] Ed Krol, (1994) "The Whole Internet", O'Reilly & Associates, Inc., 2nd edition.

[9] Ian S. Graham, (1995) "HTML SOURCEBOOK", John Wiley & Sons, Inc., 1st edition.

[10] NCSA Mosaic for the X Window System User's Guide.

[11] T. Berners-Lee, R. T. Fielding, H. Frystk Nielsen, (1995) "Hypertext Transfer Protocol - HTTP/1.0", Internet-Draft, Internet Engineering Task Force (IETF).

10.3 Publications on Networking Technology

[12] Ajit S. Thyagarajan, Stephen L. Casner, Stephen E. Deering, (1995) "Making the Mbone Real", INET'95.

[13] Claffy K., Braun Hans-Werner, Gross Andrew, (1995) "Measured Interference of Security Mechanisms with Network Performance", INET'95.

[14] Duane Wessels, (1995) "Intelligent Caching for World-Wide Web Objects", INET'95.

[15] Graham Knight, Saleem N. Bhatti, Stuart Clayman, (1995) "A Data and Telecommunications Gateway between the Internet and ISDN", INET'95.

[16] Joel F. Adam, Henry H. Houth, David L. Tennenhouse, "Experience with the VuNet: A Network Architecture for a Distributed Multimedia System", MIT Laboratory for Computer Science.

[17] K. Washburn, J.T. Evans, (1993) "TCP/IP - Running a Successful Network", Addison-Wesley, 1st edition.

[18] Kamran Ghane, (1995) "Internetworking with ATM-based Switched Virtual Networks", INET'95.

[19] Mathew Naugle, (1994) "Network Protocol Handbook", McGraw-Hill, 1st edition.

[20] Robert M. Hinden, (1994) "IP Next Generation Overview", Internet Engineering Task Force.

10.4 Publications on Multimedia Technology

[21] C. J. Lindblad, (1994) "A Programming System for the Dynamic Manipulation of Temporarily Sensitive Data", MIT LCS Technical Report 637, MIT.

[22] Christopher J. Lindblad, David J. Wetherall, and David L. Tennenhouse, "The VuSystem: A Programming System for Visual Processing of Digital Video", Telemedia Networks and Systems Group, LCS, MIT.

[23] Christopher J. Lindblad, David J. Wetherall, William F. Stasior, Joel F. Adam, Henry H. Houh, Mike Ismert, David R. Bacher, Brent M. Phillips, and David L. Tennenhouse, "ViewStation Applications: Intelligent Video Processing Over a Broadband Local Area Network", Telemedia, Networks & Systems Group, Laboratory for Computer Science, MIT.

[24] D. Wetherall, (1994) "An Interactive Programming System for Media Computation", MIT LCS Technical Report 640, MIT.

[25] David L. Tennenhouse, Joel Adam, David Carver, Henry Houh, Michael Ismert, Christopher Lindblad, Bill Stasior, David Weatherall, David Bacher and Theresa Chang, "A Software-Oriented Approach to the Design of Media Processing Environments", Telemedia, Networks & Systems Group, Laboratory for Computer Science, MIT.

[26] David R. Bacher, Christopher J. Lindblad, (1995) "Content-based Indexing of Captioned Video on the ViewStation", submitted to ACM Multimedia 95.

[27] Gavin Bell, Anthony Parisi, Mark Pesce, "The Virtual Reality Modeling Language", World Wide Web resource.

[28] Joanna Mason, Marek Czernuszenko, Dana Plepys, Thomas A. DeFanti, "CAVE-view", Electronic Visualization Laboratory, University of Illinois at Chicago.

[29] Joel F. Adam, David L. Tennenhouse, "The Vidboard: A Video Capture and Processing Peripheral for a Distributed Multimedia System", MIT LCS.

[30] K. J. Maly, H. Abdel-Wahab, R. Mukkamala, A. Gupta, A. Prabhu, H. Syed, C. S. Vemru, "Mosaic + XTV = CoReview", World Wide Web resource.

[31] Keith Andrews, Frank Kappe, and Herman Maurer, "Serving Information to the Web with Hyper-G", Institute for Information Processing and Computer Supported New Media (IICM), Graz University of Technology, Austria.

[32] Matthijs van Doorn, Anton Eli ns, "Integrating applications and the World Wide Web", Vrije Universiteit, Netherlands.

[33] Peter T. Kirstein, Mark Handley, Angela Sasse, Stuart Clayman, (1995) "Recent Activities in the MICE Conferencing Project", INET'95.

[34] Robert Voigt, Robert Barton, Shridhar Shukla, "A Tool for Configuring Multicast Data Distribution over Global Networks", World Wide Web resource.

[35] Tamara Munzer, Paul Burchard, Ed Chi, "Visualization through the World Wide Web with Geomview, CyberView, W3Kit, and WebOOGL", The Geometry Center, University of Minnesota.

[36] Tom Meyer, David Blair, Suzanne Hader, "A MOO-Based Collaborative Hypermedia System", World Wide Web resource.

[37] William F. Stasior, (1995) "Vex: A Programming System for Extracting Content Information from Unstructured Media", submitted to ACM Multimedia 95.

10.5 Publications on Security Technology

[38] Deborah Russell and G. T. Gangemi Sr., (1993) "Computer Security Basics", O'Reilly & Associates, Inc., 3 rd. edition.

[39] Hans-Werner Braun, Kimberly Claffy, Andrew Gross, (1995) "Measured interference of network security mechanisms with network performance", Proc. INET'95.

[40] James M. Galvin, Sandra L. Murphy, "Using Public Key Technology-Issues of Binding and Protection", World Wide Web resource.

[41] John W. Verity, (June 26, 1995) "HACKER HEAVEN: So many computers so few safeguards", Business Week, New York.

[42] Jon Matonis, "Digital Cash and Monetary Freedom", World Wide Web resource.

[43] Jose Kahan, (1995) “A Distributed Authorization Model for WWW”, INET'95.

[44] Nicholas Baran, (July 1995) "Internet and Beyond - Security Truce?", BYTE magazine.

[45] Stephan Crocker, Brian Boesch, Brian Boesch, Alden Hart, James Lum, "Cyber-Cash: Payment Systems for the Internet", World Wide Web resource.

[46] Toshiyuki Tsutsumi, Suguru Yamaguchi, "Secure TCP - providing security functions in the TCP layer", Nara Institute of Science and Technology.

10.6 Publications on Data Integrity

[47] Ramez Elmasri and Shamkant B. Navathe, (1989) “Fundamentals of Database systems”, Benjamin/Cummings Publishing Company, Inc., 1st edition.

[48] Richard Y. Wang, Diane M. Strong and Lisa M. Guarascio, (1993) “An Empirical Investigation of Data Quality Dimensions: A Data Consumer’s Perspective”, Total Data Quality Management (TDQM) Research Program, Sloan School of Management.

[49] Richard Y. Wang and Lisa M. Guarascio, (1991) “Dimensions of Data Quality: Toward Quality Data by Design”, IFSRC Discussion Paper #CIS-91-06.

[50] Richard Y. Wang and M. Prabhakara Reddy, (1992) "Quality Data Objects", Total Data Quality Management (TDQM) Research Program, Sloan School of Management.

[51] Richard Y. Wang, Ward Page and Peter Kaomea, (1994) "Using Dominant Quality Curves to Produce Optimal Target Information", Total Data Quality Management (TDQM) Research Program, Sloan School of Management.

[52] Yair Wand and Richard Y. Wang, (1994) "Anchoring Data Quality Dimensions in Ontological Foundations", Total Data Quality Management (TDQM) Research Program, Sloan School of Management.