

# **UCSF Radiological Informatics Research**

## **- A Progress Report -**

**February 1996**

**Laboratory for Radiological Informatics  
Rm CL-158, 530 Parnassus Avenue  
Department of Radiology  
University of California, San Francisco  
San Francisco, CA 94143-0628**

**(Tel) (415) 476-5941  
(Fax) (415) 502-3217**

## SUMMARY

During the past year, we concentrated our efforts in putting the hospital-integrated PACS in operation and developing a research program in image informatics. With respect to PACS, we connected five MR, five CT scanners from multiple sites, two computed radiography systems, two film digitizers, one US PACS module, the hospital HIS and the radiology department RIS to the PACS network. The image data is managed by a mirrored database. The PACS controller, with its 1.3 terabyte optical disk library, acquires 2.5 gigabytes digital data daily. Four 2K, five 1,600-line multiple monitor display workstations are on line in neuroradiology, pediatric radiology and intensive care units for clinical use. In addition, the PACS supports over 100 Macintosh users in the department and selected hospital sites for both images and textual retrieval through a client/server mechanism.

In the research program, we are developing a computation and visualization node in the PACS network for image informatics research. Emphasis has been on image content-based indexing retrieval. Several large-scale projects have been launched, for example, non-invasive epilepsy surgical planning using multi-modality brain images, timeline quantitation of lung nodule volumes over treatment period with spiral CT, and bone age assessment with digital hand radiographs and CR.

This report summarizes the current status of PACS and image informatics research in our laboratory including image processing, PACS, integration of database, data and knowledge extraction, and image compression. It contains preprints of papers which will appear in the Proceedings of *International Society for Optical Engineering, SPIE Medical Imaging 1996*: volumes 2707, 2709, and 2711, Newport Beach, California, February 10-15, 1996. This report also contains selected reprints published during the last year; abstracts and presentations from the 81st RSNA, November - December, 1995; the editorial and Table of Contents of the Special Invited Issue on Medical Image Database, *J. Comp. Medical Imaging and Graphics*, 1996 by Steven Wong and Bernie Huang, and the Table of Contents of the new book: *PACS in Biomedical Imaging*, 1996 by Bernie Huang.

### **This work is currently supported by:**

- PHS grant RO1 CA 51198
- NLM contract NO1-LM-4-3508
- NCHS BPA 50257C-95, National Center for Health Statistics
- CalREN grant ATMN-007, Pacific Bell
- US Army Medical R&D DAMD 17-94-J-4196
- US Army Medical R&D DAMD 17-94-J-4338
- Lumisys (equipment)
- Fuji Medical Systems USA, Inc. (equipment)
- Abe Sekkei, Japan (equipment)
- Polaroid
- Philips Medical Systems
- California Breast Research Program

## TABLE OF CONTENTS

Summary.....	i
Table of Contents.....	ii
Title Page .....	1
Laboratory for Radiological Informatics Staff.....	2
Wide Area Network in the San Francisco Bay Area.....	3
ATM WAN and LAN connecting UCSF and Mt. Zion Hospital.....	4
Department of Radiology, UCSF Network Architecture.....	5

### PREPRINTS FROM SPIE

#### Image Display

1. <b>Jun Wang, H.K. Huang</b> .....	6
In-line structure lossless digital mammographic image compression [2707-48]	
2. <b>Stephen T.C. Wong, P.S. Whaley, A. Cheong, K. Soo Hoo, J. Wang</b> .....	14
Interactive query and visualization of medical image data on the World Wide Web [2707-40]	

#### Physiology and Function From Multidimensional Images

3. <b>Stephen T.C. Wong, R. Knowlton</b> , .....	26
Use of image co-registration and visualization techniques to study relationships between MEG neurophysiology and FDG-PET metabolism in epilepsy imaging [2709-26]	
4. <b>Xiaoming Zhu, S.T.C. Wong, H.K. Huang</b> .....	27
Quantitative image measurements for outcome analysis of lung nodule treatment [2709-45]	

#### PACS Design and Evaluation: Engineering and Clinical Issues

5. <b>Joseph K. Lee, A.W.K. Wong, H.K. Huang, T.M. Bazzill, J. Zhang, K.P. Andriole</b> .....	28
An ATM distributed PACS server for ICU applications [2711-03]	
6. <b>Albert W.K. Wong, H.K. Huang, J.K. Lee, T.M. Bazzill, X. Zhu</b> .....	36
High-performance image communication network with asynchronous transfer mode technology [2711-07]	

7.	<b>Katherine P. Andriole, R.C. Brasch, C.A. Gooding, R.G. Gould, H.K. Huang</b> .....	45
	Clinical comparison of CR and screen-film for imaging the critically ill neonate [2711-35]	
8.	<b>Katherine P. Andriole, M.L. Storto, M.A. Linnolt, G. Gamsu, H.K. Huang</b> .....	50
	Impact and utilization studies of a PACS display station in an ICU setting [2711-36]	
9.	<b>Jianguo Zhang, S.T.C. Wong, K.P. Andriole, H.K. Huang</b> .....	54
	Real time multilevel process monitoring and control of CR image acquisition and preprocessing for PACS and ICU[2711-37]	
10.	<b>Stephen T.C. Wong, K. Soo Hoo, H.K. Huang</b> .....	62
	Empower PACS with content-based queries and 3D image visualization [2711-70]	
11.	<b>Todd M. Bazzill, H.K. Huang, G.R. Thoma, R. Long, M. Gill</b> .....	73
	High-speed image transmission via a communication satellite [2711-40]	
12.	<b>H.K. Huang, R.L. Arenson, A.W.K. Wong, T.M. Bazzill, A.S.L. Lou, K.P. Andriole, J. Wang, J. Zhang, S.T. Wong</b> .....	82
	Current status of the UCSF second generation PACS [2711-46]	

#### SELECTED REPRINTS

13.	<b>Shyh Liang Lou, J. Wang, M. Moskowitz, T. Bazzill, H.K. Huang</b> .....	94
	Methods of automatically acquiring images from digital medical systems. <i>Comp. Med Imag. &amp; Graphics</i> 19:369-376, 1995.	
14.	<b>Michael F. McNitt-Gray, H.K. Huang, J.W. Sayre</b> .....	102
	Feature selection in the pattern classification problem of digital chest radiograph segmentation. <i>IEEE Trans. Med. Imaging</i> 14:537-547, 1995.	
15.	<b>Xiaoming Zhu, K.-N. Lee, D.L. Levin, S.T.C. Wong, H.K. Huang, K. Soo Hoo</b> .....	114
	Temporal image database design for outcome analysis of lung nodule. <i>Comp. Med Imag. &amp; Graphics</i> . accepted, 1996.	
16.	<b>Stephen T.C. Wong, H.K. Huang</b> .....	133
	A hospital integrated framework for multimodality image base management. <i>IEEE Trans. Sys., Man., and Cybernetics</i> . In press July, 1996.	
17.	<b>Stephen T.C. Wong, H.K. Huang</b> .....	154
	Design methods and architecture issues of integrated medical image database systems. <i>Comp. Med. Imaging &amp; Graphics</i> , In press, 1996.	
18.	<b>Stephen T.C. Wong, R.C. Knowlton, R.A. Hawkins, K.D. Laxer</b> .....	179
	Multimodal image fusion for noninvasive epilepsy surgery planning. <i>IEEE Comp. Graphics and Applics</i> . January, 1996.	
19.	<b>Stephen T.C. Wong</b> .....	188
	Coping with conflict in cooperative knowledge based systems. <i>IEEE Trans. on Sys., Man., and Cybernetics</i> , to appear, 1996.	

20. **Stephen T.C. Wong, S.Tojo**..... 217  
 A deductive object-oriented database system for situated inference in law.  
*IEEE Trans. Knowledge and Data Eng.*, In press, April, 1996.

21. **H.K. Huang, R.L. Arenson, W.P. Dillon, S.L. Lou, T. Bazzill, A.W.K. Wong**..... 229  
 Asynchronous transfer mode technology for radiologic image communication.  
*AJR*: 164:1533-1536, 1995.

**RSNA ABSTRACTS AND PRESENTATIONS (RADIOLOGY VOL. 197(P),1995) 233**

**Steven TC Wong and H.K. Huang**  
 Special Invited 1996 Issue of *Computerized Medical Imaging & Graphics*  
 on Medical Image database -- Table of Contents and editorial..... 240

**H.K. Huang**  
 PACS in Biomedical Imaging , 1996, VCH Publishers, Inc. -- Table of contents..... 244

# UCSF Radiological Informatics Research

## - A Progress Report -

## Laboratory for Radiological Informatics

Department of Radiology

University of California, San Francisco

### Faculty

**Ronald Arenson, M.D.**  
Professor and Chairman  
Department of Radiology

**H.K. Huang, D.Sc.**  
Professor and Vice Chairman, Dir LRI  
Department of Radiology

**David Avrin, M.D., Ph.D.**  
Clinical Associate Professor

**Katherine Andriole, Ph.D.**  
Assistant Professor

**Robert Gould, Ph.D.**  
Professor

**Shyh-Liang (Andrew) Lou, Ph.D.**  
Assistant Professor

**Stephen Wong, Ph.D.**  
Assistant Professor

### Postdoctoral Fellows

**Fei Cao, Ph.D.**  
Postdoctoral Fellow

**Mohammed Jahangiri, M.D, MPH**  
Postdoctoral Fellow

**Xiaoming Zhu, Ph.D.**  
Postdoctoral Fellow

### Graduate Students

**Kent Soo Hoo**  
Graduate Student

**Jun Wang, M.S.**  
Ph.D. Candidate

### Staff

**Todd Bazzill, B.S.E.E.**  
PACS Engineer

**David Hoogstrate, B.A.**  
Programmer/Analyst

**Christopher Jovais**  
RIS Manager

**Joseph Lee, M.S.**  
Programmer/Analyst

**Michael Linnolt, B.S.**  
Programmer/Analyst

**Laura Snarr, B.S.**  
Administrative Assistant

**Albert Wong, B.S.**  
Software System Manager

**Lloyd Yin, B.S.**  
Programmer/Analyst

### Visiting Faculty and Fellows

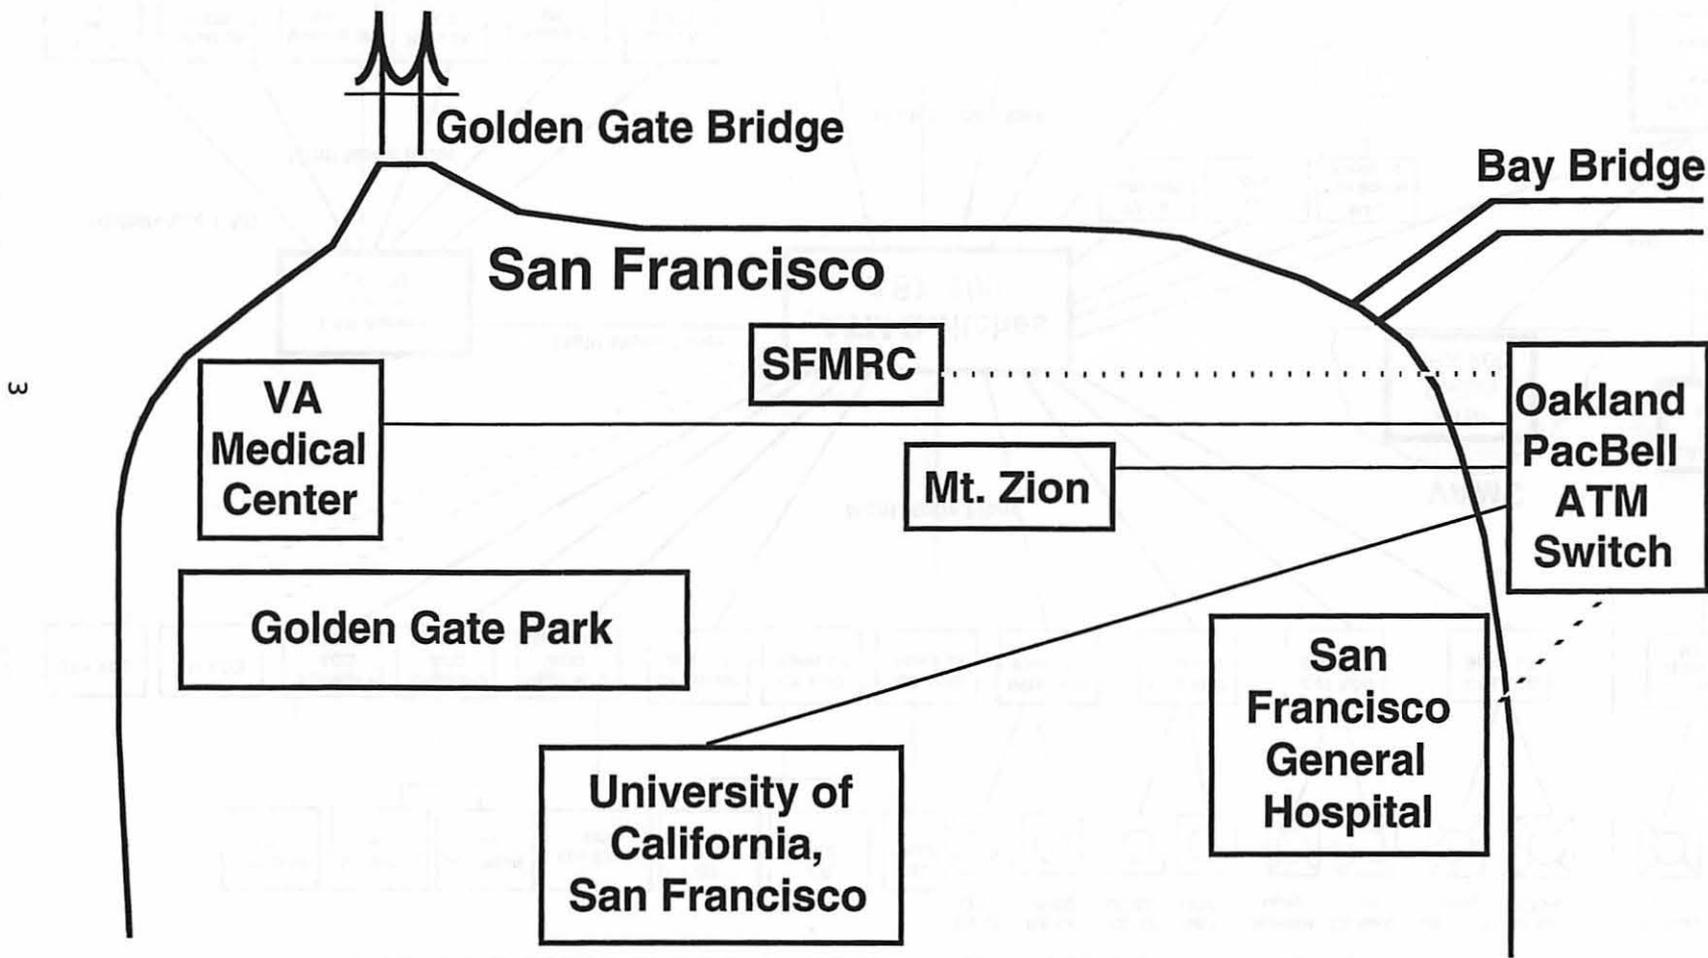
**Ewa Pietka, Ph.D.**  
Visiting Assistant Professor (Poland)

### Collaborators

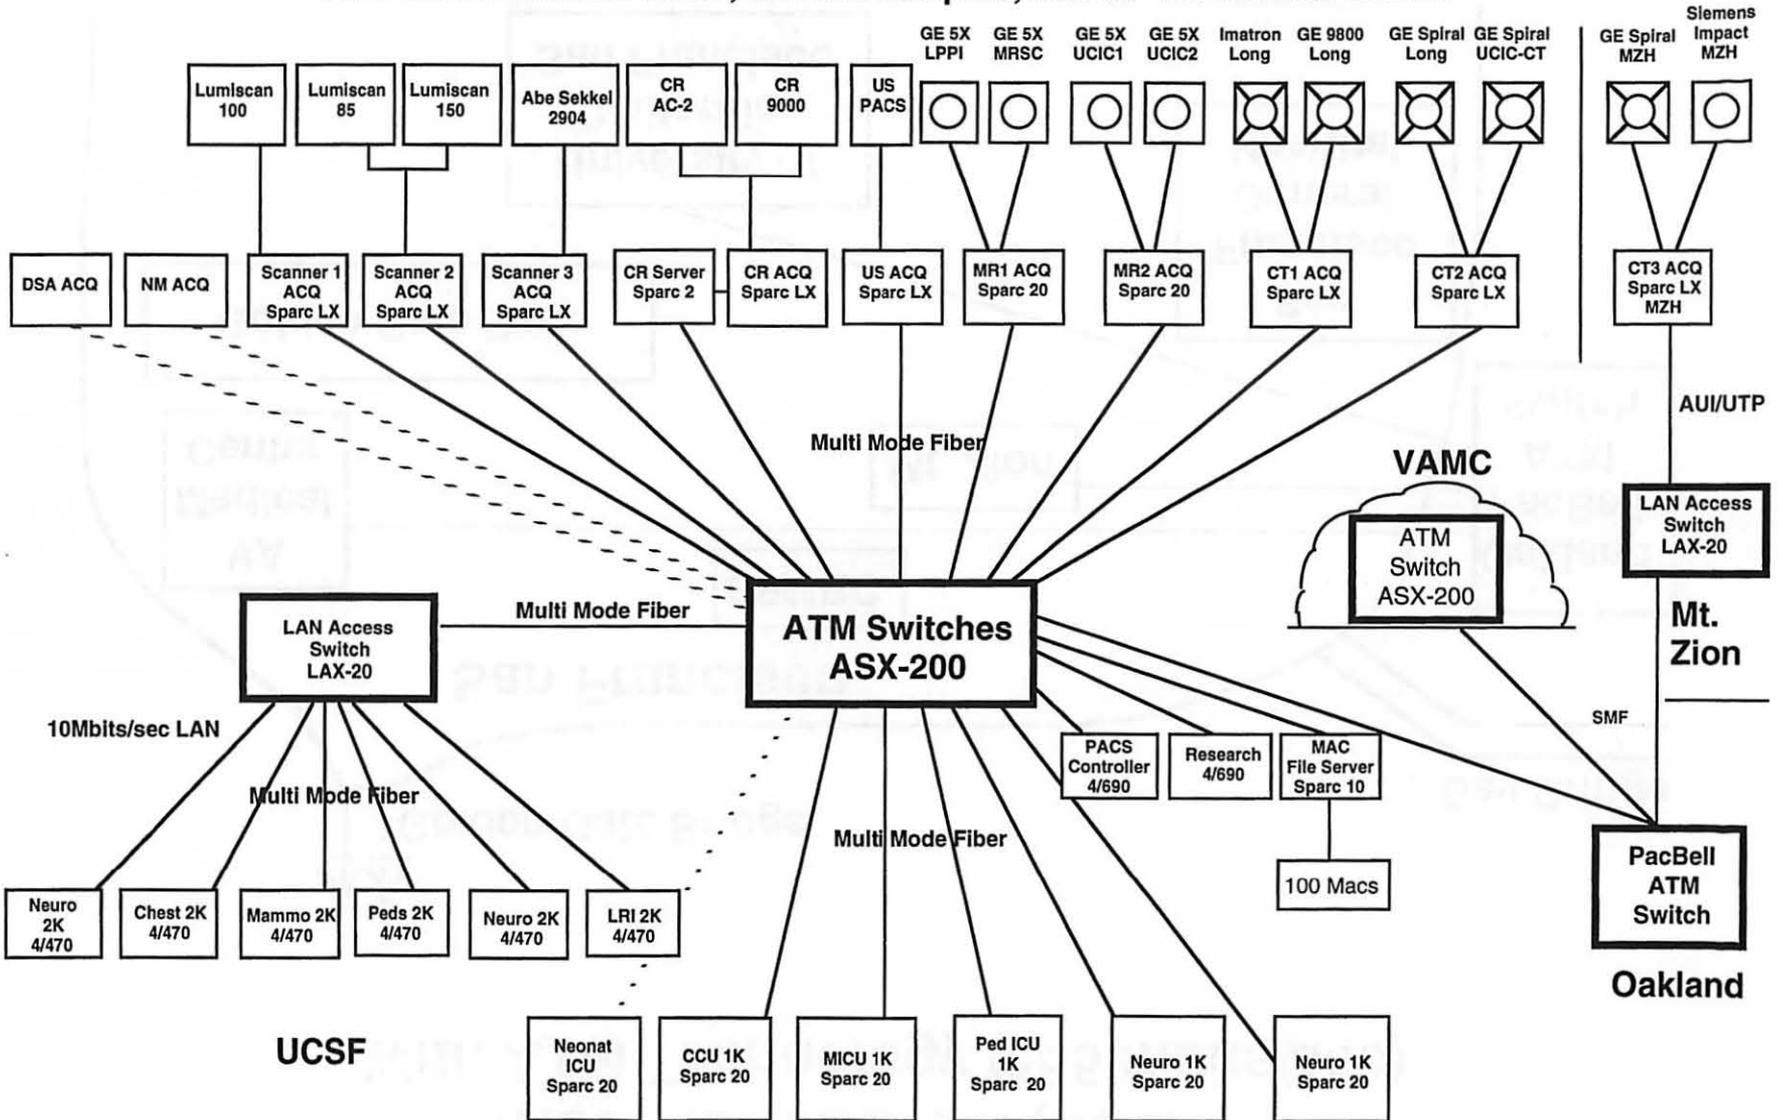
**Robert Knowlton, M.D.**  
Neurology, UCSF

