 

# Exhibit W

Intellectual Property Office
Woods Building Upper Level
100 Medical Center Way
San Francisco CA 94143-1016
Tel (415) 476-0833
Fax (415) 476-9926

University of California San Francisco A Health Sciences Campus UCSF

UNIVERSITY OF CALIFORNIA
RECEIVED
APR 11 1994
OFFICE OF
TECHNOLOGY TRANSFER

April 8, 1994

Linda Stevenson
Office of Technology Transfer
1320 Harbor Bay Parkway, Suite 150
Alameda, California 94502

Re: New Disclosure from Dr. Michael Doyle, David Martin, and Cheong S. Ang
Titled: **Embedded Program Objects in Distributed Hypermedia Systems**

Dear Linda,

Enclosed is the above referenced disclosure along with Dr. Doyle's cover letter which provides additional information to help expedite the marketing, development, and patenting of this important software invention with potential for profoundly impacting the advancement of Healthcare capabilities. Find attached responses to ROI item #s 2A & 2B, 3, 7, 8, and extra page for the third inventor's sign.

Attached as Appendix A is the manuscript submitted for Visualization '94, entitled "Integrated Control of Distributed Volume Visualization Through the World-Wide-Web"; and as Appendix B the Abstract for Dr. Doyle's paper presentation on 1/27/94, entitled "Medicine Meets Virtual Reality II: Interactive Technology & Healthcare".

Dr. Doyle's cover letter requests a formal Secrecy Agreement be established with Digital Equipment Corporation, where only a verbal agreement for nondisclosure has been covering his communications with this Company.

If I can be of any further assistance with the enclosed, please call.

Sincerely,

Anne Scott

Enclosures

Plaintiffs' Trial
Exhibit 6

E 003874

# DISCLOSURE AND RECORD OF INVENTION FORM

Note When completed, the *Disclosure and Record of Invention Form* is an important legal document. Care should be taken in its preparation Please refer to accompanying instructions If you desire assistance, call the Office of Technology Transfer (University Patent Office) at (415) 748-6600 Information contained in this document is maintained in confidence by the University Patent Office and normally will not be released to others except with attorney-client privilege, to research sponsors as required by contract, or under appropriate secrecy agreements, until a patent application is filed, the information is published, a determination not to file a patent application is made, or as may be required by law. The information contained should not be disclosed to others outside the University, except as described in section 9, without the approval of the University Office of Techology Transfer It is *not* the practice of the University Office of Technology Transfer to send your Record of Invention to other University employees for peer review

1. Short descriptive title of the invention
Embedded Program Objects in Distributed Hypermedia Systems.

2. A. Briefly summarize the invention here Include the novel features and advantages.
Please see attachment.



B. Detailed description of the invention using additional sheets as necessary and attach as appendix.

Please see attachment.

3. List the funding source(s) for the project under which this invention was made. If applicable, identify by contract or grant number and name the Principal Investigator/Supervisor of each. Please see attachment.

| Funding Source/Sponsor | Contract or Grant Number | Principal Investigator/ Supervisor |
|---|---|---|
| | | |
| | | |
| | | |
| | | |

4. For any "Inventor" named (item 13) who is not employed full-time by the University of California, please identify other employers (e.g., Veterans Administration, Howard Hughes Medical Institute, USDA), the percent of salary time funded by such other employer, and the nature of the other employment (such as research, teaching or clinical duties).
Not applicable.

5. When did you first conceive this invention?
This invention was originally conceived during a discussion between Dr. Doyle, Mr. Martin and Mr. Ang, at UCSF, on September 7, 1993.

6. What is the date of the first written record (notebook, letter, proposal, drawing, etc.) of this invention? Identify the document, page numbers involved, and location of the document.
The first written record relating to this invention was made in David Martin's research notebook on September 7, 1993.

7. When did you first successfully test this invention?
Please see attachment.

8. If you have disclosed this invention to non-UC personnel (including research sponsor) then indicate when, under what circmstances, and to whom.

a orally The internal workings of this process were discussed with Dr. Ingrid Carlbom, Head of the Visualization Group at DEC's Cambridge Research Laboratory, on February 8, 1994.

b in writing Source code for this invention, with a verbal agreement for non-distribution, was transferred, via FTP, to William Hsu, of Ingrid Carlbom's group at Digital, on February 22, 1994, in preparation for beginning a collaborative research project.

c. by actual use, demonstration, or posters The first demonstration of this invention to non-UC personnel was on November 16, 1993, to Dr. Donald Lindberg, in Room 101 of the UCSF Library Building. (Continued on attachment.)

—over—

9. Have you submitted or do you plan to submit a report, abstract, paper or thesis relating to this invention for publication, for presentation at a conference, or to a research sponsor?

Yes.

If yes, give details, including the actual or planned date of submission If a manuscript has been accepted, give the anticipated publication date. Append a copy of the latest draft manuscript available. (See instructions for the effect of publication prior to the filing of a patent application)

Please see attachment.

10. Identify any references, patent applications, or other publications of which you are aware and which you believe to be pertinent to this invention. Please attach a copy of each of these references, if available.

Please refer to the References section of the IEEE paper labelled Appendix A.

11. If any proprietary material (e g., cell line, antibody, plasmid, computer software, or chemical compound) obtained from outside your laboratory was used to develop this invention under a restrictive written or oral transfer agreement (other than a normal purchasing agreement), please attach a copy or summary of that agreement

Not applicable.

12 List companies you believe might be interested in using, developing or marketing this invention.

Silicon Graphics Corporation, Digital Equipment Corporation, Convex Computer Corporation, Microsoft, AT&T, and any company interested in the World Wide Web and the Internet.

13. Signatures, Names and Addresses of Inventors

| a) | | b) | |
|---|---|---|---|
| Signature | Date: 4-6-94 | Signature | Date: 4/6/94 |
| Print Name | Michael D. Doyle | Print Name | David C. Martin |
| Dept/ORU | Center for Knowledge Management | Dept/ORU | Center for Knowledge Management |
| Rm & Bldg | 117, Campus Library | Rm & Bldg | 112, Campus Library |
| Campus (Address if non-UC) | UCSF, Box 0840 | Campus (Address if non-UC) | UCSF, Box 0840 |
| City/State/Zip | San Francisco, CA 94143 | City/State/Zip | San Francisco, CA 94143 |
| Telephone | 415-476-8293 | Telephone | 415-476-6111 |

Note. If there are more inventors please provide signatures, names and addresses on an additional sheet of paper. Please see attachment.

14. Technically Qualified Witnesses (Two Required)—invention disclosed to and understood by

| a) | | b) | |
|---|---|---|---|
| Signature | Date: 4/6/94 | Signature | Date: 4/6/94 |
| Print Name | Benjamin F. Shang | Print Name | Marc E. Salomon |

Submit this form with ORIGINAL SIGNATURES directly to

Director—Office of Technology Transfer
XXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXX

Director, Office of Intellectual Property, UCSF
100 Medical Center Way, Box 1016
San Francisco, CA 94143
(415) 476-0833

If you do not receive an acknowledgement with 30 days, please call the University Office of Technology Transfer at (415) 748-6600.

NOTE. DISTRIBUTION OF COPIES OF A COMPLETED FORM TO THIRD PARTIES IS EXPRESSLY PROHIBITED, AS PROPRIETARY UNIVERSITY INFORMATION IS CONTAINED IN ANY COMPLETED FORM.

Effective 1/1/82
Revised 4/91

Retention. 7 yrs after last patent expires or 10 yrs after the date of the last action whichever is later



E 003876

**Attachment page 1 of 2**
**April 6, 1994**
**Invention Disclosure**
**Doyle, Martin, & Ang: "Embedded Program Objects in Distributed Hypermedia Systems."**

**2. A)** This invention comprises software methods which allow interactive program objects to be embedded within hypermedia documents in distributed hypermedia environments, such as the World Wide Web on the Internet. The first reduction to practice of this invention involved developing an extension to the HTML distributed hypermedia document type definition (DTD) and a set of enhancements to the World Wide Web (WWW) client and server programs to allow the embedding of links to self-contained interactive programs within HTML documents stored on WWW servers. When a user browsing the WWW selects such a link, that link is then decoded by a WWW client program, such as NCSA Mosaic (enhanced with our modifications), and the external program is run, mapping its display output to the correct location within the WWW client's display of the HTML document. With this invention, object linking and embedding (OLE) capabilities, such as are found in certain Microsoft Windows-based applications, can be implemented in distributed hypermedia systems and across wide area networks.

**2. B)** A detailed description of a reduction to practice of this invention is provided in Appendix A of this disclosure, which comprises an article, entitled "Integrated Control of Distributed Volume Visualization Through the Worldwide Web," authored by Cheong S. Ang, David C. Martin, and Michael D. Doyle.

**3)** This invention was developed within the UCSF Center for Knowledge Management in connection with the Center's preparation for a demonstration of the Visible Embryo Project, of which Dr. Doyle is the Principal Investigator. No external funds were used for this work.

**7)** This invention was successfully tested and demonstrated to Dr. Donald Lindberg, Director of the National Coordination Office for High Performance Computing and Communications, and Director of the National Library of Medicine, on November 16, 1993.

**8. c)** (Continued from Disclosure and Record of Invention Form.)

The first public demonstration of the invention was on January 14, 1994, in the Corporate Briefing Center, at the Mountain View, CA, headquarters of Silicon Graphics Corporation, to a group of SGI personnel. Videotape from that demonstration has been shown publicly at the following conferences: Virtual Reality Meets Medicine, San Diego Marriott, 1/28/94, and the 1994 meeting of the American Association for the Advancement of Science, San Francisco Hilton, 2/21/94.

A binary-code version of the demonstration of this invention was installed on the Reality Engine demonstration machine at Silicon Graphics' Corporate Briefing Center in Mountain View on 3/9/94, and is currently configured as one of the standard demonstration applications on that machine.

**Attachment page 2 of 2**
**April 6, 1994**
**Invention Disclosure**
**Doyle, Martin, & Ang: "Embedded Program Objects in Distributed Hypermedia Systems."**

**9)** An extended abstract (Appendix B) was published on 1/28/94 in the Proceedings of the "Virtual Reality Meets Medicine II" conference. This paper included a computer-screen image of the demonstration of this invention.

A paper detailing the reduction to practice of this invention (Appendix A) was submitted on 3/31/94 for inclusion in the Proceedings of "Visualization '94", an IEEE conference to occur on October 17-18, 1994, in Washington, D.C..

**13. c)** (Continued from Disclosure and Record of Invention Form.)

[signature] 4/7/94

Signature Date

Name: Cheong S. Ang
Department: Center for Knowledge Management
Room and Building: 115, Campus Library
Campus: UCSF, Box 0840
City/State/Zip: San Francisco, CA 94143
Telephone: 415-476-2954

********************************Nothing follows*******************************

APPENDIX A



# Integrated Control of Distributed Volume Visualization Through the World-Wide-Web

Cheong S. Ang, M S
David C. Martin, M.S.
Michael D. Doyle, Ph.D

University of California, San Francisco
Library and Center for Knowledge Management
530 Parnassus Avenue
San Francisco, California 94143-0840
contact. doyle@ckm.ucsf.edu

*Abstract*

The World-Wide-Web (W3) signals the entré of the digital library with the goal of providing immediate and ubiquitous access to digital information of any type from data repositories located throughout the world. The web's development enables not only effective access for the generic user, but also more efficient and timely information exchange among scientists and researchers. We have extended the repertoire of the web to include access to three-dimensional volume data sets with integrated control of a distributed client-server volume visualization system. This paper provides a brief background on the World-Wide-Web, an overview of the extensions necessary to support new data types and a description of the distributed visualization system.