**Jianguo Zhang, Ph.D.**  
Fellow (Assoc. Prof. Chang Chun  
Inst., Chinese Acad. Sci)

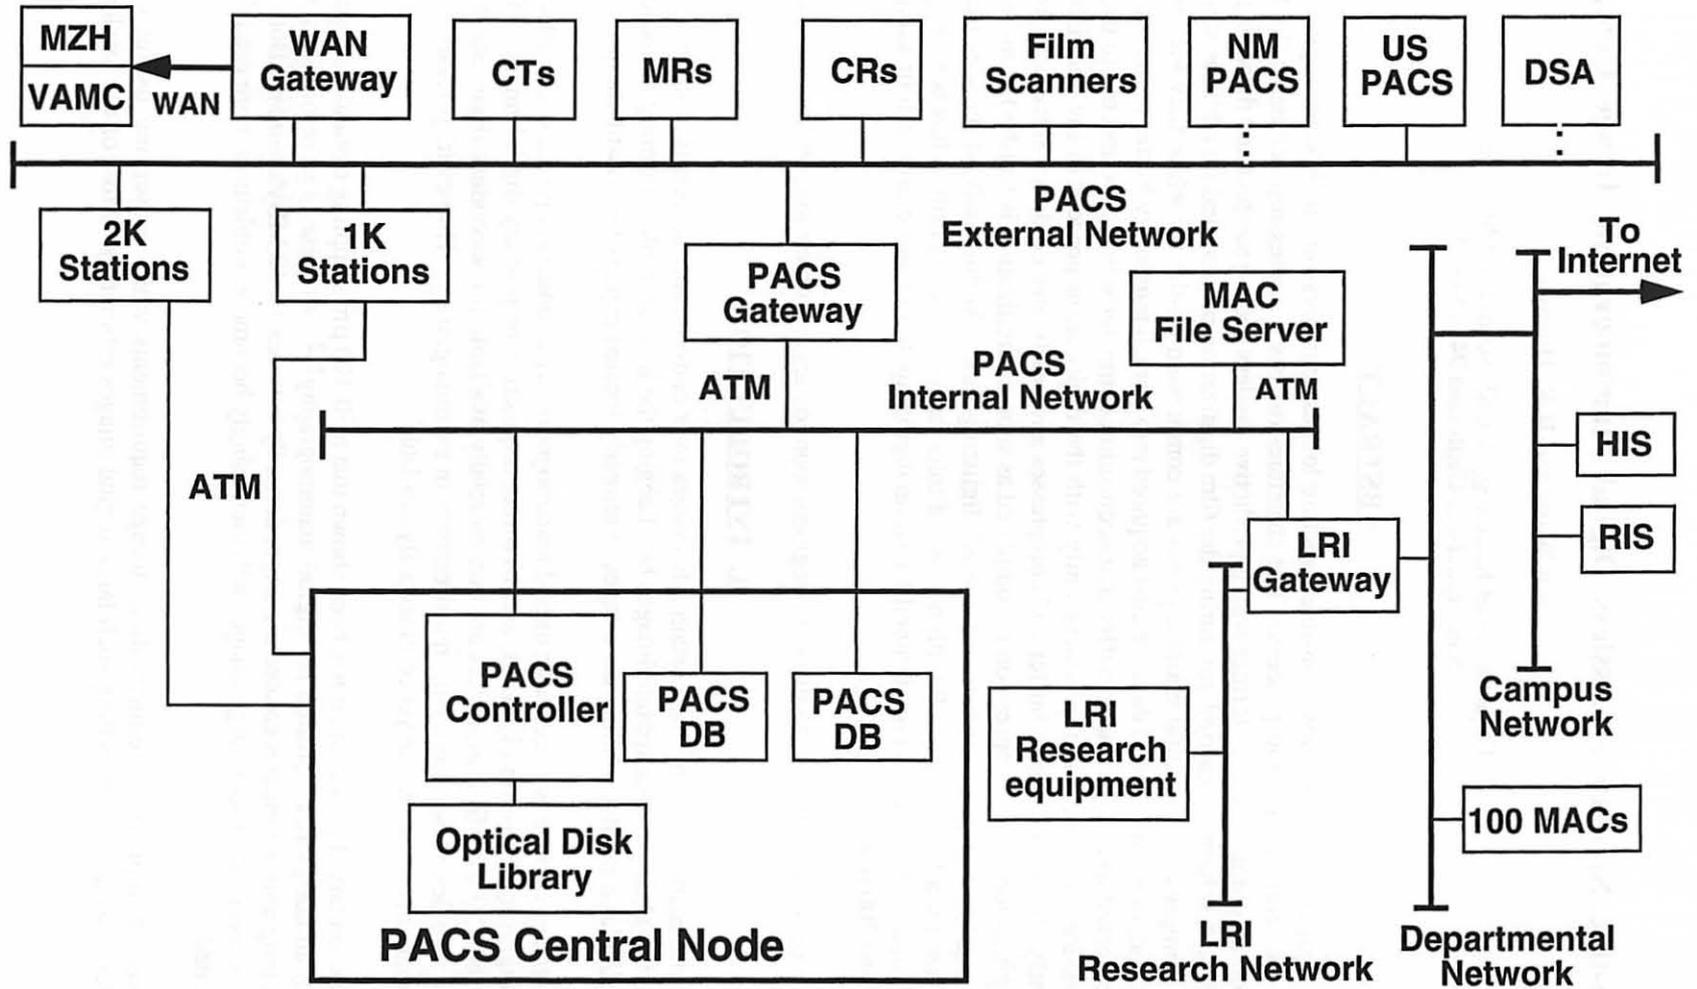
# Wide Area Network (CalREN) With ATM Technology (155 mbits/sec)



# ATM Connection at UCSF, Mt Zion Hospital, and SF VA Medical Center



# Dept. of Radiology, UCSF Network Architecture



# On-line Structure Lossless Digital Mammographic Image Compression

Jun Wang and H.K. Huang

Department of Radiology, UCSF School of Medicine  
San Francisco, California 94143-0628

## ABSTRACT

This paper proposes a novel on-line structure lossless compression method for digital mammograms during the film digitization process. The structure-lossless compression segments the breast and the background, compresses the former with a predictive lossless coding method and discards the latter. This compression scheme is carried out during the film digitization process and no additional time is required for the compression. Digital mammograms are compressed on-the-fly while they are created. During digitization, lines of scanned data are first acquired into a small temporary buffer in the scanner, then they are transferred to a large image buffer in a acquisition computer witch is connected to the scanner. The compression process, running concurrently with the digitization process in the acquisition computer, constantly checks the image buffer and compresses any newly arrived data. Since compression is faster than digitization, data compression is completed as soon as digitization is finished. On-line compression during digitization does not increase overall digitizing time. Additionally, it reduces the mammogram image size by a factor of 3 to 9 with no loss of information. This algorithm has been implemented in a film digitizer. Statistics were obtained based on digitizing 46 mammograms at four sampling distances from 50 to 200 microns.

*Key words:* mammograms, digitization, segmentation lossless compression, on-line, laser film scanners.

## 1. INTRODUCTION

Digital mammography has certain advantages over conventional screen-film mammography. These include the ability to manipulate images by changing the look-up-table, zooming and scrolling, imaging processing, and rapidly sending the images to remotely located experts for consultation.

There are two ways to generate digital mammograms: a) secondary digitization, in which conventional mammography films are digitized; and b) direct acquisition of primary digital images. Currently, high resolution (50  $\mu\text{m}$ ) film digitizers are commercially available for secondary digitization and the image quality satisfies most diagnostic requirements in mammography. However, primary full-field digital mammography units are not yet commercially available.

In secondary digitization, it has been shown that a 50-100  $\mu\text{m}$  sampling distance is required in order to achieve an adequate resolution for digital mammography<sup>1,2</sup>. At these high resolutions, the digitization time is long and the images created are large, normally between 10-40 MBytes/image. Both the storage and transmission cost of such large images will increasingly become a problem as more digital mammograms are created.

Image compression reduces data storage requirements while preserving useful information. It provides a solution to handling such large digital images efficiently. Most of the lossless compression

methods can achieve a compression ratio of around 2:1 to 3:1 for mammograms, and will allow exact reconstruction of the original image<sup>3,4</sup>. Conventionally, digitization and compression of mammograms are done separately. Compression is applied to a full digital image after it has been acquired through digitization.

This paper proposes a new structure lossless compression scheme and combines it with an on-line digitization process to increase the overall throughput of the digitization and compression operations. The structure lossless compression method segments the breast image from its background and compresses only the breast portion with a lossless predictive coding method. This method achieves compression ratios of 3:1 to 9:1. The compression process is implemented in the acquisition computer connected to the scanner, running concurrently with the digitization process. Overall time to digitize and compress an image is simply the time of digitizing.

## 2. STRUCTURE LOSSLESS COMPRESSION

Mammograms have unique characteristics compared with other types of medical images. A breast image has a regular round shape, and the background area is relatively large in a rectangular film. The background contains no anatomy information. Therefore, any changes in the background as a result of image compression and decompression will not affect the breast image. Although labels often appear in the background area, they can be stored as text in the header. They don't need be compressed and stored as part of a image.

Structure lossless compression consists of two steps. The first step segments the breast image from its background. The pixels beyond the boundary of the breast are discarded. The second step compresses the remaining portion of the image using a predictive lossless compression technique. Both segmentation and lossless compression can be applied to one line of the image data at a time.

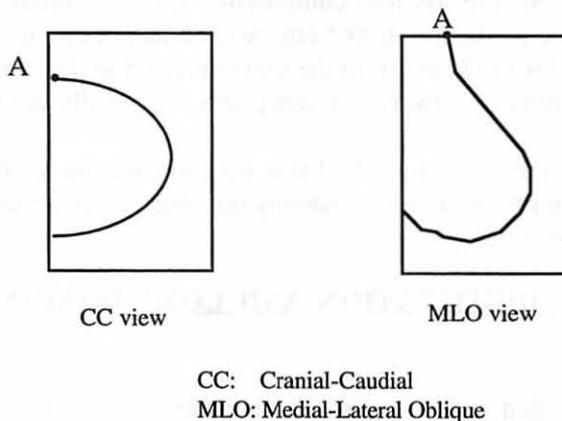


Figure 1. The boundary of mammograms

The segmentation process scans one line of the image at a time to detect the boundary of the breast image. The first step of the segmentation is to detect the top boundary point A. (Figure 1). Each line of data is scanned from top to bottom until the top boundary is detected. On every line of data, ten evenly spaced check points are selected and the gray level of each point is compared with the others to determine

if there are pixels from the breast image in this line. If so, the boundary of the breast is determined by thresholding a smoothed data which is obtained by averaging five neighboring pixels on one line of data.

The second step of segmentation completes detecting the boundary for the rest of the image. Once the top boundary point A has been detected, the boundary in the next line can be located by searching the neighboring pixel of the current boundary point. The data are first smoothed in a region which is determined by the location of the previous boundary point. Again, the exact boundary point is determined using a threshold on the smoothed data. Since breast images have a smooth boundary, the boundary points are only detected every few lines to save processing time (Figure 2(a)). The boundary points on the rest of the lines are interpolated by the two nearest detected boundary points (Figure 2(b)). The background outside the boundary is discarded.

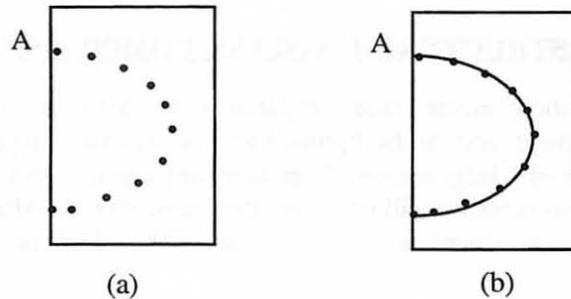


Figure 2. Boundary detection. (a) Boundary points detected using thresholding.  
(b) Connected boundary points using interpolation.

The next step of the structure-lossless compression uses the predictive coding to compress the segmented breast image. The prediction starts from the boundary of the breast and moves to the chest wall direction. The current pixel is predicted from the previous pixel in the same line. Since the neighboring pixels are correlated, the difference between adjacent pixels is generally small.

Huffman coding<sup>5</sup> with a pre-determined table is used to code the resulting differences of the pixels. The Huffman table is determined by a set of randomly selected mammograms. If an individual pixel is not in the table, it is recorded separately.

### **3. ON-LINE DIGITIZATION AND COMPRESSION PROCESSES**

#### **3.1 Digitization process**

A film digitizer is controlled via an acquisition computer. The digitizer scans (or digitizes) a film one line at a time from the left to the right, advancing line by line from the top to the bottom of the film. The data generated is first stored in an internal buffer in the digitizer. When the buffer is full, the data is transferred to the acquisition computer which has a larger size memory buffer that will hold the entire image. This scanning and transferring process continues several lines at a time, until the whole film is scanned.

During the conventional digitization process the sole responsibility of the acquisition computer is to read the data from the digitizer's memory buffer into its own memory buffer. So the acquisition computer CPU is mostly idle when no data is being transferred. The time to digitize a film is determined by the digitizer scanning speed and the film size. The scanning speed is normally 50-100 lines/second, depending on the sampling distances. To digitize an 8"x 10" mammogram, it takes about 30 seconds at a resolution of 100  $\mu\text{m}/\text{pixel}$  and 100 seconds at 50 $\mu\text{m}/\text{pixel}$ .

### 3.2 On-line compression process

We utilize the computer's idle CPU time to compress the digitized image during digitization. The compression is implemented as a second process running concurrently with the digitization process in the acquisition computer (Figure 3). Both processes share the same image buffer. A semaphore is used to control the access of the shared image memory by the two processes.

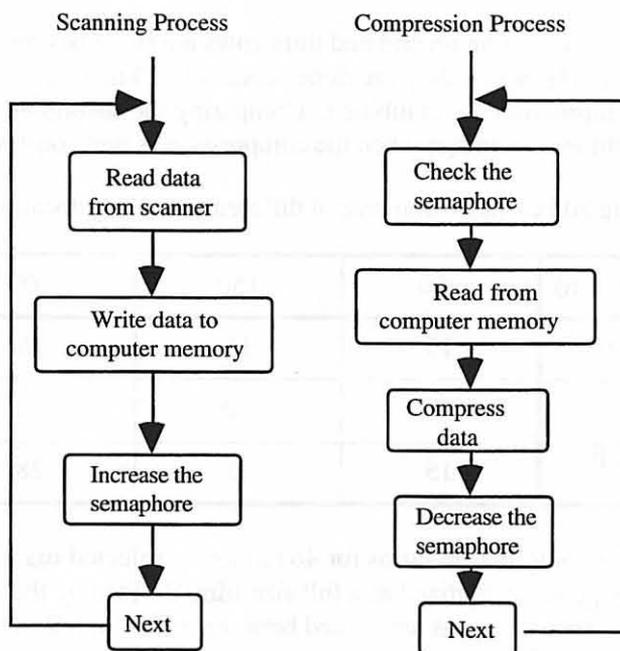


Figure 3. Data flow of the concurrent scanning and compression processes

First, scanning process reads several lines of data from scanner and writes into the acquisition memory buffer whenever there is data ready in the scanner's temporary buffer. After writing data each time, the scanning process increases the semaphore by one.

Meanwhile, the compression process checks the semaphore. If the semaphore indicates new data in the acquisition computer memory, the compression process reads these several lines of data from the image buffer, finds the boundary of the breast and compresses this breast image portion with predictive coding method as described in the previous section. The compression process then decreases the semaphore by one. If there is no new data in the memory, the semaphore blocks the memory and stops

the compression process until new data arrives in the memory. The processes continues until an entire film is scanned and the image is compressed.