## 1. Introduction

Advanced scanning devices, such as magnetic resonance imaging (MRI) and computer tomography (CT), have been widely used in the fields of medicine, quality assurance and meteorology [Pommert, Zandt, Hibbard]. The need to visualize resulting data has given rise to a wide variety of volume visualization techniques and computer graphics research groups have implemented a number of systems to provide volume visualization (e.g. AVS, ApE, Sunvision Voxel and 3D Viewnix)[Gerleg, Mercurio, Varde Wettering]. Previously these systems have depended on specialized graphics hardware for rendering and significant secondary storage for the data The expense of these requirements has limited the ability of researchers to exchange findings. To overcome the barrier of cost and provide additional means for researchers to exchange and examine three-dimensional volume data, we have implemented a distributed volume visualization tool for general purpose hardware and integrated that visualization service with the distributed hypermedia [Flanders, Broering, Kiong, Robison, Story] system provided by the World-Wide-Web [Nickerson].

Our distributed volume visualization tool, VIS, utilizes a pool of general purpose workstations to generate three dimensional representations of volume data. The VIS tool provides integrated load-balancing across any number of heterogenous UNIX™ workstations (e.g. SGI, Sun, DEC, etc...) [Giertsen] taking advantage of the unused cycles that are generally available in academic and research environments In addition, VIS supports specialized graphics hardware (e.g. the RealityEngine from Silicon Graphics) when available for real-time visualization

Distributing information that includes volume data requires the integration of visualization with a document delivery mechanism. We have integrated VIS and volume data into the W3, taking advantage of the client-server architecture of W3 and its ability to access hypertext documents stored anywhere on the Internet [Obraszka, Nickersen]. We have enhanced the capabilities of the most popular W3 client, Mosaic [Andressen] from the National Center for Supercomputer Applications (NCSA), to support volume data and have defined an inter-client protocol for communication between VIS and Mosaic for volume visualization

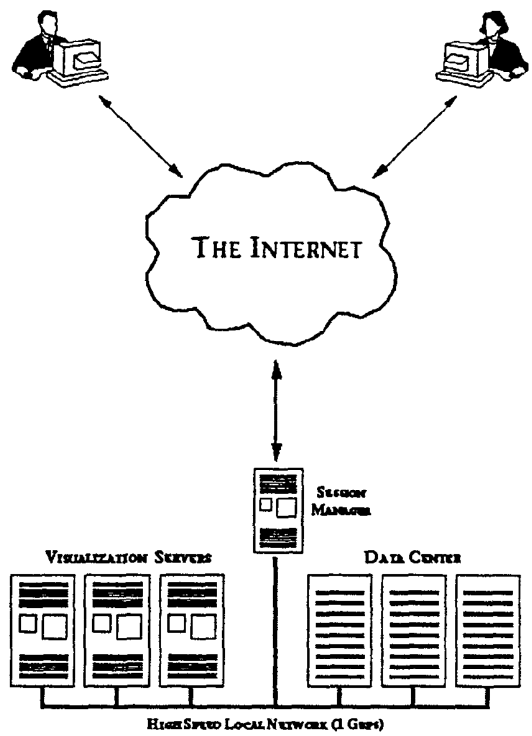


**Figure 1: VIS client/server model.**

### 1.1 The World-Wide-Web

The World-Wide-Web is a combination of a transfer protocol for hyper-text documents (HTTP) and a hyper-text mark-up language (HTML) [Nickersen]. The basic functionality of HTTP allows a client application to request a wide variety of data objects from a server. Objects are identified by a universal resource locator (URL)[Obraczka] that contains information sufficient to both locate and query a remote server. HTML documents are defined by a document type definition (DTD) of the Standard Generalized Mark-up Language (SGML). These documents are returned to W3 clients and are presented to the user. Users are able to interact with the document presentation, following hyper-links that lead to other HTML documents or data objects. The client application may also directly support other Internet services, such as FTP, Gopher, and WAIS, [Andreessen] or may utilize gateways that convert HTTP protocol requests and return HTML documents. In all interactions, however, the user is presented with a common resulting data format (HTML) and all links are accessible via URL's.

### 1.2 Mosaic

The National Center for Supercomputer Applications (NCSA) has developed one of the most functional and popular World-Wide-Web clients: Mosaic. This client is available via public FTP for the most popular computer interfaces (n.b. Motif, Windows and the Macintosh). Mosaic interprets a majority of the HTML DTD elements and presents the encoded information with page formatting, type-face specification, image display, fill-in forms, and graphical widgets. In addition, Mosaic provides inherent access to FTP, Gopher, WAIS and other network services [Andreessen].

## 1.3 VIS

VIS is a simple but complete volume visualizer. VIS provides arbitrary three-dimensional transformation (e.g. rotation and scaling), specification of six axial clipping planes (n.b. a cuboid), one arbitrary clipping plane, and control of opacity and intensity. VIS interactively transforms the cuboid, and texture-maps the volume data onto the tranformed geometry. It supports distributed volume rendering [Argrio, Drebin, Kaufman] with run-time selection of computation servers, and isosurface generation (marching cubes)[Lorenson, Levoy] with software Gouraud shading for surface-based model extraction and rendering. It reads NCSA Hierarchical Data Format (HDF) volume data files, and has a graphical interface utility to import volume data stored in other formats.

The rest of the paper is divided into five parts. Section 2 will describe VIS and how it fulfills the role as a W3 visualization tool. Section 3 is devoted to Mosaic, and will include the implementation of its interface with VIS. Section 4 presents the project results. Section 5 is about ongoing and future development, and the last section is conclusions.