If the compression is faster than the scanning, compression can be done as soon as the scanning finishes. The scanning rate is normally 40 to 100 lines/sec. We have implemented this on-line compression algorithm in a film digitizer. In this prototype system, the compression rate achieved is faster than the scanning rate.

#### 4. RESULTS

We first compared the times required to digitize and to compress an image in two situations: (1) digitization and compression processes were done separately, (2) digitization and compression were combined. A mammogram was digitized at different sampling distances ranging from 50  $\mu\text{m}$  to 200  $\mu\text{m}$  with a 50  $\mu\text{m}$  increment. The times required for each process were measured.

Table 1 tabulates the results. The second and third rows are the times required to digitize films and to compress images, respectively, when they are done separately. The fourth row shows the time required when digitization and compression are combined. Comparing the second and the fourth row, there is no extra time required to compress an image when the compression is done on-line.

Table 1. The scanning and compression time at different sampling distances

Sampling Distance ( $\mu\text{m}$ )	200	150	100	50
Scanning (sec)	15	15	28	103
Compression (sec)	7	10	22	102
Scanning with On-line Compression (sec)	15	15	28	103

We also compared the compression ratios for 46 randomly selected mammograms. The compression ratio is defined as full image size (digitized as a full size film) divided by the compressed image size. The compression ratios for 46 mammograms are ranged between 3.2 : 1 to 8.9 : 1 with an average ratio of 5.65 : 1 and a standard deviation of 1.46.

The large variation of the compression ratio is due to the variation of the area occupied by the breast image. A larger breast occupies more film area which reduces the background region, so the portion that can be discarded is less and therefore the compression ratio is smaller. Figure 4 shows an example with a larger breast that occupies most of the film. The compression ratio achieved is about 3.1 : 1. Figure 5 shows an example of a small size breast that occupies small film area, and the compression ratio achieved is about 6.3 : 1.

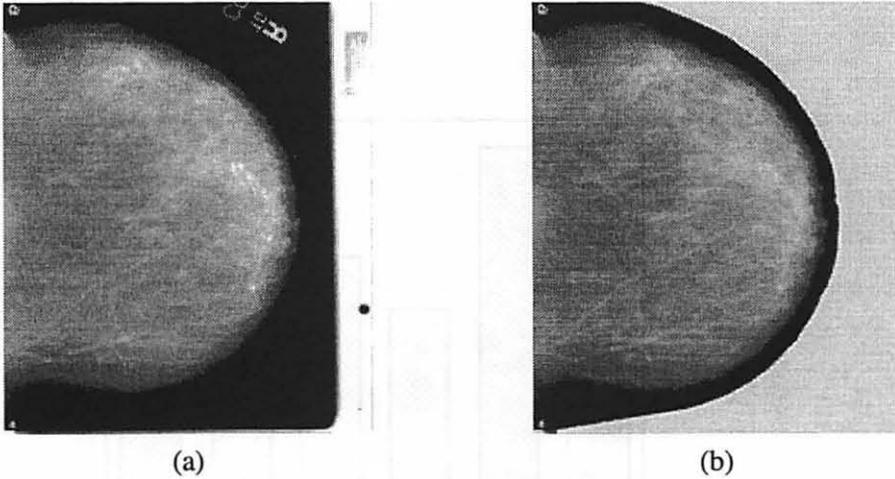


Figure 4. (a) A mammogram with a large breast that occupies most of the film.  
 (b) Segmented image. The compression ratio is 3.1 : 1.

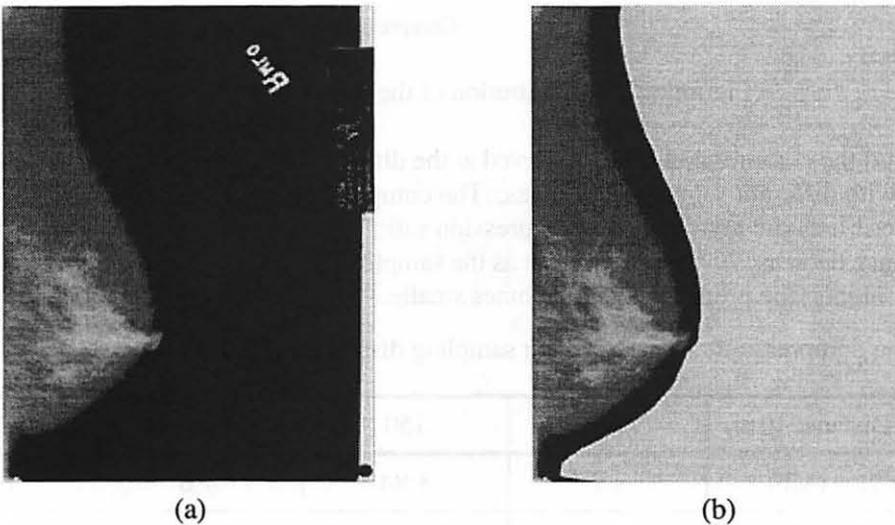


Figure 5. (a) A mammogram with a small breast that occupies small portion of the film.  
 (b) Segmented image. The compression ratio is 6.3 : 1.

The distribution of compression ratios for 46 mammograms is shown in Figure 6. About 70% of the images have compression ratios between 4 : 1 and 7 : 1. About 20 % of the images have compression ratios larger than 7 : 1. Only 10% of the images have compression ratios between 3 : 1 and 4 : 1.

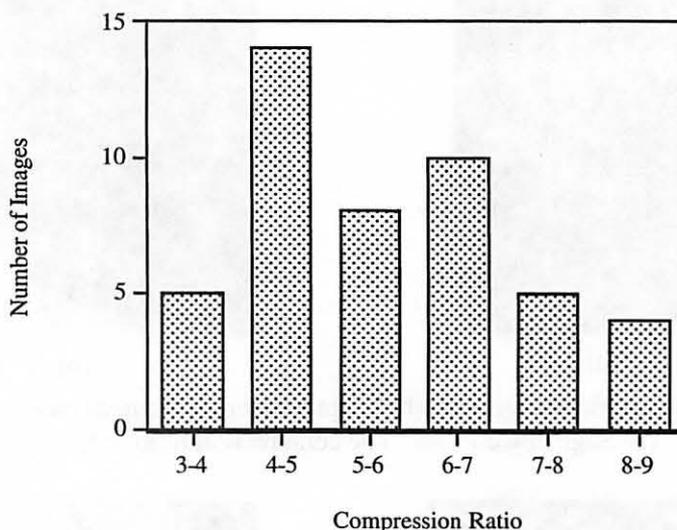


Figure 6. The distribution of the compression ratios

We compared the compression ratios achieved at the different sampling distances. Five mammograms were digitized with different sampling distances. The compression ratios were averaged at each sampling distance. Table 2 lists the results. The compression ratio increases slightly as the resolution increases (sampling distance decreases). This is because as the sampling distance decreases, the correlation between pixels becomes higher, the predicted error becomes smaller, and the compression ratio is higher.

Table 2. The compression ratio at different sampling distances

Sampling Distance ( $\mu\text{m}$ )	200	150	100	50
Image Size (Mb)	2.2	3.84	8.6	34.6
Compression Ratio	7.12	7.24	7.39	7.83

## 5. CONCLUSIONS

We developed a structure lossless compression method for digital mammograms based on segmentation and predictive coding. The compression ratios achieved are between 3 : 1 to 8 : 1 with no information loss, which are much higher than conventional lossless compression.

We have also implemented the compression on a film digitizer which compresses the data on the fly while it is acquired. Digitizing mammograms with on-line compression takes less time compared with scanning and compression processes done separately. In fact, on-line compression during digitization does not increase overall scanning time.

## **6. ACKNOWLEDGMENT**

This work was supported in part by U.S. Army Medical Research Acquisition Activity grant number DAMD17-94-J-4196.

## **7. REFERENCES**

1. Karssemeijer N, Frieling JTM, Hendriks JHC, Spatial Resolution in Digital Mammography. Investigative Radiology, Vol. 28(5), pp. 413-419.
2. Nawano S, Evaluation of Digital Mammography in Diagnosis of Breast Cancer. Journal of Digital Imaging, Vol. 8 (1) Suppl 1, pp. 67-69, 1995.
3. Maeder AJ, Mammogram Compression Using Adaptive Prediction. Proceedings of SPIE Medical Imaging, Vol. 2431, pp216-223, 1995.
4. Saipetch P, Panwar RK, Ho BKT, Ma M, Lossless Image Compression Using Wavelet-like Filter Banks and Arithmetic Coding. Proceedings of SPIE Medical Imaging, Vol. 2431, pp501-507, 1995.
5. Huffman DA, A Method for the Construction of Minimum-redundancy Codes. Proceedings of the IRE, Vol. 40(10) pp. 108-1101, 1952.
6. Wang J and Huang H.K. Real-Time Structure Lossless Mammographic Image Compression. Radiology Suppl. Vol. 197(p), pp393 , 1995 (Abstract).

# Interactive Query and Visualization of Medical Images on the World Wide Web

Stephen TC Wong<sup>1,2</sup>, Pablo Whaley<sup>4</sup>, CS Ang<sup>3</sup>, Kent Soo Hoo<sup>2</sup>, Jun Wang,<sup>1</sup> HK Huang<sup>1,2</sup>

<sup>1</sup>Department of Radiology, UCSF, San Francisco, CA 94143-0628

<sup>2</sup>Bioengineering Graduate Program, UCSF/UC Berkeley, San Francisco, CA 94143-0640

<sup>3</sup>Library and Center of Knowledge Management, San Francisco, CA 94143-0840

<sup>4</sup>Department of Mathematics, Univ. of California, Berkeley, CA 94530

## ABSTRACT

The World Wide Web (WWW) is becoming the predominate force for global information dissemination. The Web browsers are originally designed for the client computers to navigate and display hypermedia documents and image bitmaps stored at the server machines. Their capabilities must be extended with database query and interactive visualization before the Web could be useful for medical imaging community. This paper presents the system design and tools for interactive query and visualization of medical images on the Web. Examples from breast and brain imaging applications will be used to illustrate the operations and capabilities of such tools.

## 1. INTRODUCTION

Current methods for medical image communication and visualization are based on dedicated computer networks and expensive display stations. These workstations often require specialized graphics hardware for rendering, large local disk for data storage, and customized graphical user interface (GUI) for image retrieval and display. The expense of these requirements has limited the ability of the researchers and scientists to exchange findings as well as the clinicians and educators to share teaching cases. To overcome the barrier of cost and the usability of client software in various platforms, we investigated software tools that can allow multiple Web clients to access image databases at the server and to visualize multidimensional medical images on general purpose computers with the inexpensive Web browsers. Currently no such tool exists.

The hypothesis of our research is that interactive query and visualization software for medical images can be designed and implemented using the current technology of hypermedia computing, graphics, and databases, and can be accessed freely via the World Wide Web (WWW). Before the WWW, the Internet was one big distributed text only information system, where queries were carried out from the command prompt, and requires considerable effort by the users to find and access appropriate resources. Now, with hypertext and browsers such as Mosaic and Netscape, the WWW provides the user with a nearly unlimited informational system consisting of not only text but full color image bitmaps, video, and sound.

Our system prototype has been tested with two medical imaging applications. They include a computer aided instruction (CAI) program of breast imaging with 2-D mammographs and related modalities and the software tools for visualizing 3-D brain images from an epilepsy imaging database. In the medical imaging domain, the ability to query on-line image databases and to render 2-D and 3-D medical images could prove to be invaluable for biomedical research, education, and clinical practice. For example, existing methods of medical education, such as pamphlets, brochures and various on-line services, are passive and textual providing little or no opportunity for visual or interactive learning. CAI of multimedia medical data on the Web will provide a more stimulating and cost-effective means of learning.

The remainder of this paper is organized as follows: Section 2 provides an overview of the World Wide Web and states its limitations for database query and interactive visualization. In Section 3, we present the distributed client/server architecture of our prototype system. Section 4 then describes the CAI application for breast imaging while Section 5 presents the interactive volume visualization techniques for 3-D neuroimages.

## 2. THE WORLD WIDE WEB

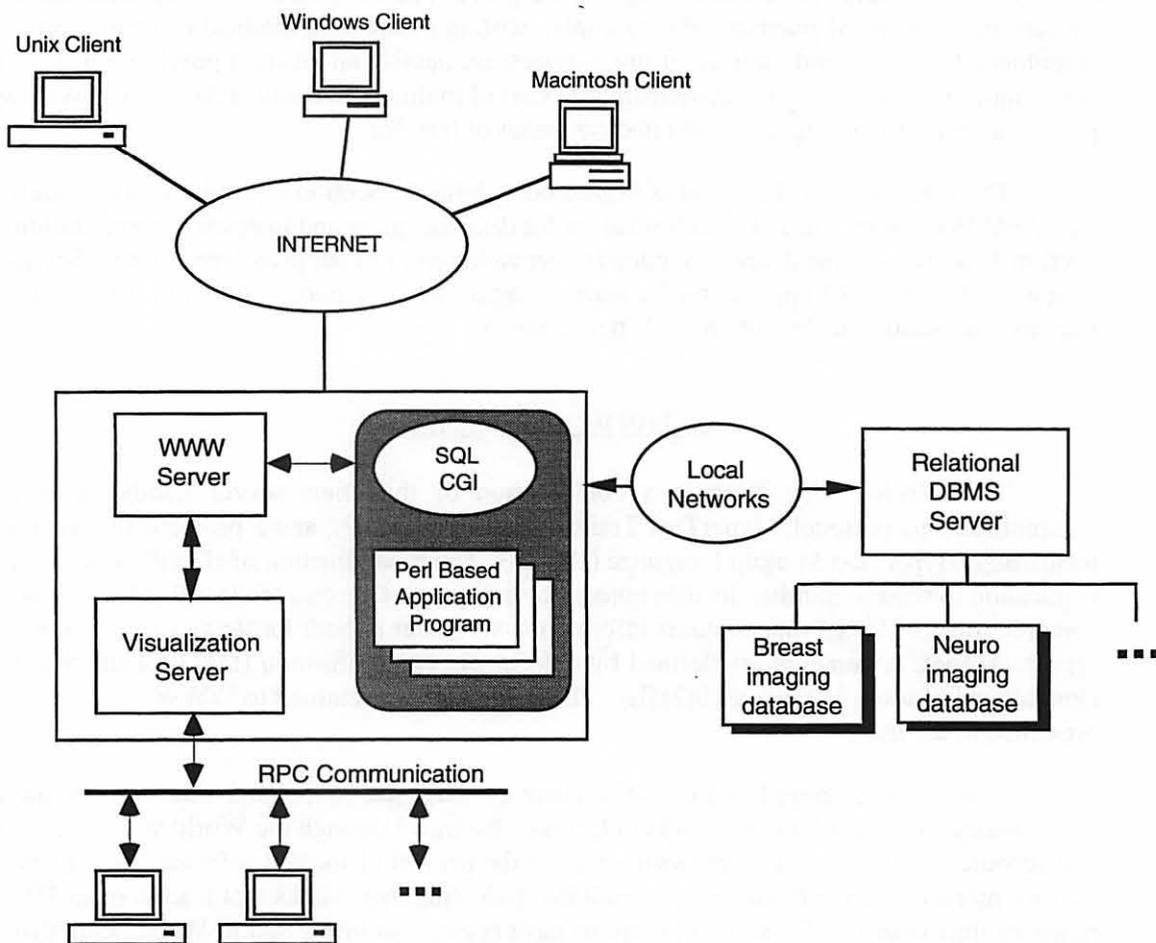
The World Wide Web is a combination of the client-server handshaking and communications protocol, HyperText Transfer Protocol (HTTP), and a protocol for document formatting, HyperText Markup Language (HTML). The basic function of HTTP allows a client application to request multimedia data objects from a server. Objects are identified by a universal resource locator (URL) that contains information sufficient to both locate and access a remote server. HTML documents are defined by a document type definition (DTD) of the Standard Generalized Markup Language (SGML). The documents are returned to WWW clients and are presented to the user.

The two basic ingredients of a Web client are a navigation tool and a document browser. The primary function of a navigator is to facilitate the travel through the World Wide Web from one resource to another. The browser supports the perusal of located information, allows the user to interact with the document presentation, following hyper-links that lead to other HTML pages or data objects. These two functions have become so integrated in Web clients that all navigator/browser are now referred to by the single term, browser. However, it is instructive to remember that these two functions serve distinct operational needs. The client browsers may also integrate other Internet services, such as File Transfer Protocol (FTP), Gopher, and WAIS.<sup>1</sup>

The utility of the Web is well demonstrated by recent surveys.<sup>2,3</sup> The Web became the dominant Internet resource as measured by both Packet count and volume, surpassing FTP, Telnet, and Gopher. By April 1995, the number of packets moved by the Web is 21.4%, well ahead of both FTP (14%) and Telnet (7.5%). Further, a wide range of robust Web clients are now available for Windows and Macintosh computers. This enables the use of economic, desktop computers for multimedia information access and presentation.

As Web users become more adroit, their expectations continue to increase and they will demand a wide range of functionality from their clients. This situation applies to medical imaging community where the access to large image databases, such as picture archiving and

communication systems (PACS), and the ability to visualize and render 3-D images would have the potential to improve the communication and sharing of information among the researchers, scientists, and clinicians. To achieve this, new software capabilities must be added into the Web servers.



A network of heterogeneous Unix workstations

Figure 1. The distributed client/server architecture for medical imaging on the Web.

### 3. SYSTEM ARCHITECTURE

The interactive query and visualization software is based on the distributed client/server architecture shown in Figure 1. The web client is any Windows PC, Unix, or Macintosh desktop in the Internet running a Web browser, e.g. Mosaic and Netscape. There are three servers in this architecture: the Web server, the database server, and the visualization server. All servers are running on Unix computers.

The current database server is a relational database management system (DBMS) organizing and storing multiple imaging databases. The DBMS used in our implementation is Sybase (Sybase Inc., Emeryville, CA).<sup>4</sup> A relational database is simply a series of tables made up of fields (columns) and tuples (rows) where image and textual data are stored and categorized. The query to the DBMS is based on the standard Structured Query Language (SQL). The DBMS server resides in a different physical computer from the Web server and has large disk space for storing various image databases. The image database server in the Laboratory for Radiological Informatics manages four databases of different specialties, namely, breast imaging (mammograms), epilepsy imaging (magnetic resonance imaging and positron emission tomography), thoracic imaging (Spiral computed tomography), and hand imaging (computed radiography and digitized x-rays). It resides in a 4 processor SUN SPARC 690 with 20 gigabytes of disk space, connecting to the Web server in the same local area network. Additional database servers for less popular object-oriented and object-relational databases can also be added if necessary.