## 2. VIS: A Distributed Volume Visualization Tool

VIS is a highly modular distributed visualization tool, following the principles of client/server architecture (figure 1), and consisting of three cooperating processes: VIS, Panel, and VRServer(s). The VIS module handles the tasks of transformation, texture-mapping, isosurface extraction, Gouraud shading, and manages load distribution in volume rendering. VIS produces images that are drawn either to its own top-level window (when running stand-alone) or to a shared window system buffer (when running as a cooperative process). The Panel module provides a graphical user-interface for all VIS functionality and communicates state changes to VIS. The VRServer processes execute on a heterogenous pool of general purpose workstations and perform volume rendering at the request of the VIS process . The three modules are integrated as shown in figure 3 when cooperating with another process. A simple output window is displayed when no cooperating process is specified.

### 2.1 Distributed Volume Rendering

Volume rendering algorithms require a significant amount of computational resources. However, these algorithms are excellent candidates for parallelization. VIS distributes the volume rendering among workstations with a "greedy" algorithm that allocates larger portions of the work to faster machines [Bloomer]. VIS segments the task of volume rendering based on scan-lines, with segments sized to balance computational effort versus network transmission time. Each of the user-selected computation servers fetches a segment for rendering via remote procedure calls (RPC), returns results and fetch another segment. The servers effectively compete for segments, with faster server's processing more segments per unit time, ensuring relatively equal load balancing across the pool. Analysis of this distribution algorithm [Giertsen, 93] shows that the performance improvement is a function of both the number of segments and the number of computational servers, with the optimal number of sections increasing directly with the number of available servers. Test results indictate that performance improvement flattens out between 10 to 20 segments distributed across an available pool of four servers. Although this algorithm may not be perfect, it achieves acceptable results.

### 2.2 Cooperative Visualizaton

The VIS client, together with its volume rendering servers, may be lauched by another application as a visualization server. The two requirements of cooperation are a shared window system buffer for the rendered image and support for a limited number of inter-process messages. VIS and the initiating application communicate via the ToolTalk service, passing message specifying the data object to visualize, options for visualization, and maintaining state regarding image display. The VIS Panel application appears as a new top-level window and allows the user control of the visualization tool.

## 3. Visualization with Mosaic

We have enhanced the Mosaic W3 browser to support both a three-dimensional data object and communication with VIS as a cooperating application (figure 2). Mosaic provides the user with the ability to locate and browse information available from a wide variety of sources including FTP. WAIS, and Gopher. HTTP servers respond to requests from clients, e.g. Mosaic, and transfer hypertext documents. Those documents may contain text and images as intrinsic elements and may also contain external links to any arbitrary data object (e.g. audio, video, etc...). Mosaic may also communicate with other Internet servers, e.g FTP, either directly – translating request results into HTML on demand – or via a gateway that provides translation services. As a W3 client, Mosaic communicates with the server(s) of interest in response to user actions (e.g. selecting a hyperlink), initiating a connection and requesting the document specified by the URL. The server delivers the file specified in the URL, which may be a HTML document or a variety of multimedia data files (for example, images, audio files, and MPEG movies) and Mosaic uses the predefined SGML DTD for HTML to parse and present the information. Data types not directly supported b Mosaic are displayed via user-specifiable external applications and we have extended that paradigm to both include three-dimensional volume data as well as to integrate the external application more completely with Mosaic.

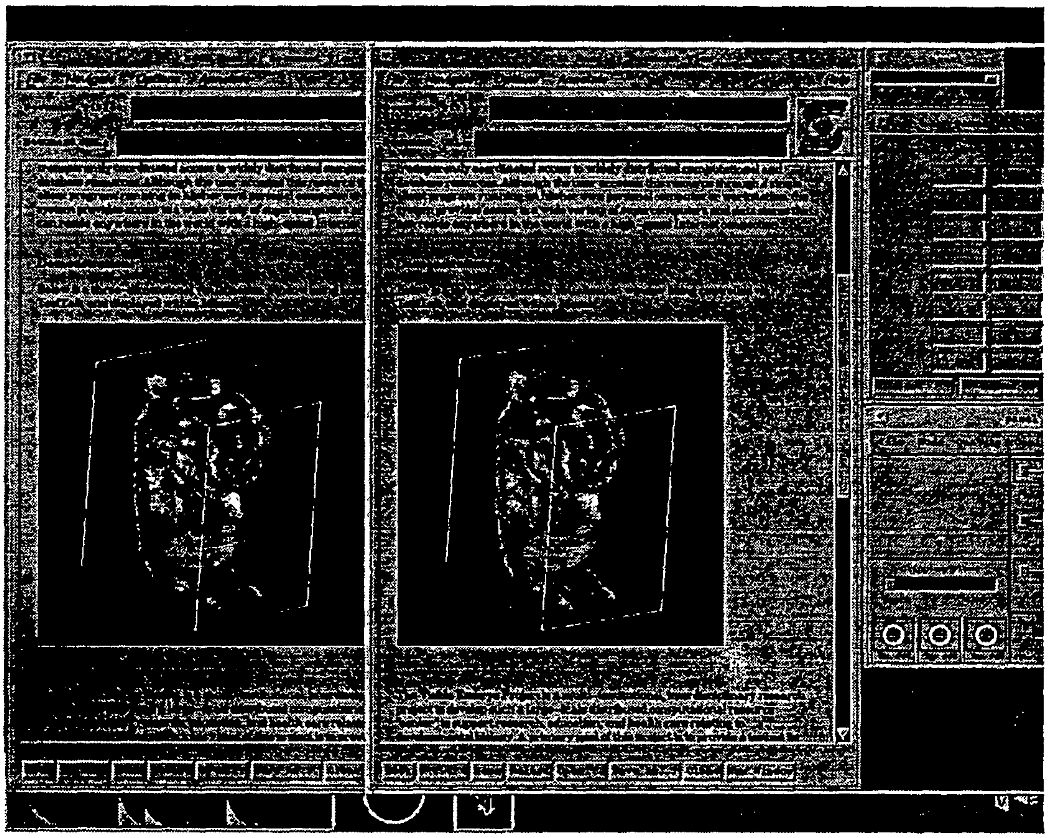

**Figure 2: Mosaic utilizing VIS to embed an interactive visualization within an HTML document. Here the data has been rotated 6 degrees in a "cloned" Mosaic window to create a stereo pair.**