For the Web server to retrieve image data from the relational database server, a CGI (common gateway interface) script was written using Sybperl.<sup>5</sup> Sybperl is an SQL API (application programming language) that maps a subset of Sybase SQL library commands to the interpretation language Perl (Practical Extraction Report Language).<sup>6</sup> This enables the development of Perl-based application software to access and retrieve specific datasets from the underlying imaging database. For communicating with other kinds of DBMS, appropriate CGIs must be added.

The visualization server allows the manipulation and visualization of medical images. It has the following routines. PolyMap<sup>7</sup> enhances the users ability to "interact" with the software by making predefined regions of interest on the image "clickable". For example, a region of an image that is predefined using the PolyMap will become outlined in blue when the user passes over it with the mouse. This region can be linked to a "close-up" image, providing the user with an enhanced view of that specified region. PolyMap is similar to the Imap function provides available on the Web except that it can be used to define arbitrary boundary of the image. This tool has been used extensively in the interactive teaching of 2D mammographic images, which is discussed in Section 4. At the moment, this routine is compatible with the Mosaic browser only. Work is progressing to extend the mapping for the Netscape browser.

The other part of the visualization server consists of a distributed volume visualization tool, VIS, that can be used render and manipulate 3D volume data retrieved from the image database server. Volume rendering algorithms require a significant amount of computational resources and time. To speed up the process, VIS distributes the volume rendering among workstations via remote procedure calls (RPC). This part of the project complements the effort of the Center for Knowledge Management at UCSF in bringing volume visualization on the Web.<sup>8</sup> To reduce communication overhead, the visualization server and the Web server reside with the same Unix computer (SUN SPARC LX, SUN Microsystems, Mountain View, CA). Section 5 describes the functions of VIS and illustrates the use of this tool to visualize 3D MR head images.

In addition, the Web server contains software programs written in Perl that integrates all of the components above for particular applications. Perl programming language was chosen over others, such as C, because it is an interpreter language, providing the user with built-in functions for easy manipulation of processes, files, and networking tasks, and is available in the public domain. The next section describes a computer aided instruction program written in Perl for learning breast imaging cases on the Web.

#### 4. IMAGING DATABASE QUERY FOR INTERACTIVE LEARNING

The long term objective of this CAI program is to provide an interactive learning environment for breast imaging teaching file on the World Wide Web. Figure 2 shows the generic algorithm used in the CAI program for breast cancer education. The implementation of this algorithm separates program control from medical data. The program control or the instruction sequences is defined by a Perl program with four reusable subroutines and embedded HTML statements. The medical data are organized in tables of the remote database server and can be retrieved by these routines using SQL CGI. The data includes breast images, textual description, question listings, and correct answer of each question. In addition, the user can stop the learning process at any time and the session will be logged into a local record file. The user then can start from the point that he or she left off in the last session. In this way, one can learn the teaching materials at one's own pace.

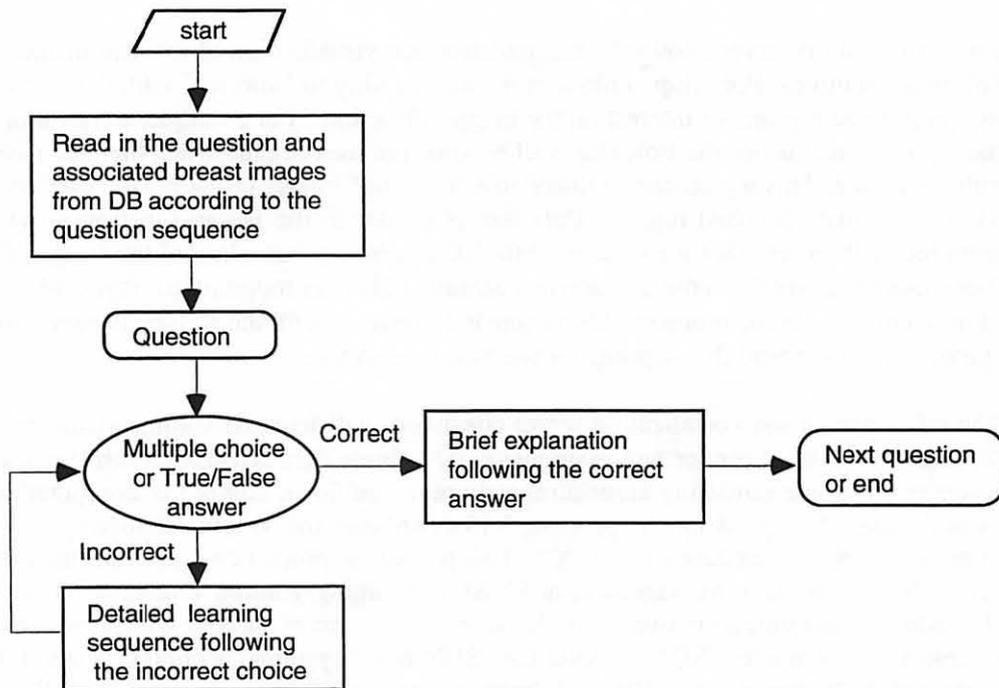


Figure 2. A general algorithm for CAI program for Breast Cancer education.

As explained in Figure 2, the Breast Cancer education program consists of a CAI algorithm. This algorithm is composed of a basic set of learning sequences which define dynamic and precise system responses uniquely determined by the user input. These learning sequences are carried out by a series of subroutines, each of which is executed by “if-else” statements which match user input with a predefined variable.

As an example, in one of the multiple choice or true/false questions, once the user chooses an answer, the program queries the database to check the correctness of that answer and prompts the user with the appropriate result and explanation. The resulting data then activates a subsequent question and answer sequence and the interactive learning process continues. This algorithm can be easily modified and expanded for other applications because of the flexibility of Perl. In total, the main program is responsible for creating a dynamic learning sequence to provide the user with an interactive learning experience. It not only provides them with answers to questions and enables them to move at their own pace, but also provides image visualization, which is driven by a similar learning sequence as illustrated in Figure 3.

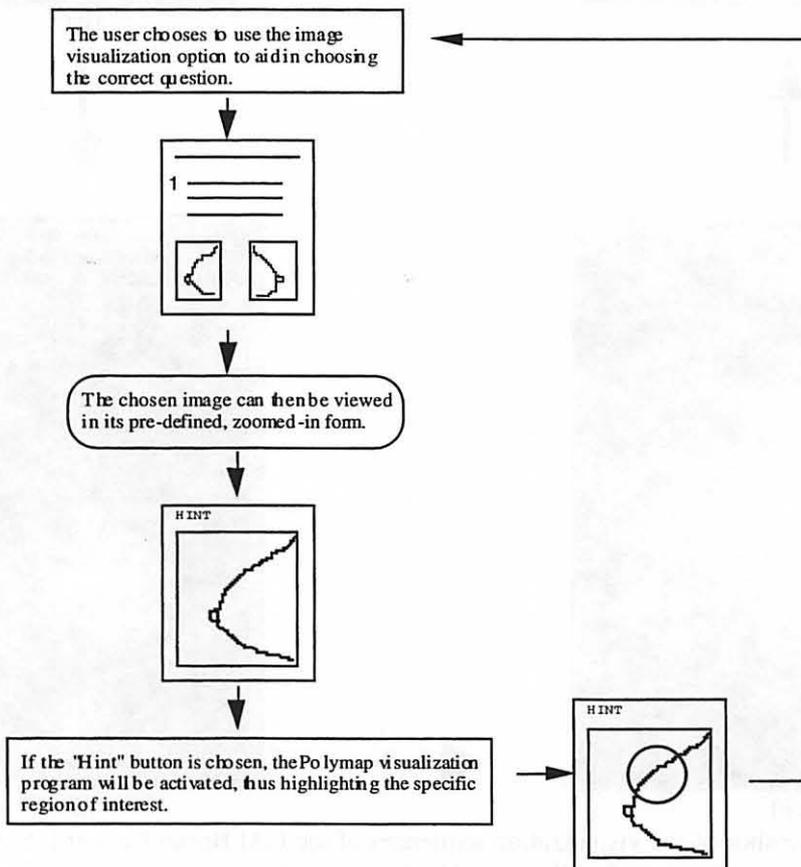
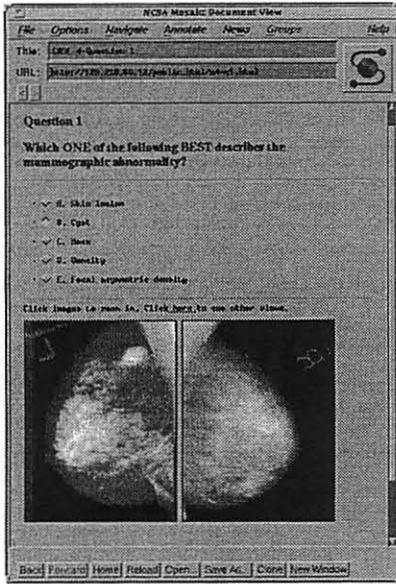
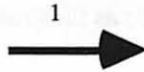


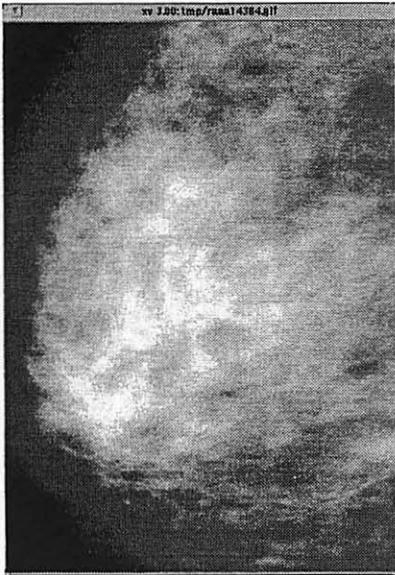
Figure 3. A general algorithm for the image query and visualization component of the CAI Breast Cancer education program.



(a)



(b)



(c)



(d)

Figure 4. A snapshot of the visualization sequences of the CAI Breast Cancer Education program on the WWW, using the algorithm illustrated in Figure 3. The breast image on (b) is digitized with 200  $\mu\text{m}$  resolution while the breast image on (d) is of 75  $\mu\text{m}$  resolution. All versions of breast images are stored in the breast imaging database (see Figure 1).

## 5. INTERACTIVE VOLUME VISUALIZATION OF MEDICAL IMAGES

We extended the capability of the visualization server to render and display 3D volume data stored in the neuroimaging database. The front end of the client program remains the same: enhanced Mosaic as the main user interface, but with VIS, a visualization application that consists of several concurrent processes: a control panel, a Panel, a visualization engine (Vis), and one or more volume rendering servers, VRservers. Enhanced Mosaic is a version of NCSA Mosaic modified to support embedding of program objects in HTML documents. Panel and Vis are two separate processes which communicate via ToolTalk, an interprocess communication mechanism that operates across networks. VRservers handle the volume rendering operation of VIS. Since VRservers are connected to Vis via RPC, the computationally intensive volume rendering process can potentially be distributed to many machines on the network.

A user of the visualization program will first enter the keys characterizing some volume datasets. The query is submitted to the HTTP server which, through Syberl CGI, queries the relational DBMS server for the volume dataset of user interest. If a matching of dataset is found, an HTML page will be returned to the enhanced browser, with the URL of the dataset embedded in the "EMBED" tag on the page. The browser then starts up VIS, the application registered for data of MIME type application/x-vis, to handle the volume dataset. VIS started in the above embedded mode will display computed images directly in a designated window on the HTML page in the browser. Data manipulation which includes functions such as rotation, scaling, axial clipping, and arbitrary clipping, is done through Panel which relies on Vis for computation. VIS interactively transforms the wire-frame representation of the volume data, and texture-maps the volume data onto the transformed geometry. It also supports distributed volume rendering with run-time selection of computational servers, and isosurface generation using the marching cubes algorithm.

The data file format currently supported is NCSA Hierarchical Data Format (HDF), but VIS provides a graphical interface utility for importing raw volume data. The front-end programs of this visualization system are available for SunOS 4.1.x, and IRIX 5.x. The enhanced NCSA Mosaic may be downloaded at the embedding technology licensees home page, <http://www.eolas.com/>. Figure 5 shows our distributed client/server architecture with the WWW visualization modules. Except for the modules in rectangular boxes, which may be thought of as machine boundaries, each of the modules can potentially be running on an individual machine on the network. Figure 6 presents an example of the visualization of an MRI head volume dataset from a Web-based client computer. The VIS application can be downloaded from <ftp://ftp.library.ucsf.edu/pub/vis>. More information regarding the visualization system can also be obtained from the authors.

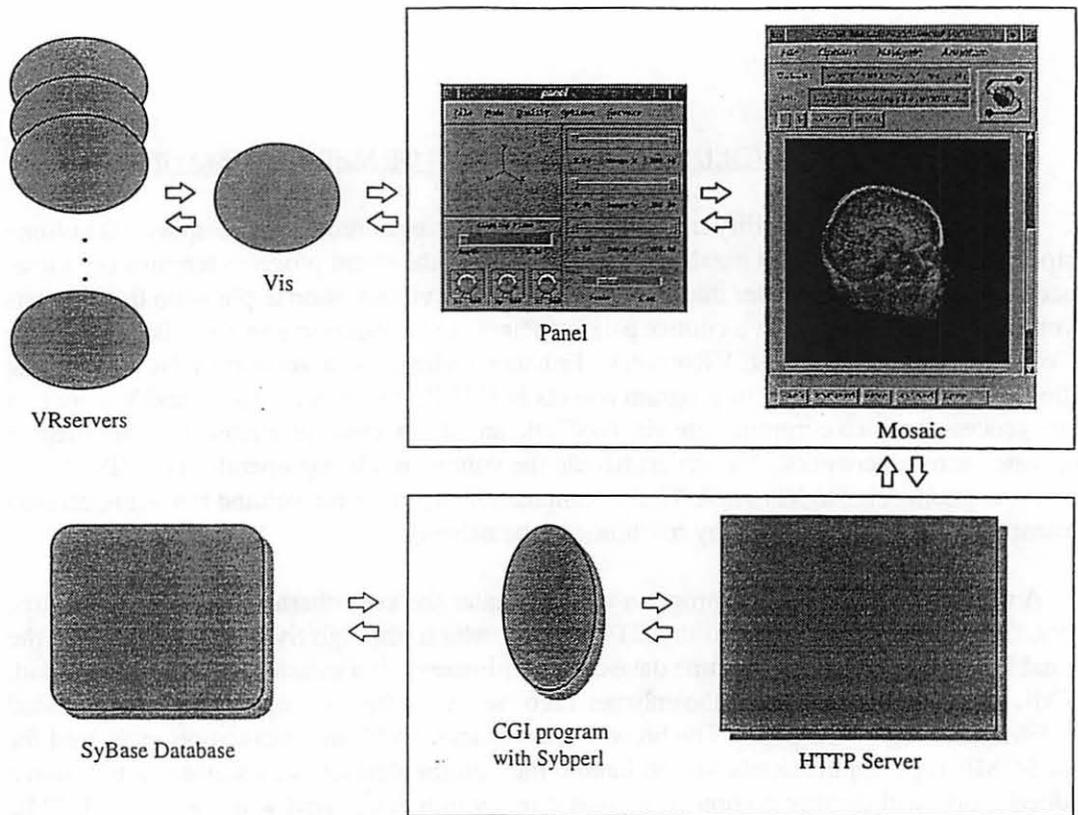


Fig 5: WWW Visualization System. Except for the modules in rectangular boxes, which may be thought of as machine boundaries, all other modules can potentially be running on an individual machine on the network.

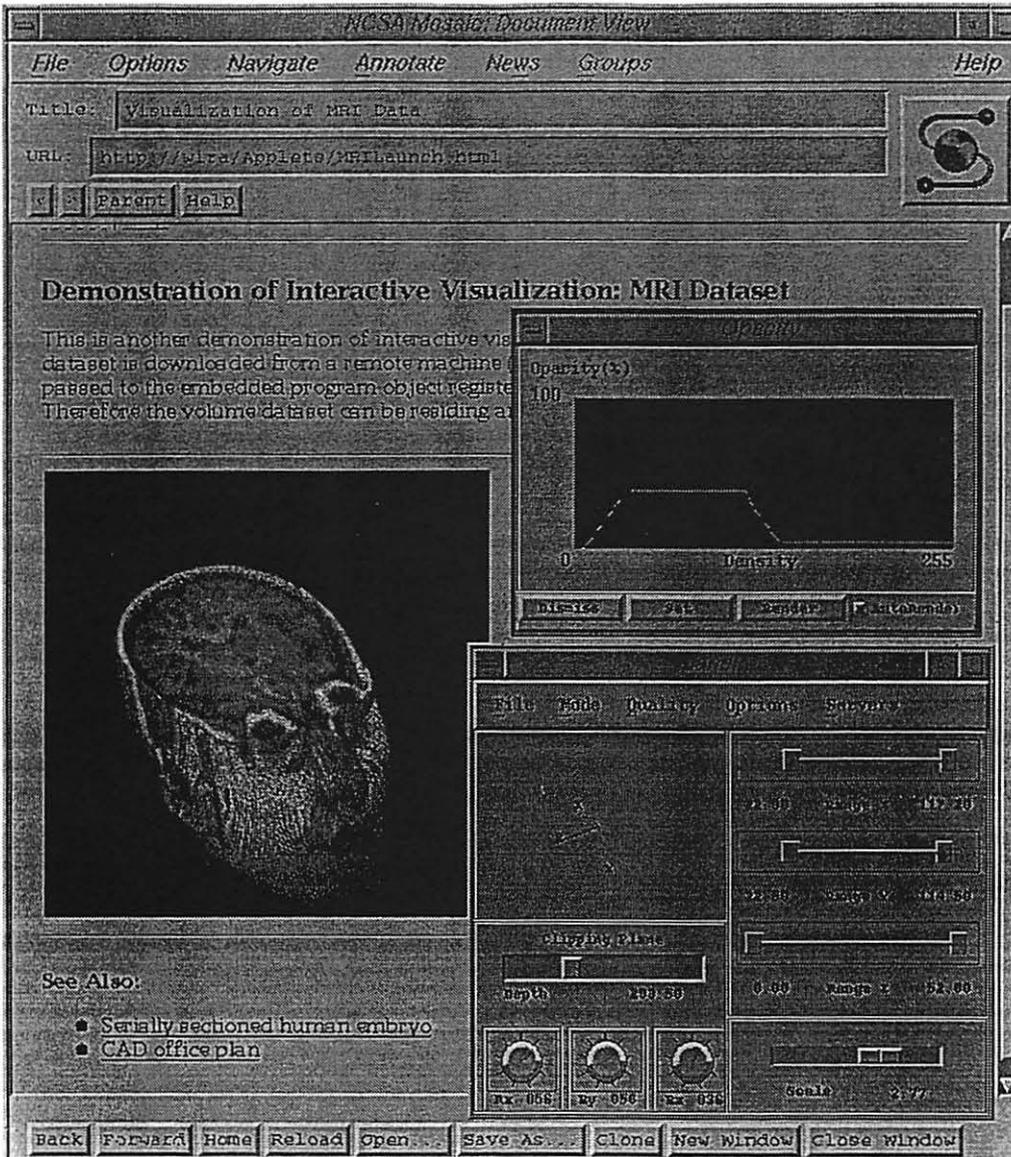


Fig 6 : Visualization of an MRI head volume dataset through the WWW.