E 003882

EOLASTX-0000000023

### 3.1 Mosaic 3D Image support

We have extended the HTML DTD to support three dimensional data via the introduction of a new SGML element: IMG3D. This element provides information to the presentation system (i.e. Mosaic) about the content that is referenced in the document. The IMG3D element is defined in the HTML DTD as follows:

```
<!ELEMENT    IMG3D EMPTY>
<!ATTLIST    IMG3D VOLUME CDATA #REQUIRED
             WIDTH NUMBER #REQUIRED
             HEIGHT NUMBER #REQUIRED
             IMAGE CDATA #IMPLIED>
```

which is translated as "SGML document instance element tag IMG3D containing no content; three required attributes: volume data set, window width and height; and an optional image". In a HTML document, a 3D image element would be represented as:

```
<IMG3D       VOLUME="http://www.library.ucsf.edu/.../data/Embryo.hdf"
             WIDTH=400
             HEIGHT=400
             IMAGE="http://www.library.ucsf.edu/.../images/Embryo.gif">
```

which may be interpreted as "create a 3D visualization window of width 400 pixels, height 400 pixels, and visualize the data embryo.hdf located at the HTTP server site www.library.ucsf.edu".

### 3.2 Interface with Mosaic

The VIS/Mosaic software system consists of three elements: VIS, Mosaic, and Panel (which provides GUI for data manipulation). Currently, VIS communicates with Mosaic via ToolTalk™, but the system will work with any inter-client communication protocol. Whe Mosaic interprets the HTML tag IMG3D, it creates a drawing area widget in the document page presentation and creates a shared window system buffer to receive visualization results. In addition, Mosaic launches the VIS application, specifying the location of the data object to render and identifying the shared image buffer. When the VIS process begin execution, it verifies its operating parameters and launches the Panel application that in turn presents the user with the control elements for manipulating the visualization and manages the communication between VIS and Mosaic.

In addition to coordinating communications, the Panel application attempts to balance the rendering load across the selected VRServer(s), sending rendering parameters to the selected server(s), and integrating the resulting image segments(s). Thus the scenario following a user's action on the Panel will be (1) Panel sends rendering requests to VRServer(s) with the parameters from its GUI, (2) Panel fetches the returned image data, then writes it to the pixmap, (3) Panel sends a message to notify Mosaic upon completion, and (4) Mosaic bit-blots the image in the pixmap into its DrawingArea widget. This will be addressed in more details under section 3.2. The configuration of this software system is depicted in figure 3.

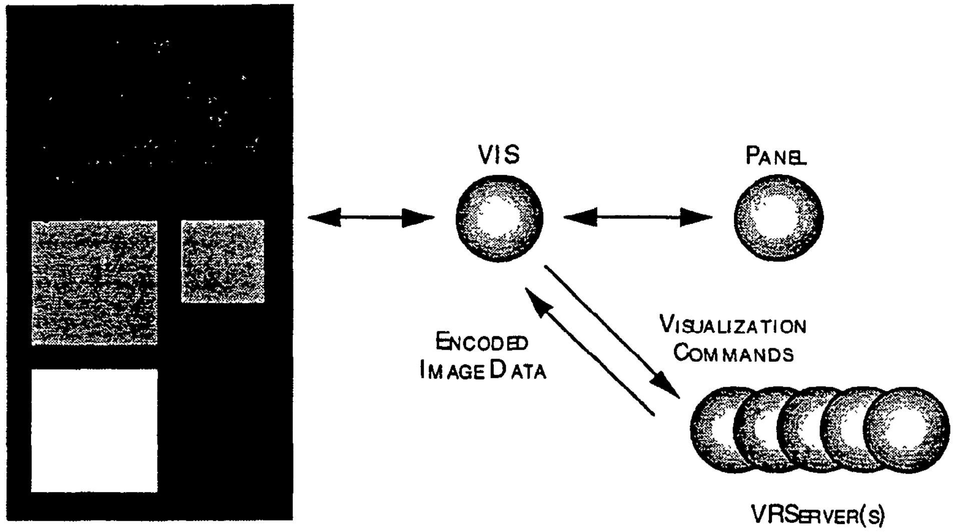


**Figure 3: Interprocess communication among Mosaic, VIS and distributed rendering servers.**

### 3.3 Interclient communication

We recognized the minimum set of communication protocols between Mosaic and Panel:

(a) Mosaic notifies Panel when the window id is no longer valid. This happens when the user navigates from the current page to the previous (with BACK button), or to another document through a link.

(b) Panel notifies Mosaic when the scene needs to be updated. This occurs when the user manipulates the data through the Panel's GUI, which may be any actions (rotations, scaling, changes in clipping planes, mode switching, for instance, from texture mapping to volume rendering) that result in rerendering of the picture.

The first case of (a) does not impose any problem since Mosaic can simply terminate VIS (this solution also applies when Mosaic is quitting). The second case is harder to handle because Mosaic caches the page, and in principle, we expect VIS to remain running because it should not be loading the HDF data, which may be quite a few megabytes, again. In the previous case, we simply have Mosaic kill the child process that launched Panel, which may subsequently send a terminate message to the VIS servers. In the latter case, Mosaic will send a message to Panel requesting it to unrealize/unmap its window. Then when the user comes back to the 3D object page, Mosaic may send another message to Panel ordering it to realize the GUI window, and update the window id of the DrawingArea that Panel has been sending message to (ironically Mosaic does not cache the widgets/windows on the cached page).

In (b), Panel will send an update message directly to the DrawingArea window id provided by Mosaic. Upon receiving the redrawing message, the DrawingArea widget calls its registered refreshing function, which simply bit-blots the pixmap Panel updated into the window. The protocols are summarized in Table 1.

| MESSAGES | MOSAIC NOTIFIES VIS | VIS RESPONDS TO MOSAIC |
|---|---|---|



E 003884

EOLASTX-0000000025

| | | |
|---|---|---|
| Refresh | Compute Image | Done Drawing |
| PageCaching | Hide Panel | Panel Hidden |
| BacktoCachedPage | New Window ID | ID Changed, Panel Shown |
| DestroyPage | Terminate | Terminating |

**Table 1: Mosaic/VIS IPC communication.**

## 4. Results

The results of the above implementation are very encouraging. The Mosaic/VIS sucessfully allows users to visualize HDF volume datasets from various HTTP server sites. Fig 2 shows a snapshot of the W3 visualizer. Distributing the volume rendering loads results in a remarkable speedup in image computations. However, we did not do any detail performance analysis, because the results of such an analysis would not be reproducable for varying network traffics loads, and fluctuating server workstation loads. The server pool we tested the system consists of heterogenous workstations: a SGI Indigo2 R4400/150MHz, two SGI Indy R4000PC/100MHz, a DEC Alpha 3000/500 with a 133MHz Alpha processor, two Sun SparcStations 10, and two Sun SparcStations 2, which were located arbitrarily on a Ethernet network. To our knowledge this is the first demonstration of the embedding of interactive control of a client/server visualization application within a multimedia document in a distributed hypermedia environment, such as the World Wide Web.

## 5. Ongoing/Future work

We have begun working on several extensions and improvements on the above software system:

### 5.1 MPEG Data Compression

The data transferred between the visualization servers and the clients constitue the exact byte streams computed by the servers, packaged in the XDR machine independent format. One way to reduce network transferring time would be to compress the data before delivery. We propose to use the MPEG compression technique, which will not only perform redundancy reduction, but also a quality-adjustable entropy reduction. Furthermore, the MPEG algorithm performs interframe, beside intraframe, compression. Consequently, only the compressed difference between the current and the last frames is shipped to the client.

### 5.2 Generalized External-Application-to-Mosaic-Document-Page Display Interface

The protocols specified in Table 1 are simple, and general enough to allow most image-producing programs be modified to display in the Mosaic document page. Our undertakings will be to extend the protein database (PDB) displaying program, and the xv 2D image processing program, to be Mosaic PDB visualization server, and Mosaic 2D image processing server.

### 5.3 Multiple Users

With multiple users, the VIS/Mosaic distributed visualization system will need to manage the server resources, since multiple users utilizing the same computational servers will slow the servers down significantly. The proposed solution is depicted in Fig 4. The server resource manager will allocate servers per VIS client request only if those servers are not overloaded. Otherwise, negotiation between the resource manager and the VIS client will be necessary, and, perhaps the resource manager will allocate less busy alternatives to the client.

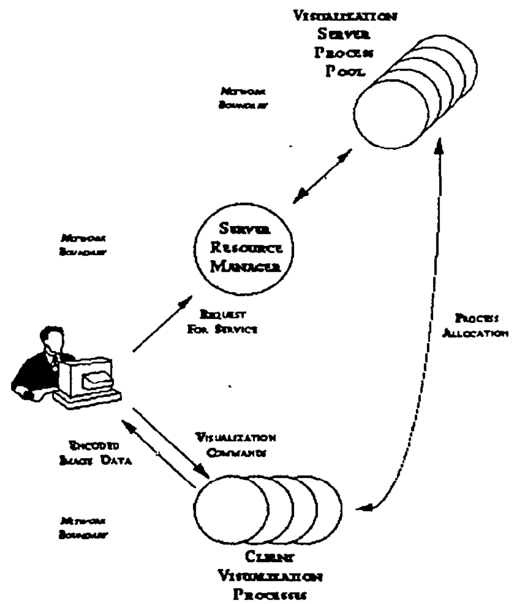


**Fig 4: Server Resource Management**

### 5.4 Load Distributing Algorithm

Since the load distributing algorithm in the current VIS implementation is not the most optimal load distribution solution, we expect to see some improvement in the future implementation, which will be using sender-initiated algorithms, described in [Shivaratri].

## 6. Conclusions

Our accomplishment takes the technology of networked multimedia system (especially the World Wide Web) a step further by proving the possibility of adding new interactive data types to the $W^3$ servers and clients, and coordinating the execution of the applications that handle them with the $W^3$ clients. The addition of the 3D volume data object in the form of HDF to the $W^3$ is welcomed by many medical researchers, for it is now possible for them to view volume datasets without a high-cost workstation, and accessing datasets in the $W^3$ fashion, through hypertext and hypergraphics links within an HTML page. As for the researchers who would like to share their findings with the world, they merely have to run a $W^3$ server that serves the HDF files, which in turn, can be easily created with the various import facilities available freely from NCSA.

## 7. References


Andreessen, Marc, "NCSA Mosaic Technical Summary", FTPed from ftp.ncsa.uiuc.edu, Revision 2.1, May 8, 1993.

Argiro, V. "Seeing in Volume", Pixel, July/August 1990, 35-39.

Avila, R., Sobierajski, L. and Kaufman A., "Towards a Comprehensive Volume Visualization System", Visualization '92 Proceedings, IEEE Computer Society Press, October 1992, 13-20.

Bloomer, John, "An RPC Case Study: Networked Ray Tracing", Power Programmingwith RPC, O'Reilly & Associates, Inc., Sept 92, 409-451.

Brinkley, J.F., Eno, K., Sundsten, J.W., "Knowledge-based client-server approach to stuctural information retrieval: the Digital Anatomist Browser", Computer methods and Programs in Biomedicine, Vol. 40, No. 2, June 1993, 131-145.

Broering, N. C., "Georgetown University, The Virtual Medical Library," Computers in Libraries, Vol. 13, No. 2, February 1993, 13.

Drebin, R. A., Carpenter, L. and Hanrahan, P., "Volume Rendering", Computer Graphics, Vol. 22, No. 4, August 1988, 64-75.

Flanders, B., "Hypertext Multimedia Software: Bell Atlantic DocuSource", Computers in Libraries, Vol 13, No. 1, January 1993, 35-39.