## 6. SUMMARY

In this paper, we presented our interactive query and visualization software that is the first of its kind for medical imaging applications. This client/server information system integrates the technologies of hypermedia computing, relational databases, and computer graphics to provide efficient medical information dissemination on the Web. We have also shown that the WWW successfully interacts with the individual components necessary to facilitate interactive query and visualization of medical images. The “logic” and “textual data” components of the Breast Cancer education program are fully functional. The visualization server performs well for 3D brain images and work is in progress to query a multimodality image database for epilepsy.

The Web experience is so new that it is difficult to predict all of the features that will be integrated into Web servers or client browsers for medical imaging applications. This reported work is just the beginning of investigating the WWW for disseminating and viewing medical images. Many technical issues remain. For example, how to ensure data security and authenticity once the patient information are opened to the ubiquitous Web. Or, the use of more expressive, interactive programming language, such as Java, in developing efficient application programs. Efficient image compression algorithms must also be developed to optimize the transfer of large image data through the Web.

## ACKNOWLEDGMENTS

This research is funded by the US Army Medical R&D Command DAMD17-94-J-4338, National Science Foundation, NSF-IRI 9527982 and the UCSF Summer Undergraduate Research Program, 1995.

## REFERENCES

1. A Ford. *Spinning the Web: How to provide information on the Internet*. VNR, New York, 1995.
2. J Pitkow., et al. The GVU Center's Third WWW user survey, 1995.  
URL=[http://www.cc.gatech.edu/gvu/user\\_surveys/survey-04-1995](http://www.cc.gatech.edu/gvu/user_surveys/survey-04-1995).
3. Backbone Traffic Distribution Statistics. April, 1995.  
<http://www.cc.gatech.edu/guv/stats/NSF/merit.html>.
4. A Berson and G Anderson. SYBASE and client/server computing. McGraw-Hill, 1995.
5. Sybperl was developed by Michael Keppler.
6. E Quigley. *Perl by Example*, Prentice Hall PBR, New Jersey, 1995.
7. PolyMap was developed by CS Ang in the Center of Knowledge Management, UCSF.
8. CS Ang, DC Martin, and MD Doyle. Integrated control of distributed volume visualization through the World Wide Web, Technical Report, Center of Knowledge Management, UCSF, June, 1995.

Use of Image Co-registration and Visualization Techniques To Study Relationships Between MEG Neurophysiology and FDG-PET Metabolism in Epilepsy Imaging

Stephen TC Wong, Robert C Knowlton

(University of California, School of Medicine, San Francisco, CA 94143)

Many epilepsy patients who resist to anti-epileptic drugs are candidates for surgical treatment by locating and resecting the tissues of epileptogenic foci. However, determination of the epileptogenic zone is the most challenging step in the evaluation of patients for epilepsy surgery. Magnetoencephalography estimates the localization of interictal spike sources with a high degree of accuracy. In this presentation, we report our image registration and analysis experience with a large array biomagnetometer and a high resolution positron emission tomography (PET) scanner in the presurgical localization of spike sources in the evaluation of a group of patients representing a large spectrum of localization related epilepsies.

Patients chosen for study were selected randomly from a total of 65 patients evaluated for surgery last year at the UCSF Comprehensive Epilepsy Center. EEG was simultaneously recorded with MEG. A single equivalent current dipole model was used to estimate localization of spike sources. Magnetic source images (MSI) from interictal spike localizations were created by displaying dipole sources on each patient's volumetric MRI. MSI localization of spike disturbances was compared with video EEG (scalp and intracranial), MRI, and coregistered PET. Localization of the epileptogenic zone for analysis was based on ECoG, region of surgical resection, pathology, and surgical outcome.

All these multimodal neurologic data are collected and integrated into an image database system running Sybase SQL server (Sybase Inc., Emeryville, CA) on a SUN SPARC 690 MP machine (Sun Microsystems, Mountain View, CA). A Unix based client workstation (SUN SPARC LX) with an object oriented multimedia graphical user interface is developed to access and query the on-line multimedia data stored in the SQL server. Additional image analysis and visualization modules (VIDA<sup>TM</sup>) have been incorporated into the client GUI for decision aids.

Results of our computer-aided analysis indicates that dipole sources of interictal spikes were able to be estimated in 19 of 31 patients. In the group as a whole correlation of MSI with final localization of the epileptogenic zone was higher than either video EEG, MRI, or high resolution PET. MSI and PET provided the highest correlation of any combination of non-invasive modalities.

In conclusions, MSI can non-invasively predict localization of the epileptogenic zone as determined by ictal EEG, MRI, PET, and ECoG. When high resolution PET is concordant with MSI localization of epileptogenic zone can be predicted accurately even when intracranial ictal EEG recordings are ambiguous. Finally, MSI can aid in the optimal placement of intracranial electrodes. The electronic database management and registration of image and signals will enhance the analysis and productivity of such image technology research.

**Keywords:** Registration, PET, MEG, Neurophysiology, Glucose Metabolism, and Epilepsy

### Quantitative image measurements for outcome analysis of lung nodule treatment

Xiaoming Zhu, Ki-Nam Lee, STC Wong, and HK Huang (Univ. of California, San Francisco, CA 94143-0628)

In this study, we designed and implemented a temporal image database for outcome analysis of lung nodules based on spiral CT images. The software package is composed of three parts. They are, respectively, a database management system which stores patient image data and nodule information; a user-friendly graphical user interface (GUI) which allows a user to interact with the image database; and image processing tools that are designed to segment out lung nodules in the CT image with a simple mouse click anywhere inside a nodule. The image database uses the relational Sybase database system. Patient images and nodule information are stored in separate tables. Software interface has been designed to allow a user to retrieve any patient study from the picture archiving and communication system (PACS) into the image database.

Graphical user interface is developed with X-window in Motif development environment. It can display medical images in 16 bits by allowing a user to select window and level values to map the medical images into 8 bits for display. A user can choose different window and level settings for different anatomic features. Image segmentation is based on the pixel values of the lesion which is much higher than its air-like surroundings. The software automatically inspects adjacent slices to look for three dimensional (3D) information of a nodule. Segmentation results are presented to the user for visual inspection for any possible errors. Should there be any, which can be due to the attachment of the nodule with other tissues or thoracic walls, the software allows the user to manually edit it. Output from the imaging process tools, such as the volume of the nodule and its center of mass, is saved into a separate table in database for later retrieval for outcome analysis.

Accuracy of imaging processing tools has been verified in a phantom study. The average relative error was found to be within 5% for all nodules of size 6 mm and above. The software has some problem to make accurate volume estimates for nodule of 3 mm and smaller. Fortunately those nodules do not make significant overall volume contribution for lung metastasis. To insure the same nodule will not be counted twice, our display software keeps track of all nodules that have been examined and masks them with a red color. The temporal image database allows clinicians to collect and to process consecutive CT images of the same patient for quantitative information for outcome analysis during therapeutic treatment. This is possible due to the availability of the digital image data from PACS. The software is currently under clinical evaluation.

# **AN ATM DISTRIBUTED PACS SERVER for ICU APPLICATION**

**Joseph K. Lee, Albert W.K. Wong, H.K. Huang,  
Todd Bazzill, Jianguo Zhang, Kathy Andriole**

**Laboratory for Radiological Informatics  
Department of Radiology  
University of California, San Francisco  
San Francisco, CA 94143-0628**

## **ABSTRACT**

In order for PACS (Picture Archiving and Communications System) to better serve our intensive care units (ICUs), we, at University of California, San Francisco, (UCSF), have designed and developed a client/server application that is specifically tailored to provide fast, reliable access to our PACS data from diagnostic viewing stations in the ICUs. One of our utmost design criteria is to ensure consistent delivery of high speed, high performance data throughput, and yet, the system should be cost-effective and render minimal maintenance.

As high technology advances, we are able to utilize powerful mass storage device such as raid disk, which serves as a central image repository, to store images and data. We are also able to utilize Asynchronous Transfer Mode (ATM) technology, which is regarded as the prevailing technology for reliable, high speed data communications, to transfer large imagery data sets across systems and networks.

This paper describes the design and mechanism of how ICU viewing stations take advantages of sharing a high performance raid disk, and ATM technology in data transfer for timely delivery of images in a clinical setting.

## I. INTRODUCTION

As the computing trend evolves from the 80s into the 90s, a noticeable, dramatic revolution in technologies, both software and hardware, emerges. High volume storage device not only becomes available and readily accessible within reach, it also integrates easily with systems and applications in a seamless fashion. High speed, powerful workstation that delivers superb performance generally is regarded as the hardware of choice to run complex, real-time applications. And finally innovative networking technology that provides fast data transfer speed, with scaleable bandwidth and great flexibility not just becomes the prevailing trend for interconnecting computers and networks, but offers the ultimate solution for distributing information across the globe.

Our ICU application is developed by utilizing these latest technologies. The infrastructure consists of three major components. 1) High performance workstations are used for image display and application management. 2) A raid disk, which is capable of storing multi-giga bytes of data, is designated for massive storage of images. 3) An ATM network is established for data communications. In the following sections, we will discuss our implementation strategy in greater details.

## II. MATERIALS AND METHODS

### The Raid Disk

The raid disk is an array of disks integrates remarkably with some top-notch features. It provides raid, hot spares, a high speed fiber channel, and a dual-ported controller for system failover. Because of their modular, scaleable design, adding storage capacity and performance can be easily and cost-effectively installed when the system expands. To ensure data will not be lost, the raid disk provides raid levels 5, 1, and 0+1. For additional flexibility, it also provides raid 0 and independent disks. With the hot-spare option, a built-in spare goes in operation automatically when a faulty disk is detected, thus rebuilding data and restoring system to full speed. And with the warm-plug feature, replacing a faulty disk does not require system downtime, thus eliminating the timely reboot process.

## The ATM Technology

Asynchronous Transfer Mode, or ATM, is a multiplexing technology based on the switching of small fixed length packets of data called cells. Each cell is a 53-byte data including a 5-byte header which identifies the cell's route through the network. The remaining 48 bytes contains user data. Another advantage offered by ATM is its superb data transfer speed. At ATM OC3 the maximum data transmission bandwidth is at 155 Mega bits per second. Because of the high bandwidth and uniform switching of ATM networks, it can transfer large volumes of multimedia data such as images, audio, video and, as well as graphics and text from application to application at a much higher speed. Another distinctive feature in the ATM technology is its open architecture in the sense that it is not destined with a single physical medium or speed. It allows traffic from multiple sources to be switched to multiple destinations by fast ATM switch. Expanding local area networks can simply be interconnecting multiple ATM switches.

## The System Architecture

Our system infrastructure is based on the client/server computing principle. In a client/server architecture, the server is considered centralized and is designated to provide services upon request from its clients. Clients are generally running its own operating system software and applications. Only when services are requested, clients will then connect and communicate to the server. A network that connects the server and the clients provides the image and data communications.

The ICU server consists of a high performance workstation with a raid disk for central storage of images and reports. Dual portrait monitors are used for display. The following table specify the hardware configuration of the ICU server.

Table 1: ICU Server Configuration

Sun Sparc Workstation 20
128 Mega Bytes Memory
Sun Raid Disk: 30 Giga Bytes
2 Turbo GX Plus 4M Graphic Cards
ForeSystems SBA-200 ATM Sbus Adapter card
(2) 24-inch 1280 X 1600 portrait monitors

The ICU review station, which deploys as the client, also consists of a high performance workstation with an internal disk for local storage. Dual portrait monitors are used for display. The following table specify the hardware configuration of the ICU review station.

Table 2: ICU Review Station Configuration

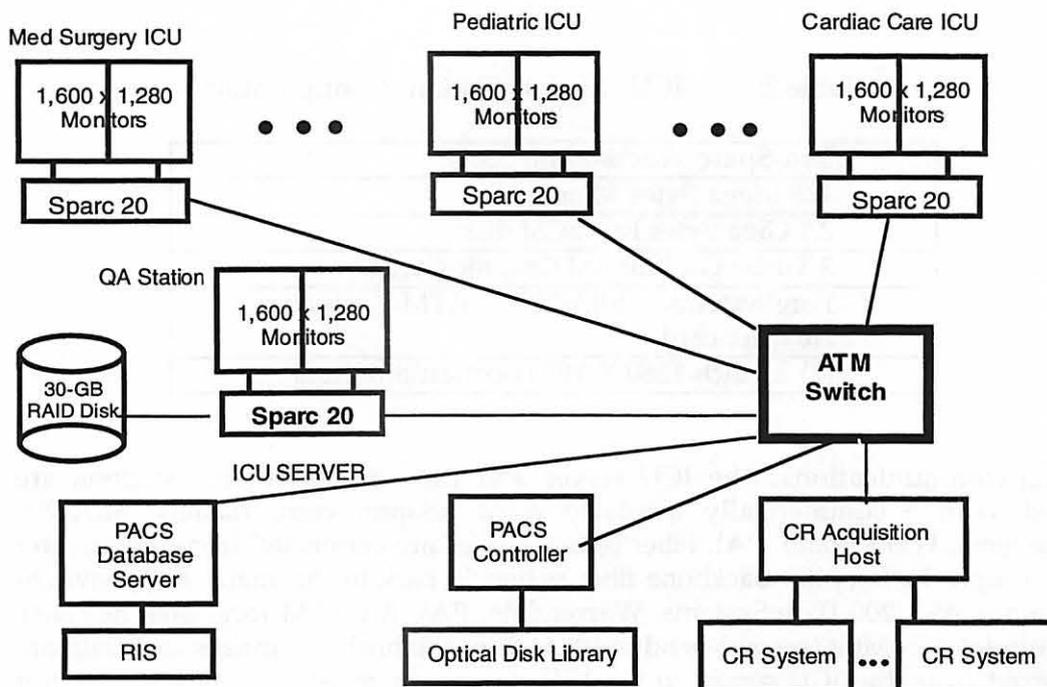
Sun Sparc Workstation 20
128 Mega Bytes Memory
2.1 Giga Bytes Fast SCSI disk
2 Turbo GX Plus 4M Graphic Cards
ForeSystems SBA-200 ATM Sbus Adapter card
(2) 24-inch 1280 X 1600 portrait monitors

For data communications, the ICU server and each of the review stations are installed with a commercially available ATM adapter card, namely SBA-200 (ForeSystems, Warrendale, PA). Fiber optical cables are connected from the adapter card, through the hospital backbone fiber network, back to the main ATM switch-ForeRunner ASX 200 (ForeSystems, Warrendale, PA). An ATM local area network that provides 155 Mbit/second bandwidth is thus established; images and data are transferred from the ICU server to the ICU stations in real-time using the higher speed available through ATM.

The PACS controller, which is also connected to the ATM local area network, provides automatic image routing to the ICU server. With four central processing units (CPUs) running on a Sun 690 machine (Sun Microsystems, Mountain View, CA), the PACS controller allows multiple processes running simultaneously with minimal shared CPU time. Upon receiving an image, PACS will validate the image, archive a copy to the optical disk library (1.3 Tera bytes, Cranel, OH), update the mirrored Sybase databases (Emeryville, CA), and then based on the routing code, send the image to the appropriate diagnostic review workstation, or to the ICU server. It also handles real-time image and diagnostic report retrieval requests from display stations.

The following table shows a schematic view of the ATM distributed ICU server and the PACS infrastructure.

Table 3. An ATM Distributed PACS Server for the ICUs



### The Image Acquisition Process

At UCSF, our ICUs include pediatric, cardiovascular, neurosurgery, medical surgery, cardiac care, and neonatal. Except neonatal, each ICU uses Computed Radiograph (CR) for conventional radiography. The technologist enters a two letter code designating a particular ICU when the CR reads the imaging plate. During the acquisition process, the CR image is first converted to DICOM 3 format (Digital Imaging Communications in Medicine, an industry standard in medical imaging communications), it then validates the patient ID, and runs a special utility to remove any white background and sends to PACS. Upon receiving the CR image, PACS archives its original image to the optical disk library, updates the mirrored Sybase databases, and routes a copy of the image to the ICU server. A successful acknowledgment is also sent back to the acquisition host so that the local image in the host computer can be purged.

## The ICU Application

Currently our ICU server is serving three units which include pediatric, medical surgery and cardiac care. The ICU server, which contains thirty high performance, low profile disk arrays, is capable of storing up to 30 Giga bytes of data and can be easily expanded as needed. Instead of configuring a large global database, which can bring all three ICU applications to a halt should the global database becomes inoperable, the ICU server maintains a separate database for each ICU station. In other words images are stored centrally, and yet, in a distributed database environment.

The process of maintaining separate databases for each ICU is relatively simple. The ground work involves configuring three database servers, each of them is responsible for all the activities of a particular ICU. When images arrive at the ICU server, based on the routing code, the ICU server will insert a patient's demographic data and its pertaining images to the appropriate database. Upon request from a ICU workstation, the image and textual data are sent to the client workstation through ATM transmission. The advantage of having a local database for each ICU is that even when one database goes down, the other ICUs can continue to operate without any interruption.

At each ICU workstation, a local directory that contains all the current patients of that ICU is maintained. A user can request images of a patient by first selecting the patient ID from the patient directory. Upon receiving the request, the ICU server routes the patient's images and demographic data back to the ICU workstation. Images are transferred in their original sizes to the workstation's memory cache, but are interpolated to 1,600 X 1,200 for the two display monitors. While one monitor maintains the most current image, the other monitor can be used to review the remaining set of images.

To manipulate images, many user-friendly, easy-to-use functions are available. Some of the most frequently used features include window and level, zoom, rotation, and image layouts. Through a pop-up window, a patient's demographic data and study information can be obtained. Real time access to historical images and diagnostic reports are also available. The image retrieval button allows a user to enter a patient ID and retrieve all the previous examinations from PACS. The report button allows instantaneous access to the Sybase database which stores all the diagnostic reports. When a report request is initiated, a report window will be open, and a user can browse through all the available reports accordingly. Because of direct access, historical images and diagnostic reports can be forwarded from PACS to an ICU workstation without severely penalizing the performance of the ICU server.