Gelerg, L., "Volume Rendering in AVS5", AVS Network news, Vol. 1, Issue 4, 11-14.

Giertsen, C. and Petersen, J., "Parallel Volume Rendering on a Network of Workstations", IEEE Computer Graphics and Applications, November 1993, 16-23.

Jäger, M., Osterfeld, U., Ackermann, H, and Hornung, C., "Building a Multimedia ISDN PC", IEEE Computer Graphics and Applications, September 1993, 24-33.

Kaufman, A., Cohen, D., and Yagel, R., "Volume Graphics", Computer, July 1993, 51-64.

Kiong B., and Tan, T., "A hypertext-like approach to navigating through the GCG sequence analysis package", Computer Applications in the Biosciences, Vol. 9, No. 2, 1993, 211-214.

Levoy, M., "Display of Surfaces from Volume Data", IEEE Computer Graphics and Applications, Vol. 8, No. 5, May 1988, 29-37.

Lorensen, W., Cline, H.E., "Marching Cubes: A High Resolution 3D Surface Construction Algorithm", Computer Graphics, Vol. 21, No. 4, July 1987, 163-169.

Mercurio, F., "Khoros", Pixel, March/April 1992, 28-33.

Narayan, S., Sensharma D., Santori, E.M., Lee, A.A., Sabherwal, A., Toga, A.W., "Animated visualization of a high resolution color three dimensional digital computer model of the whole human head", International Journal of Bio-Medical Computing, Vol 32, No. 1, January 1993, 7-17.

Nickerson, G., "WorldWideWeb Hypertext from CERN", Computers in Libraries, Vol. 12, No. 11, December 1992, 75-77.

Obraczka, K., Danzig, P, and Li, S., "Internet Resource Discovery Services", Computer, Vol. 26, No. 9, September 1993, 8-22.

Pommert, A., Riemer, M., Schiemann, T., Schubert. R., Tiede, U., Hoehne, K-H, "Methods and Applications of Medical 3D-Imaging", SIGGRAPH 93 course notes for volume visulization, 68-97.

Robison, D., "The Changing States of Current Cites: The Evolution of an Electronic Journal", Computers in Libraries, Vol. 13, No. 6, June 1993, 21-26.

Shivaratri, N.G., Krueger, P., and Singhal, M., "Load Distributing for Locally Distributed Systems", Computer, December 1992, 33-44.

Singh. J, Hennessy, J. and Gupta A., "Scaling Parallel Programs for Multiprocessors: Methodology and Examples", Computer, July 1993, 42-49.

Story, G., O'Gorman, L., Fox. D, Schaper, L. and Jagadish, H.V., "The RightPages Image-Based Electronic Library for Alerting and Browsing", Computer, September 1992, 17-26.

VandeWettering, M., "apE 2.0", Pixel, November/December 1990, 30-35.

Woodward, P., "Interactive Scientific Visualization of Fluid Flow", Computer, Vol. 26, No. 10, June 1993, 13-25.

Zandt, W.V., "A New 'Inlook' On Life", UNIX Review, Vol 7, No. 3, March 1989, 52-57.

# Medicine Meets Virtual Reality II: Interactive Technology & Healthcare

## The Virtual Embryo: VR Applications in Human Developmental Anatomy

Authors:

Michael D. Doyle, Ph.D.
Director, Center for Knowledge Management
University of California, San Francisco
530 Parnassus Ave, Box 0840
San Francisco, CA 94143-0840
(415)476-9742
doyle@ckm.ucsf.edu

Adrianne Noe, Ph.D.
Curator, Human Developmental Anatomy Center
National Museum of Health and Medicine
Armed Forces Institute of Pathology
Walter Reed Army Medical Center, Bldg. 54
6825 Alaska Ave. N.W.
Washington, D.C. 20306-6000
(202)576-0401
NOE@padstars.afip.osd.mil

Ingrid Carlbom, Ph.D.
Director, Visualization Group
Cambridge Research Laboratory
Digital Equipment Corp.
1 Kendal Square, Bldg. 700
Cambridge, MA 02139
(617)621-6629
carlbom@crl.dec.com

Cheong Ang, M.S.
Programmer, Innovative Software Systems Division
Center for Knowledge Management
University of California, San Francisco
530 Parnassus Ave, Box 0840
San Francisco, CA 94143-0840
(415)476-8293
ang@ckm.ucsf.edu

David Martin, M.S.
Head, Innovative Software Systems Division
Center for Knowledge Management
University of California, San Francisco
530 Parnassus Ave, Box 0840
San Francisco, CA 94143-0840
(415)476-8293
martin@ckm.ucsf.edu

## The Virtual Embryo: VR Applications in Human Developmental Anatomy

If virtual reality applications in biology and medicine are going to fufill their promise of recreating the subtleties and complexities of biological environments, they must draw from a rich assortment of heterogeneous and highly dense databases containing information about both the structures and functions of the systems to be simulated. This presents a problem for workers in virtual reality, for there are few readily available resources for accurate and detailed data on human biological structure and physiology. Several national-level research initiatives, such as the Visible Human Project and the Human Brain Project, are taking great strides towards creating information resources that can fill this gap. These projects are creating, for the first time, canonical datasets on adult human anatomic topography and topology which will serve as reference datasets for many areas of research into the applications of computer technology to biomedicine. The Visible Embryo project is an attempt to create such information resources for the fields of developmental and molecular biology. It also represents an integrated approach to developing applications of this technology in the areas of image databases, software tools, educational applications, and both basic and clinical sciences. A critical component of this research involves the development of virtual reality applications for research and education in the basic science of human developmental anatomy, as well as surgical simulation tools for practitioners of pediatric medicine.

The term "metacenter" was coined several years ago by Larry Smarr, the Director of the National Center for Supercomputing Applications, to propose network-based computing and information "centers that actually represent transparently coordinated access to widely distributed computational and database resources. Although the user accesses this resource through a single client program running on his/her own personal computer or workstation, the system provides access to a wide variety of supercomputers and data stores that can reside anywhere in the world with a high-speed Internet connection.