### III. RESULTS AND DISCUSSION

We have deployed three ICU workstations in three intensive care units in the hospital. For the past five months, performance is superior without any interruption from hardware, software or the ATM network. Physicians, on the average, are able to request and review the first image within 1.5 to 2.0 seconds.

The advantages of having a high performance ICU server to serve the ICUs are manifest and enormous. First of all, the ICU application is considered much simpler than other radiology services. A server separates the ICUs from other applications allows the ICU server to have a simpler system architecture design. Coupled with the innovative technologies of disk array and ATM, a high performance server can provide a timely delivery of images to the ICUs, and at the same time eliminate redundancy in storage components for each workstation. With the availability and affordability of raid disk and ATM networking tools, each workstation can receive its first image in about 1.5 to 2.0 seconds from the server, which is desirable in clinical operation. Finally, since all images are deposited in the disk array in the server, quality assurance and maintenance can be easily and readily performed at the server site without having to physically visit each ICU workstation.

Because of the popularity and success of the prototype ICU application, our plan is to continue to support the current ICU workstations. The future direction is to implement additional ICU workstations and set up ATM network in the remaining three ICUs. The ultimate goal is to have the ICU server to provide digital radiology service to all the ICUs.

### VI. REFERENCES

1. Wong STC, Lemke HU, Huang HK, "Application of ATM to distributed PACS environment", SPIE Medical Imaging- PACS Design and Evaluation 1995; 2425:430-439
2. Lemke HU, Wong STC, Huang HK, "System design criteria for ATM oriented broadband medical services", European Symposium- Advanced Networks and Services, March 1995

3. Albert W. K. Wong, H. K. Huang, Ronald L. Arenson, Joseph K. Lee, "Digital Archive System for radiologic images", RadioGraphics 1994; 14:1119-1126

4. H.K. Huang, Ronald L. Arenson, Albert W.K. Wong, Todd Bazzill, Andrew S.L. Lou, Katherine Andriole, Jun Wang, Jianguo Zhang, Stephen Wong, "Current Status of the UCSF Second Generation PACS", SPIE Medical Imaging Proceeding 1996

# High-performance image communication network with asynchronous transfer mode technology

Albert W.K. Wong, H.K. Huang, Joseph K. Lee  
Todd M. Bazzill, Xiaoming Zhu

albert\_wong@radmac1.ucsf.edu  
Laboratory for Radiological Informatics, Department of Radiology  
University of California, San Francisco  
San Francisco, CA 94143-0628

## ABSTRACT

Asynchronous transfer mode (ATM) technology has been implemented within our radiology department's hospital-wide PACS as well as in a wide area network (WAN) connecting affiliated hospitals. This paper describes our implementation strategies and the network performance observed in a clinical setting.

The image communication network for our PACS is composed of two network interfaces: ATM (OC-3, 155 Mbps) and Ethernet (10 Mbps). This communication network connects four major campus buildings and two remote hospitals, providing intra- and interbuilding communication for radiologic images including CT, MR, CR, US, and digitized screen-film images. The network links these modalities via their acquisition computers to the PACS controller and to display workstations. The ATM serves as the primary network for transmission of radiologic images and relevant data within the entire PACS. The standard Ethernet is used as a backup network for ATM. It interconnects all PACS components including radiologic imaging systems, acquisition computers, display workstations, the PACS controller, the database servers, and the RIS and HIS.

Our communication network operates on a 24 hrs/day, 7 days/week basis. Performance of the ATM network was evaluated in terms of disk-to-disk, disk-to-memory, and memory-to-memory transmission rates. The average memory-to-memory transmission rate over the wide area ATM network was 8.3 MBytes/s, which corresponds to transferring a 40-slice (or, 20-MByte) CT examination to a remote site in less than 3 seconds. With the emerging ATM technology, we believe that ATM-based digital communication network is a suitable choice for large-scale PACS involving both LAN and WAN.

**KEY WORDS:** picture archiving and communication systems (PACS), asynchronous transfer mode (ATM), communication network, local area network, wide area network, network performance.

## 1. INTRODUCTION

We have previously developed a large-scale picture archiving and communication system (PACS) for clinical use in a radiology department<sup>[1]</sup>. The image communication network implemented in this PACS was composed of three types of network interfaces: Ethernet, Fiber Distributed Data Interface (FDDI), and the 1-Gbps bandwidth UltraNet (Computer Network

Technology, Oakland, Calif). This three-tiered network provided independent pathways for data transmission, resulting in significant improvement of transferring high-volume radiologic images among individual PACS components<sup>[2,3]</sup>. However, this high-capacity digital network lacked wide area network (WAN) connectivity, and was therefore limited its functionality to local area network (LAN) applications.

The second-generation PACS<sup>[4,5]</sup> developed at UCSF has adapted asynchronous transfer mode (ATM) technology to its image communication network that provides both LAN and WAN connections for the PACS. This comprehensive digital network, composed of high-speed ATM and standard Ethernet, interconnects four major UCSF campus buildings: Moffitt Hospital, Long Hospital, the Ambulatory Care Center, and the Campus Library, and two remote affiliated hospitals: Mt. Zion Hospital (MZH) and San Francisco Veterans Administration Medical Center (SFVAMC), in the San Francisco Bay area<sup>[6,7]</sup>. The network has been in operation since March, 1995. It provides intra- and interbuilding communication for radiologic images throughout the entire PACS. This paper describes our implementation strategies and the network performance observed in a clinical setting.

## 2. UCSF PACS IMAGE COMMUNICATION

Currently, the UCSF PACS acquires 2.5 GBytes of radiologic images daily from four magnetic resonance (MR) imagers, three computed tomographic (CT) scanners, two computed radiographic (CR) systems, one ultrasound (US) PACS module, and two laser film scanners (LS). These images are sent from their acquisition computers via the PACS communication network to the PACS controller, where they are archived to optical disks for long-term storage and distributed to display workstations for clinical review. The communication network uses ATM as the primary network, with Ethernet and T1 as the backups for LAN and WAN, respectively. Failure of the ATM automatically triggers the communication network to reconfigure itself so that images and relevant data will be transmitted over the Ethernet or the T1 lines.

### 2.1 The networks

#### ATM

The ATM network interconnects four major campus buildings and links UCSF to the two remote affiliated hospitals, MZH and SFVAMC. The network operates at a bandwidth of 155 Mbps (OC-3) and is used for image and data communication among the acquisition nodes, the database servers, the PACS controller, the intensive care units (ICU) image server, the Mac file server, and the display workstations. The network interconnects these components via the local ATM switches as well as the public ATM switch, constituting an integrated digital imaging network that provides high-speed communication for radiologic images and relevant data throughout the entire PACS. Figure 1 shows the ATM WAN connections between UCSF, MZH, and SFVAMC, and Figure 2 shows the configuration of the ATM LAN at UCSF.

ATM is a connection-oriented switching and multiplexing technology that transmits information including video, voice and data using 53-byte cells (5-byte header and 48-byte information). ATM can operate at different optical carrier (OC) levels such as OC-1 (51 Mbps), OC-3 (155 Mbps), or OC-12 (622 Mbps). It allows the same high-bandwidth LAN applications to be transported across the WAN.

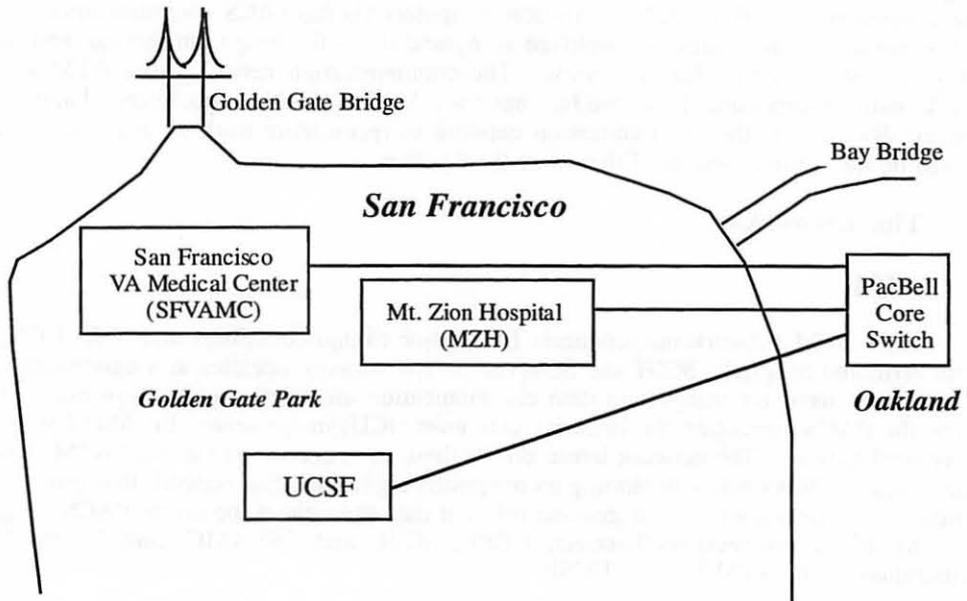
## Ethernet

The standard Ethernet is used for communication between radiologic imaging devices and their corresponding acquisition computers. In addition, the Ethernet network interconnects all PACS components including acquisition nodes, display workstations, database servers, the PACS controller, the ICU image server, and the Mac file server. This global Ethernet serves as a backup network for the high-speed ATM LAN. Figure 3 shows the configuration of the PACS Ethernet network at UCSF.

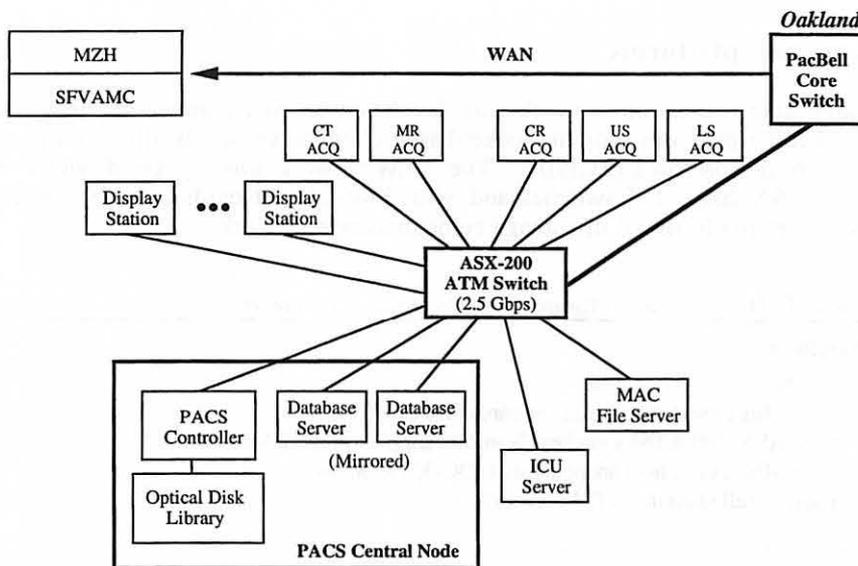
Ethernet is specified at a bandwidth of 10 Mbps. The multipoint characteristics of the bus topology used by Ethernet provides multi-access to the PACS computers so that any new node can be tapped into the existing Ethernet LAN without reconfiguring the network. One disadvantage of using Ethernet is that the transmission rates of individual communication processes drop significantly while multiple processes are running simultaneously.

## T1

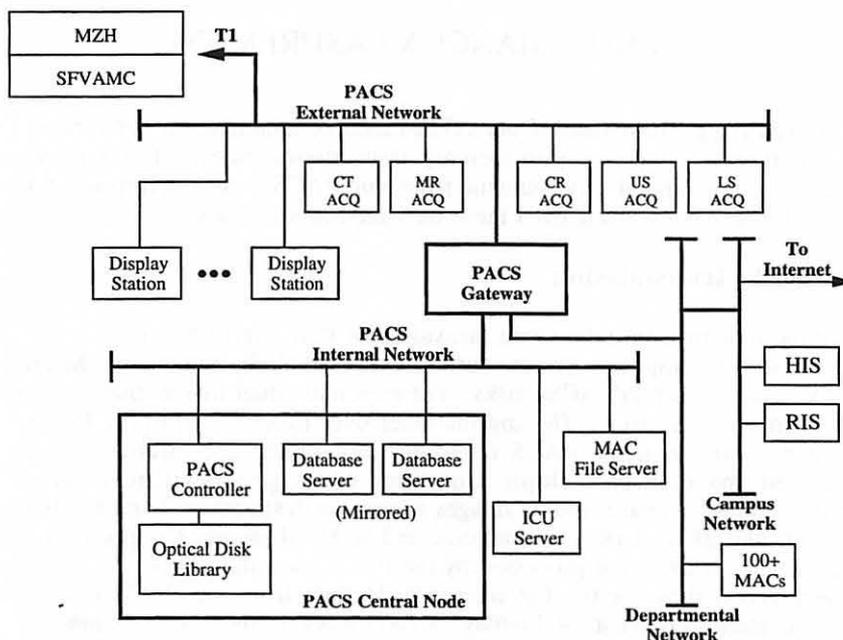
T1 operates at a bandwidth of 1.544 Mbps and is used as a backup network for the ATM WAN. This T1-link connects UCSF to MZH and SFVAMC, which guarantees WAN connectivity of the PACS in the event of ATM failure.



**Figure 1.** ATM connections between UCSF, MZH, and SFVAMC in the San Francisco Bay area. All connections are established via Pacific Bell's core switch in Oakland, Calif.



**Figure 2.** Configuration of the PACS ATM network. All acquisition nodes and display workstations, the database servers, the PACS controller, the ICU image server, and the Mac file server are interconnected via a Fore Systems' ASX-200 ATM switch. ATM WAN connections between UCSF, MZH, and SFVAMC are established through Pacific Bell's core switch in Oakland, California.



**Figure 3.** Configuration of the PACS Ethernet network. All acquisition nodes and display workstations are connected to the PACS external network, and the database servers, the PACS controller, the ICU image server, and the Mac file server are connected to the PACS internal network. These two networks (Class C) are both connected to the campus Ethernet (Class B) via the PACS gateway. The departmental network is a subnet of the campus network. T1 lines are used to connect UCSF to MZH and SFVAMC.

## 2.2 Development platforms

The image communication network was developed on Sun computers (Sun Microsystems, Mountain View, Calif.) platform running socket-based client/server applications using transmission control protocol/Internet protocol (TCP/IP). The ATM network was equipped with Fore Systems' (Warrendale, PA) ASX-200 ATM switches and SBus host adapter cards. Table 1 summarizes the hardware and software platforms of this image communication network.

**Table 1.** Hardware and software platforms of the image communication network

---

Hardware:

- Sun Sparc computer systems
- Sun SBus Ethernet host adapter cards (10Mbps)
- Fore ASX-200 ATM switches (Non-blocking, 24 ports, 2.5 Gbps)
- Fore SBus ATM host adapter cards (OC-3, 155 Mbps)
- Pacific Bell Oakland ATM core switch

Software:

- SunOS/Solaris operating systems
- C programming language
- TCP/IP transmission protocols
- Unix socket communication
- Client/server applications

---

## 3. PERFORMANCE MEASUREMENTS

We measured the performance of our ATM-based communication network in terms of disk-to-disk, disk-to-memory, and memory-to-memory transmission rates. These measurements were conducted in a real-time clinical environment from our PACS. A description of the methods of measurement and the results obtained from these measurements follows.

### 3.1 Disk-to-disk transmission

Disk-to-disk transmission rates were measured for transferring real-time CR images from the PACS controller's small computer systems interface (SCSI) disks via the ATM and the Ethernet networks to the ICU image server's SCSI disks. For each individual image, the image server received two copies of the image, one over ATM and the other over Ethernet, from the PACS controller. In this performance measurement, the PACS controller represented a centralized image management engine that accepted images from multiple acquisition nodes, processed image header information and updated PACS global database, routed images to various destinations (e.g., the ICU image server, the Mac file server, and 2K and 1K workstations), and archived images to optical disk library. Over 10 GBytes of data transactions were processed by the PACS controller daily.

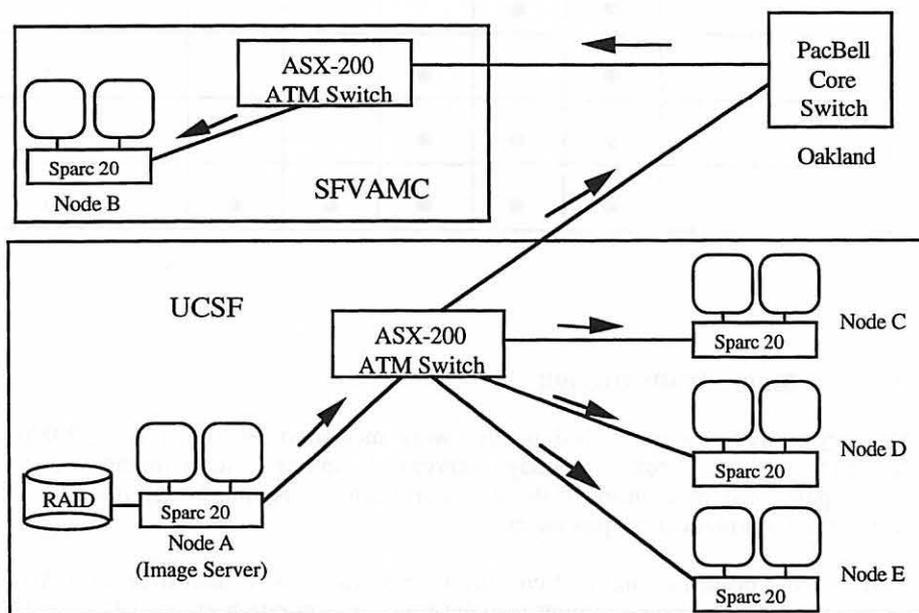
The transmission time for the CR images to be sent from the PACS controller to the ICU image server were measured on a 24 hrs/day, 7 days/week basis. Approximately one month of transaction data was analyzed. Results from the performance measurement (Table 2) showed that, even with the constraint of slow disk input/output (I/O) speed, the ATM network still outperformed Ethernet by almost 200%. This calculation was based on the measured transmission rates of ATM (1.247 MBytes/s) and Ethernet (452 KBytes/s).

**Table 2.** ATM vs Ethernet disk-to-disk transmission rates

	ATM	Ethernet
Average Throughput (KBytes/s)	1,247	452

### 3.2 Disk-to-memory and memory-to-memory transmissions

We used our ICU PACS module<sup>[8,9]</sup> as a test bed to measure disk-to-memory and memory-to-memory transmission rates of the ATM network. In the experiment, five dual-monitor, 1,600 x 1,280-resolution ICU workstations running motif-based applications on Sun Sparc 20 computers were interconnected via two ASX-200 ATM switches (Figure 4). Node A represented the ICU image server that distributed CR images over ATM to the display workstations (Nodes C, D, and E) located at the ICU ward. This image server, configured with 30 GBytes of high-performance redundant array of inexpensive disks (RAID), served as a central storage of images for all the ICU workstations. Node B represented a remote display workstation at SFVAMC that was connected to the image server at UCSF via Pacific Bell's core switch.



**Figure 4.** Diagram shows computers and equipment used in the disk-to-memory and memory-to-memory performance measurements. Five dual-monitor, 1,600 x 1,280-resolution display workstations (nodes A, B, C, D, and E) running ICU applications were used with two ASX-200 ATM switches. Node A represented the ICU image server configured with a 30-GByte RAID disk. Node B at SFVAMC was connected to UCSF via Pacific Bell's core switch in Oakland.

### Disk-to-memory transmission

Disk-to-memory transmission rates were measured for transferring eighty CR images, each 8 MBytes in size, from the image server's RAID disks over ATM to individual computer system memory of the ICU workstations. The socket-based client/server communication processes used in the performance measurement allocated 64-KByte TCP window size and 128-KByte buffer size for the data transfers.

Table 3 shows the results from point-to-point and concurrent transmissions of data over the LAN and the WAN. The average point-to-point transmission rate between UCSF and SFVAMC was 5.2 MBytes, which corresponds to transferring an 8-MByte CR image from a local magnetic storage over the WAN to a remote display memory in less than two seconds. The aggregate throughput for the image server to transfer images simultaneously to the four ICU workstations was 7.4 MBytes/s.

**Table 3.** Disk-to-memory transmission rates

Number of Concurrent Processes	Sender	Receiver				Aggregate Throughput (MBytes/s)
	A	B	C	D	E	
1	●	●				5.2
1	●		●			5.8
2	●	●	●			6.9
4	●	●	●	●	●	7.4

### Memory-to-memory transmission

Memory-to-memory transmission rates were measured for transferring 2,000 data packets, each 128 KBytes in size, from the image server's computer system memory over ATM to the individual computer system memory of the ICU workstations. Again, 64 KBytes of TCP window size were used for the communication processes.

Results from point-to-point and concurrent transmissions of data over the LAN and the WAN are shown in Table 4. The average point-to-point transmission rate between UCSF and SFVAMC was 8.3 MBytes, which corresponds to transferring a 40-slice (or, 20-MByte) CT examination over the WAN to a remote site in less than 3 seconds. The aggregate throughput for the image server to transfer data simultaneously to the four ICU workstations was 16.0 MBytes/s (or, 128 Mbps), which is equivalent to 82.6% of ATM's 155-Mbps signaling rate.

**Table 4.** Memory-to-memory transmission rates

Number of Concurrent Processes	Sender	Receiver				Aggregate Throughput (MBytes/s)
	A	B	C	D	E	
1	●	●				8.3
1	●		●			10.4
2	●	●	●			13.4
4	●	●	●	●	●	16.0

#### 4. DISCUSSION

We have successfully implemented ATM technology within our radiology department's hospital-wide PACS as well as in a WAN connecting affiliated hospitals. The comprehensive digital network, composed of ATM and Ethernet, provides high-speed yet reliable communications for radiologic images and relevant data over both LAN and WAN throughout the entire PACS.

From the results of our measurements, we believe that the transmission rates of ATM can satisfy most radiologic image communication requirements and hence a digital communication network with ATM technology would be a suitable choice for large-scale PACS involving both LAN and WAN. Ethernet, on the other hand, is inadequate to be used for image communication, particularly in a multimodal PACS environment, because of the shared bandwidth characteristics inherited from the bus topology that Ethernet uses.

Factors affecting the throughputs of an ATM network include: (1) processing overhead of the network protocols (e.g., TCP/IP); (2) low-speed disk I/O of the computer system; and (3) I/O buffering within the computer system, the ATM switch, and the applications. Among these factors, further improvements in network transmission protocols and in disk I/O speed will be the major efforts to optimize the throughput performance of ATM.

#### 5. ACKNOWLEDGMENT

This work was partially supported by CalREN grant ATMN-007 (Pacific Bell) and NLM Contract N01-LM-4-3508.

## 6. REFERENCES

1. H.K. Huang, R.K. Taira, S.L. Lou, A.W.K. Wong, et al, "Implementation of a large-scale picture archiving and communication system", *J. Comp. Med. Imaging and Graphics*, Vol 17(1):1-11, 1993.
2. A.W.K. Wong, B.K. Stewart, S.L. Lou, K.K. Chan, H.K. Huang, "Multiple communication networks for a radiological PACS", *SPIE Med. Imaging Proceedings*, Vol 1446:73-80, 1991.
3. A.W.K. Wong, H.K. Huang, "Subsystem throughputs of a clinical picture archiving and communication system", *J. Digital Imaging*, Vol 5(4):252-261, 1992.
4. H.K. Huang, R.L. Arenson, S.L. Lou, A.W.K. Wong, K.P. Andriole, et al, "Second generation picture archiving and communication systems", *SPIE Med. Imaging Proceedings*, Vol 2165:527-535, 1994.
5. R.L. Arenson, D.E. Avrin, A. Wong, R.G. Gould, H.K. Huang, "Second generation folder manager for PACS", *Comp. Appl. to Assist Radiology*, 601-605, 1994.
6. H.K. Huang, R.L. Arenson, W.P. Dillon, S.L. Lou, T. Bazzill, A.W.K. Wong, "Asynchronous transfer mode technology for radiologic image communication", *Am. J. Roentgenology*, Vol 164:1533-1536, 1995.
7. A.W.K. Wong, H.K. Huang, T.M. Bazzill, J.K. Lee, I. Bobkov, M.J. Moskowitz, "Performance characteristics of a digital imaging network with asynchronous transfer mode technology", *Radiology*, Vol 197(P):259, 1995.
8. H.K. Huang, A.W.K. Wong, T.M. Bazzill, K.P. Andriole, J.K. Lee, J.G. Zhang, "Asynchronous transfer mode - distributed PACS server for intensive care unit application", *Radiology*, Vol 197(P):547, 1995.
9. J.K. Lee, A.W.K. Wong, H.K. Huang, T.M. Bazzill, J. Zhang, K.P. Andriole, "ATM distributed PACS server for ICU application", *SPIE Med. Imaging Proceedings*, 1996 (In press).

## Clinical comparison of CR and screen-film for imaging the critically-ill neonate

Katherine P. Andriole, Robert C. Brasch, Charles A. Gooding, Robert G. Gould, Pierre A. Cohen, Ingo R. Rencken, and H.K. Bernie Huang

Department of Radiology  
University of California at San Francisco  
San Francisco, CA 94143-0628

### ABSTRACT

A clinical comparison of computed radiography (CR) versus screen-film for imaging the critically-ill neonate is performed, utilizing a modified (hybrid) film cassette containing a CR (standard ST-V) imaging plate, a conventional screen and film, allowing simultaneous acquisition of perfectly matched CR and plain film images. For 100 portable neonatal chest and abdominal projection radiographs, plain film was subjectively compared to CR hardcopy. Three pediatric radiologists graded overall image quality on a scale of one (poor) to five (excellent), as well as visualization of various anatomic structures (i.e., lung parenchyma, pulmonary vasculature, tubes/lines) and pathological findings (i.e., pulmonary interstitial emphysema, pleural effusion, pneumothorax). Results analyzed using a combined *kappa* statistic of the differences between scores from each matched set, combined over the three readers showed no statistically significant difference in overall image quality between screen-film and CR ( $p=0.19$ ). Similarly, no statistically significant difference was seen between screen-film and CR for anatomic structure visualization and for visualization of pathological findings. These results indicate that the image quality of CR is comparable to plain film, and that CR may be a suitable alternative to screen-film imaging for portable neonatal chest and abdominal examinations.

**KEYWORDS:** computed radiography, neonatal ICU, comparative clinical studies

### 1. INTRODUCTION

Previous studies by ourselves<sup>1,2</sup> and others<sup>3-5</sup> have supported the notion that the use of a hybrid cassette containing both a computed radiography (CR) imaging plate and conventional screen-film for the single-exposure, simultaneous acquisition of images can facilitate the direct clinical comparison of conventional screen-film and CR radiographs. And it is believed that the dual-image recording method is preferable to two separate exposures, one each for standard plain film and CR, when a more objective, controlled clinical comparison study is desired. Such a hybrid cassette is utilized here to evaluate the clinical utility of a CR imaging plate versus plain film for neonatal intensive care unit (ICU) applications.

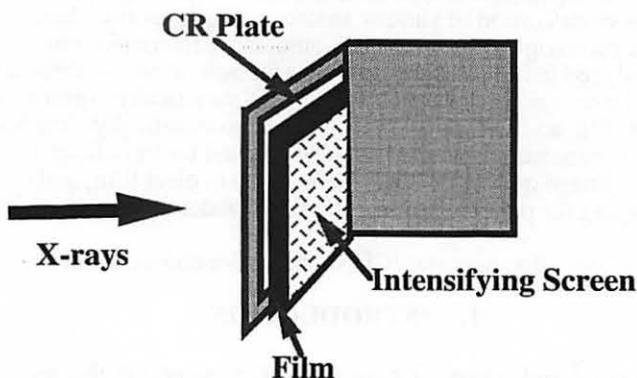
Among the motivating factors behind this examination are the fact that at our institution, there is a 52-bed Newborn Intensive Care Unit (NBICU) on the fifteenth floor of the main hospital, while the pediatric radiology reading room is on the third floor. Currently, film is duplicated for the ICU for immediate review up in the unit, and the original film is sent to radiology for primary diagnosis. Undue delay in this service can negatively impact on clinical care. If this study determines that CR is comparable to screen-film, then the current practice of film duplication for ICU review could be eliminated. The use of CR in place of screen-film would give the neonatologist immediate access to patients' images on a local PACS (picture archiving and communication system) ICU display station, and the pediatric radiologist would have immediate access to the same images in their reading room.

The latest version standard ST-V (Fuji Medical Systems U.S.A., Inc.) photostimulable phosphor CR imaging plate was used in a modified film cassette containing both a CR plate and conventional screen-film for a dual-image recording technique similar to that described by Chotas *et al.*<sup>3</sup> and further examined by ourselves<sup>1,2</sup> and Wilson *et al.*<sup>5</sup>. Note that the ST-V 8 inch by 10 inch imaging plates are scanned with a density of 10 pixels/mm. For 100 portable neonatal chest and abdominal examinations, plain film was subjectively compared to CR hardcopy, all recorded from a single simultaneous acquisition exposure. Hardcopy images

were masked to disguise the origin of the film and were presented to the readers in random order for scoring. Three pediatric radiologists graded each image for overall image quality on a scale of one (poor) to five (excellent). The readers also graded the visualization of various structures in the chest including lung parenchyma, pulmonary vasculature, soft tissues, bone trabeculae, spine detail, tubes, and lines. The visualization of pathological findings including pulmonary interstitial emphysema, pleural effusion, interstitial thickening, pneumothorax, pneumomediastinum, pneumoperitoneum, GI intramural air, and portal venous air were also evaluated.

## 2. MATERIALS AND METHODS

The hybrid cassette or dual detector configuration used for image acquisition in this clinical comparison study is diagrammed in Figure 1. It consists of a film cassette with the front intensifying screen removed and replaced by a CR plate, reversed in the cassette so that the phosphor side is facing the film which is the middle piece of the hybrid cassette, followed by the back intensifying screen.



**Figure 1.** The Hybrid Cassette consists of a conventional screen-film cassette with the front intensifying screen removed and replaced with a CR imaging plate (reversed in the cassette), followed by the film and the back intensifying screen.

The specific screen-film cassette system used was an 8 inch by 10 inch Du Pont carbon fiber film cassette, along with a Du Pont Quanta Fast Detail / Cronex 10 screen-film system. The Quanta Fast Detail is a blue/UV-emitting rare earth (Niobium-activated yttrium tantalite) intensifying screen which when matched with Cronex 10 film, creates a medium (250) speed, low relative noise screen-film system with a (high) spatial resolution of 10 lines/mm<sup>6</sup>. All exposures were performed on the same GE amx4 portable X-ray machine using portable neonatal chest technique factors (50-60 KVp, 1 MAS), and plain film was processed under normal dark room conditions using a Kodak RPX-omat Model M6B film processor.

The computed radiography components utilized included a Fuji 8 inch by 10 inch carbon fiber CR cassette with the standard Fuji type ST-V CR imaging plate which has a scanning density of 10 pixels/mm. Each photostimulable phosphor plate contains europium doped barium-fluoro-halide-halide ( $\text{BaFBrI:Eu}^{2+}$ ) which emits in the blue-green (400 nm) spectral range<sup>7</sup>. CR plates were read by the same Fuji FCR 9000 laser reader and CR film was processed by a Fuji CR-LP414 laser printer.

The specific detector parts of the hybrid cassette used from each system included, the Du Pont carbon fiber film cassette with only the back Quanta Fast Detail intensifying screen in place, Cronex 10 film, and the Fuji type V plate. (Refer to Figure 1.) Both the film and the storage-phosphor plate were placed into and removed

from the cassette under dark room conditions. The exposed film was processed in the usual way and the exposed plate was removed from the hybrid film cassette, placed into a CR cassette and processed by the laser scanner. Three display formats result from the detector configurations described above: standard analog film, hardcopy CR (digital image on film) and softcopy CR (digital data viewable on a display monitor). The experimental technique followed, emulated that utilized in our hospital for portable pediatric chest X-ray examinations, with exposures in the 50 to 60 KVP range at 1 MAS.

Over a four month period, all weekday morning films which were requested for Newborn ICU patients were imaged using the hybrid cassette with the ST-V plate. The X-ray exposure used clinically for the hybrid configuration was slightly higher than that used for the conventional 250 speed screen-film system currently utilized at our institution for portables in the newborn ICU. The dosage was, however, well within the reasonable, permissible limits. In fact, the exposure difference corresponded to one step in KVP, which may be within the range of variation in techniques between different X-ray technologists. Daily films were read from the plain film from the hybrid cassette. Hardcopy images were masked to disguise the origin of the film (either CR or plain film) and were presented to three pediatric radiologists in random order for scoring.

Radiologists rated the ability to see the structures of the chest including lung parenchyma, pulmonary vasculature, soft tissues, bone trabeculae, and detail of the spine, as well as nasogastric and endotracheal tubes. They also graded the visualization of pathological findings including pulmonary interstitial emphysema (PIE), pleural effusion, interstitial thickening, edema, pneumothorax, pneumomediastinum, pneumoperitoneum, gastrointestinal (GI) air, and portal venous air, and the overall image quality using a scale of one to five as follows:

<u>Score</u>	<u>Evaluation</u>	<u>Description</u>
1	unsatisfactory	very poor image, structure not visible, nondiagnostic
2	poor	structure can be visualized but quality is unacceptable for diagnosis
3	satisfactory	study is adequate for diagnosis, fair quality
4	good	good image, acceptable quality
5	excellent	structure is beautifully visualized
NA	not applicable or absent	

### 3. RESULTS

A representative sampling of pathological categories and/or findings was included in the test set, obtained over a typical four month period. Among the 100 total matched image sets included in the study, all but 15 had at least one tube or line present and most were standard AP (anterior-to-posterior, referring to the X-ray source-to-patient-to detector configuration) portable chest examinations. Table 1 details the test set makeup.

<u>VIEWS</u>		<u>PATHOLOGICAL CATEGORY/FINDING*</u>			
AP Chest	80	Normal	20	No Findings	24
AP Abdomen	5	RDS	38	Pulmonary Edema	21
Chest + Abdomen	14	S/P Surgery	18	Atelectasis	20
Lateral Chest	1	BPD	10	PIE	16
		Anasarca	1	Pleural Effusion	15
		Rickets	5	Interstitial Thickening	8
		CHD	4	Pneumothorax, -med	5
		Infection/Cyst	3	Cardiomegaly	1
				Hypoplastic Lung	1

**Table 1.** Test set makeup. \* Note: images may contain more than one finding. AP: Anterior-to-Posterior, BPD: Broncho-Pulmonary Displasia, CHD: Congenital Heart Disease, -med: Pneumomediastinum, PIE: Pulmonary Interstitial Emphysema, RDS: Respiratory Distress Syndrome, S/P: Status Post.

Results obtained using a combined *kappa* statistic analysis of the differences between a reader's scores for each of the matched sets of images, combined over each of the three readers are shown in Table 2, for the overall image quality evaluation and for visualization of each anatomic structure. Note that the *p* values show no statistically significant difference between screen-film and CR in overall image quality or for visualization of any of the anatomic structures including tubes and lines. A similar trend is seen in preliminary analysis of visualization of the pathological categories and findings listed in Table 1.

<u>OVERALL IMAGE QUALITY</u>	<i>p</i> = 0.19
<u>ANATOMIC STRUCTURE VISUALIZATION</u>	
Lung Parenchyma	<i>p</i> = 0.10
Pulmonary Vasculature	<i>p</i> = 0.17
Soft Tissue	<i>p</i> = 0.28
Bone Trabeculae	<i>p</i> = 0.16
Spine Detail	<i>p</i> = 0.13
Tube/Line	<i>p</i> = 0.11

**Table 2.** Combined kappa statistic analysis of score differences rating overall image quality and visualization of anatomic structures, tubes and lines.

#### 4. DISCUSSION

The results obtained in this study lead to the conclusion that clinically there is no difference between computed radiography and screen-film for this application to imaging the critically-ill neonate. Thus, CR may be a suitable alternative to screen-film radiography for bedside neonatal ICU images; particularly when assessed with the advantages gained through the use of a digital modality, such as the ability to rapidly distribute images to display workstations in the intensive care unit or anywhere within the picture archiving and communication system (PACS); the ability to manipulate the digital image through window and level, inverse and other image processing and enhancement procedures; and the ability to archive digital images for future retrieval for clinical, research and teaching purposes. Based on the results of this clinical comparison between screen-film and CR, UCSF will begin routine use of CR in the newborn intensive care unit.