The Visible Embryo Project is a multi-institutional, interdisciplinary research project to develop such a large-scale distributed computational resource "center" to support research, education, and health care relating to developmental biology. A primary goal of this project is to provide a testbed for the development of new technologies, and the refinement of existing ones, for the application of high-speed, high-performance computing and communications to current problems in biomedical science.

Sets of serial microscopic cross-sections through human embryos, within the collection of the National Museum of Health and Medicine, will be digitized and processed to create volumetric reconstructions of normal human embryonic anatomy. During the five years of this initial project, the entire contents of the Museum's Carnegie Collection of Human Embryology, over 600 embryos, will be digitized, reconstructed and archived, together with case histories, scientific articles, research notes, didactic descriptions, and other data contained within the collection. This massive database will be housed at the Museum at Washington D.C., while teams of researchers at more than 20 universities and companies around the United States will access widely distributed super computing

E 003890

resources to develop visualization, analysis and telecollaboration software tools, educational materials, virtual reality simulations, basic science investigations, and clinical research projects based upon the data contained within the collection.

This project will serve the dual purpose of providing a testbed for new technology development in high performance computing and communications, as well as creating powerful new tools for the developmental biology research community. New advances in visualization technology are beginning to allow investigators to break through previous technical limitations and discover universally-applicable rules for pattern formation and shape development in organisms. By applying these new technologies to the existing archives of cross-sectional image information that exist in the literature and in collections around the world, we can tap into an enormous amount of new information that can be extracted from these databases.

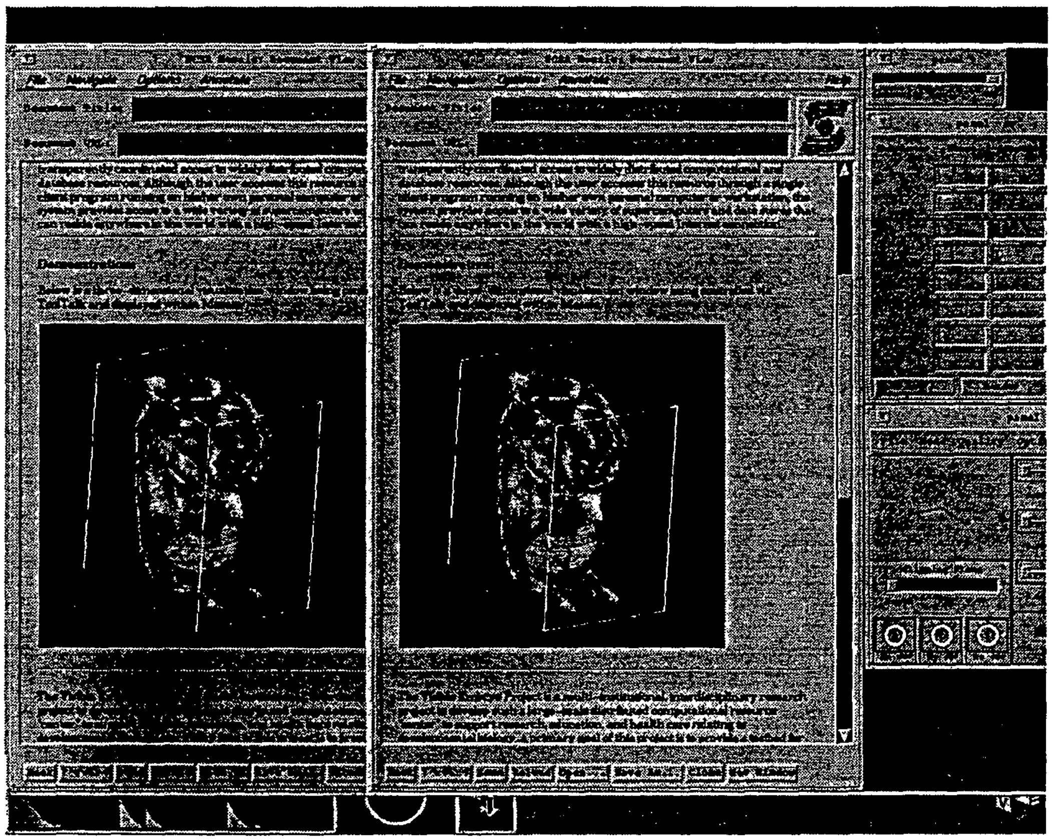

*This image illustrates UCSF's integration of Mosaic documents with remote access to real-time interactive volume visualization tools. The two embryo images are slightly rotated to form a stereo-pair image which can be viewed in 3D by allowing your eyes to diverge slightly so that three images apear, and then focusing in on the middle image (it takes practice). The interactive visualization controls are to the right.*

The task of integrating access to such massive information and computational resources is nontrivial. Just one embryo from the 650 serially-sectioned specimens in the collection can yield as much as a terabyte of anatomical volume data (a 20mm specimen, sectioned at 5 microns and digitized at a resolution of 8000x8000 pixels/section at 36bits RGB produces 1.073 TB of voxel data). It is clear that no single workstation, or even supercomputer, can manipulate, process and analyze such a large quantity of data as a single unit, much less perform computational operations on a database of hundreds of such datasets. For this reason, the Visible Embryo team at UCSF has developed tools to allow integrated Internet access (through NCSA Mosaic) to remote volume visualization engines which can distribute computation across a large number of graphics supercomputers connected by high-speed networking. This allows the integration, through NCSA Mosaic and the World Wide Web, of text-based, image, audio, and video data with real-time interactive control of high-performance visualizations embedded within Mosaic documents. We are also exploring the potential of using this technology for delivering interactive access to virtual reality applications through the Internet. The success of these efforts will allow widespread access to highly accurate and complex virtual reality simulations of human developmental anatomy, using inexpensive workstations and personal computers.