Current continuing studies include a comparison of the matched clinical sets of screen-film and CR film, to CR softcopy viewing on a high-resolution 2K by 2K (2048 pixels by 2048 pixels) display workstation which uses a Sun 4/470 CPU for the patient demographic user interface, two 2048 by 2500 line Megascan (currently EMed) monitors, and a Storage Concepts parallel transfer disk. Several image pre-processing tasks are involved in preparing a CR chest image from its original raw data format to that compatible for softcopy display at the viewing workstation. These include: background removal to eliminate the potentially distracting unexposed (white) background due to X-ray collimation at the time of exposure; image orientation correction to determine the original image object orientation, and rotate the image, if necessary, to a standard viewing position; default optimal look-up table calculation for initial display parameters based on the anatomical region, the image histogram and the minimum and maximum gray levels; and auto-contrast enhancement for different tissue densities (i.e., muscle, fat, mediastinum, lung, bone) for further display window and level optimization based on piecewise-linear look-up tables<sup>8</sup>. For softcopy viewing, a reader is allowed to interactively window and level the image if so desired, from the default values, as well as invert the gray levels (white to black and black to white), magnify or rotate the image.

Additional ongoing related studies include comparison of image quality evaluation assessed by radiologists of differing experience levels and image quality evaluation assessed by non-radiologist ICU clinicians.

## 5. ACKNOWLEDGMENTS

The authors would like to thank Katrina Lam, R.T. (UCSF), Wyatt M. Tellis (UC Berkeley) and James W. Sayre, Ph.D. (UCLA) for their contributions to this study.

## 6. REFERENCES

1. Andriole KP, Gould RG, Bazzill TM, and Huang HK. "Physical parameter analysis of a high-resolution CR imaging plate versus conventional screen-film radiography." *79th Scientific Assembly and Annual Meeting of the Radiological Society of North America*, Chicago, November 28 - December 3, 1993.
2. Andriole KP, Gooding CA, Gould RG, and Huang HK. "Analysis of a high-resolution CR imaging plate versus conventional screen-film for neonatal ICU applications." *SPIE, Medical Imaging 1994: Physics of Medical Imaging*, Newport Beach, CA, Vol. 2163, pp. 89-97, February 13-18, 1994.
3. Chotas HG, Dobbins JT, Floyd CE, and Ravin CE. "Single-exposure conventional and computed radiography image acquisition." *Investigative Radiology*, Vol. 26, No. 5, pp. 438-445, May 1991.
4. Sanada S, Doi K, Su XW, Yin FF, Giger ML, and MacMahon H. "Comparison of imaging properties of a computed radiography system and screen-film systems." *Medical Physics*, Vol. 18, No. 3, pp. 414-420, May/June 1991.
5. Wilson AJ and West OC. "Single-exposure conventional and computed radiography: the hybrid cassette revisited." *Investigative Radiology*. Vol. 28, No. 5, pp. 409-412, May 1993.
6. Du Pont Diagnostic Imaging Systems Document. E. I. du Pont de Nemours and Co. Inc., Photosystems and Electronic Products Department, Wilmington, DE.
7. Personal communication, M. Kobayashi. Fuji Medical Systems Inc., Stamford, CT, Nov. 1993.
8. Andriole KP, Wong STC and Huang HK. "Parallel programming implementation of computed radiography preprocessing tasks." *SCAR'94: 12th Conf. of Computer Applications in Radiology and 8th Conf. on Computer Assisted Radiology*, Winston-Salem, NC, pp. 304-305, June 12-15, 1994.

## **Impact and utilization studies of a PACS display station in an ICU setting**

Katherine P. Andriole, M. Luigia Storto, Gordon Gamsu, and H.K. Bernie Huang

Department of Radiology  
University of California at San Francisco  
San Francisco, CA 94143-0628

### **ABSTRACT**

An assessment of changes in health-care professional behavior as a result of the introduction of a PACS (picture archiving and communication system) display station to an adult medical-surgical intensive care unit (ICU) is investigated via pre- and post-PACS evaluations. ICU display station utilization and the impact on clinical operations are also examined. Parameters measured both pre- and post-PACS ICU display station placement include the number of films per patient day, the number of clinician reviews of a patient's images per day and the percentage of images on which the unit interacts with a radiologist. The elapsed times from the time of exposure to the time of: review by the referring physician, radiologist-unit interaction and clinical action based on image information are also measured. The results of this investigation suggest that the introduction of a PACS display station in the ICU may reduce the number of exams per patient day, decrease the elapsed time from the time of exposure to the time of review by the unit clinician, and improve the time to clinical action. Note, however, that it does not appear to change the percentage of total images on which the unit interacts with a radiologist.

**KEYWORDS:** PACS, display station, technology assessment, intensive care unit (ICU), clinical evaluation

### **1. INTRODUCTION**

Positive impact of a new technology on the clinical environment in which it is introduced is often assumed but difficult to demonstrate. Several groups, most notably the University of Pennsylvania and Duke University, have examined the effect of a digital imaging network on physician behavior in an intensive care unit, time comparison of ICUs with and without digital viewing systems and strategies for experimental design of ICU PACS station technology evaluation studies.<sup>1-9</sup>

The purpose of this investigation was to measure, present and analyze objective data documenting the direct results of placing a PACS display station in an ICU on: physician and radiologist behavior, the time to clinical action and the number of exams performed per patient day. These measures, along with display station utilization studies, elucidate the impact of this enabling PACS technology on patient care and on health-care professional behaviors, leading to conclusions about cost-effectiveness. Additionally, such investigations may speed technology improvement through end-user feedback, redesign and further practical clinical testing.

### **2. MATERIALS AND METHODS**

All unit portable examinations at our institution are performed with computed radiography (CR). Parameters measured both prior to placement of a PACS display workstation in the ICU (pre-PACS) and after placement of the display station in the unit (post-PACS) are summarized in Table 1. They included the number of films per patient day, the number of clinician reviews of a patient's images per day and the percentage of images on which the unit interacted with a radiologist. The elapsed times from the time of exposure to the time of: review by the referring physician, radiologist-unit interaction and clinical action based on image information were also measured.

### MEASURED PARAMETERS

# Films per Patient Day

# Clinician Reviews of a Patient's Images per Day

% Images on which the Unit Interacts with Radiology

### ELAPSED TIME FROM TIME OF EXPOSURE TO TIME OF

Image Review by Non-Radiologist Physician

Radiologist-Unit Interaction

Clinical Action based on image information

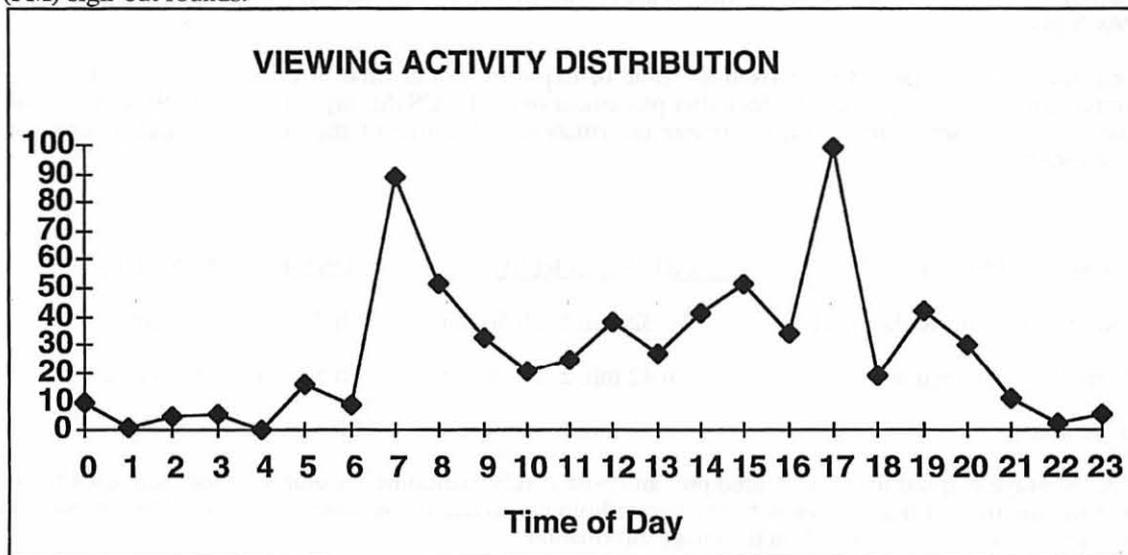
**Table 1.** Parameters and elapsed times measured both pre- and post-PACS ICU display station placement.

Parameter measurements were obtained utilizing a motion activated video camera placed in the chest reading room, aimed at the alternator dedicated entirely to the medical-surgical ICU under study. Other data collected came from review of patients' charts and day sheets in the unit, and unit clinician/nurse daily survey forms. Additionally, in the post-PACS phase, some measurements were obtained via a computer workstation function utilization tracking program.

The medical-surgical ICU display station used in this investigation was an ISG workstation (ISG Technologies, Inc. Ontario, Canada). It consists of two medium resolution (1280 by 1600 line) monitors. Patient demographics and radiology reports (both preliminary and final) appear on the left monitor, as does the most recent image of a selected patient, resident in the unit. Previous or historical images are displayed on the right monitor and can be paged through in full screen resolution or viewed in half-, quarter-, eighth-, etc. minified versions. Display functionality includes interactive window and level (or contrast and brightness) modification, zoom (magnification), pan or scroll, image rotation, and gray scale reversal.

### 3. RESULTS

The graph in Figure 1 depicts typical ICU display station utilization or viewing activity in the adult medical-surgical unit over a 24 hour period. Note that the workstation was accessed at every hour of the day except at 4 AM, with heavy usage during pre-morning (AM) rounds, during morning rounds and during afternoon (PM) sign-out rounds.



**Figure 1.** ICU display station viewing activity distribution over a typical 24 hour period.

The average results for the parameter measurements of the number of films per patient day, the number of clinician reviews of a patient's images per day and the percentage of time the unit interacts with a radiologist are shown in Table 2, for both the pre-PACS phase (prior to the placement of the PACS display station in the ICU) and the post-PACS phase (after placement of the PACS workstation in the ICU).

Note that the number of films per patient day prior to placement of the PACS display station up in the ICU averaged one to two images per day, while after placement of the PACS workstation up in the ICU the average number of films per patient was slightly less. Also, pre-PACS measurements of the number of clinician reviews of a patient's images per day averaged one to three viewings per day, with the post-PACS measurement averaging roughly the same number of viewings of a patient's images per day or slightly increased viewing. Note, however, that the percentage of images on which the unit consulted with a radiologist was roughly the same for the pre- and post-PACS phases.

<u>AVERAGE</u>	<u>PRE-PACS RESULTS</u>	<u>POST-PACS RESULTS</u>
# Films / Patient Day	1.30 ± 0.60	1.09 ± 0.69
# Clinician Reviews of Images / Day	2.18 ± 1.35	2.00 ± 4.80
% Time Unit Interacts with Radiologist	27.4 %	28.9 %

**Table 2.** Average results for parameter measurements pre- and post-PACS, including the number of films per patient day, the number of clinician reviews of a patient's images per day and the percentage of images on which the unit interacts with a radiologist.

The average results for the elapsed times from the time of exposure: to the time of image review by the non-radiologist clinician, to the time of unit-radiology interaction and to the time of clinical action based on information ascertained from a patient's images are shown in Table 3, for both the pre-PACS phase and the post-PACS phase.

Note that the average elapsed time from the time of exposure to the time of image review by the non-radiologist clinician was significantly less after placement of the PACS display station up in the unit, as was the time to clinical action based on the image information. The time of the radiologist-unit interaction remained roughly the same.

<u>AVERAGE ELAPSED TIME</u>	<u>PRE-PACS RESULTS</u>	<u>POST-PACS RESULTS</u>
Image Review by Non-Radiologist	2 h 32 min ± 1 h 50 min	1 h 35 min ± 56 min
Radiologist-Unit Interaction	1 h 42 min ± 2 h 3 min	1 h 56 min ± 1 h 23 min
Clinical Action	3 h 21 min ± 3 h 24 min	2 h 6 min ± 1 h 46 min

**Table 3.** Average elapsed times measured pre- and post-PACS, including the elapsed time from the time of exposure to: the time of image review by the non-radiologist clinician, the time of unit-radiology interaction and the time of clinical action based on the image information.

#### 4. DISCUSSION

In this study to assess the changes in health-care-professional behavior as a result of the introduction of a PACS display station in the medical-surgical intensive care unit, and to assess ICU workstation utilization and the impact on clinical operations, via pre-and post-PACS evaluations, the following may be concluded. Introduction of a PACS display station in the ICU may reduce the number of exams per patient day, decrease the elapsed time from the time of exposure to the time of review by the unit clinician, and improve the time to clinical action. Note, however, that it does not change the percentage of total images on which the unit interacts with a radiologist.

#### 5. ACKNOWLEDGMENTS

The authors would like to thank Jonathan Go of ISG Technologies, Inc., Ontario, Canada, for his contributions to this study. This work was supported in part by a donation from Fuji Medical Systems U.S.A., Inc.

#### 6. REFERENCES

1. Arenson RL, Seshadri SS, Kundel HL, *et al.* "Clinical evaluation of a medical image management system for chest images." *AJR* 150:55-59, January 1988, pp. 55-59.
2. DeSimone DN, Kundel HL, Arenson RL, Seshadri SS, *et al.* "Effect of a digital imaging network on physician behavior in an intensive care unit." *Radiology* 169, October 1988, pp. 41-44.
3. Humphrey LM, Fitzpatrick K, Atallah N, and Ravin CE. "Time comparison of ICUs with and without digital viewing systems." *SPIE, Medical Imaging 1992: Physics of Medical Imaging*, Newport Beach, CA, Vol. 1654, February 1992, pp. 402-407.
4. Kundel HL, Seshadri SS, Carey B, Kishore S, *et al.* "Display station utilization in a PACS serving the medical intensive care unit of the Hospital of the University of Pennsylvania." *SPIE, Medical Imaging 1992: Physics of Medical Imaging*, Newport Beach, CA, Vol. 1654, February 1992, pp. 424-427.
5. Kundel HL, Seshadri SS, Langlotz C, Kishore S, *et al.* "Strategies for the clinical evaluation of PACS: experience with experimental design in a medical intensive care unit." *SCAR'94: 12th Conf. of Computer Applications in Radiology and 8th Conf. on Computer Assisted Radiology*, Winston-Salem, NC, June 12-15, 1994, pp. 614-619.
6. Kundel HL, Seshadri SS, Shile PE, Polansky M, *et al.* "Evaluation of PACS in a medical intensive care unit: the effect of computed radiography." *SPIE, Medical Imaging 1994: Physics of Medical Imaging*, Newport Beach, CA, Vol. 2165, February 1992, pp. 481-487.
7. Horii S, Kishore S, Feingold E, Stevens JF, *et al.* "Intensive care unit workstation usage: digitized film versus phosphor plate imaging." *SPIE, Medical Imaging 1995: Physics of Medical Imaging*, San Diego, CA, Vol. 2435, February 1995, pp. 286-293.
8. Langlotz CP, Cleff B, Even-Shoshan O, Bozzo M, *et al.* "The incremental cost of PACS in a medical intensive care unit." *SPIE, Medical Imaging 1995: Physics of Medical Imaging*, San Diego, CA, Vol. 2435, February 1995, pp. 294-301.
9. Kundel HL, Seshadri SS, Langlotz CP, Horii S, *et al.* "Prospective comparison of the usage of conventional film and PACS based computed radiography for portable chest x-ray imaging in a medical intensive care unit." *SPIE, Medical Imaging 1995: Physics of Medical Imaging*, San Diego, CA, Vol. 2435, February 1995, pp. 302-309.

# Real time multilevel process monitoring and control of CR image acquisition and preprocessing for PACS and ICU

Jianguo Zhang, Stephen Wong, Katherine P. Andriole,  
Albert Wong, and H. K. Huang

Laboratory for Radiological Informatics, Department of Radiology,  
University of California, San Francisco, CA94143-0628

## ABSTRACT

The purpose of this paper is to present a control theory and a fault tolerance algorithm developed for real time monitoring and control of acquisition and preprocessing of computed radiographs for PACS and Intensive Care Unit (ICU) operations. This monitoring and control system uses the event-driven, multilevel processing approach to remove computational bottleneck and to improve system reliability. Its computational performance and processing reliability are evaluated and compared with those of the traditional, single level processing approach.

Key words: Multilevel process structure, monitoring and control, computed radiography, acquisition, preprocessing, PACS and ICU.

## 1. INTRODUCTION

Projection radiography accounts for over 60% of current radiography procedures. As digital imaging becomes more important in radiology practice [1], the need to convert images from projection radiography into digital format becomes apparent. Computed Radiography (CR), as one of the major digital radiography methods, aims to replace screen/film radiographs and is becoming more widely used in medical centers [2]. Raw images generated by CR systems must be processed in many complex computational steps before storing into the centralized archive of PACS (picture archive and communication systems) or transmitting to remote display stations. The robustness of these acquisition, preprocessing steps and the quality of processed images are essential for incorporating CR systems into the clinical PACS framework.

Some processes of CR acquisition and preprocessing run slower than the others, causing computational bottleneck. Subsequently, this delays the transmission of images from the CR system to display stations, like in the ICU for which the quick response to user query is critical. Furthermore, some of these processes manipulate images based on the header information entered by the operators [3]. However, illegal input parameters, software bugs, or system errors often crash the CR acquisition computer. In this paper, we present a control theory and a fault tolerance algorithm to circumvent such problems. The software implementation of the control theory and the algorithm is based on the event-driven, multilevel processing structure. The automated software has been used to provide real time monitoring and control of CR image acquisition and preprocessing in the UCSF PACS and ICU operations. It has been proved to improve processing time, minimize user intervention, and speed up the previously time-consuming quality assurance procedure. The methodology can be extended to streamline and enhance image acquisition and preprocessing of other digital modalities with similar sequential computational structure as CR.

## 2. PROCEDURE OF CR ACQUISITION AND PREPROCESSING

The department of radiology at UCSF operates two CR systems (Fuji 9000 and Fuji AC-2) for all ICUs and portable examinations generating about 1.2 gigabits of image data daily. Fig. 2.1 shows a diagram of the network structure of CR systems, PACS, and ICUs. These large volumes of acquired image data must be processed before they are archived into the PACS and sent to ICU viewing stations through an OC3-c, 155 Mb/s Asynchronous Transfer Mode (ATM) network.

Figure 2.1 lists all computational processes involved in the acquisition and preprocessing of CR. The functions of these routines are as follows: (1) **DASM\_server** process acquires images from a Fuji CR system; (2) **communication** process transfers data between the two Sun workstations; (3) **reformatting** process converts and formats the acquired image data from different sizes to a uniform  $2k \times 2k$  dimension; (4) **background removal** process eliminates redundant background signals in the acquired CR image caused by the collimation effect; (5) **orientation correction** process align and rotate the CR image to a position suitable for viewing in display station; (6) **look\_up\_table (LUT) adjustment**

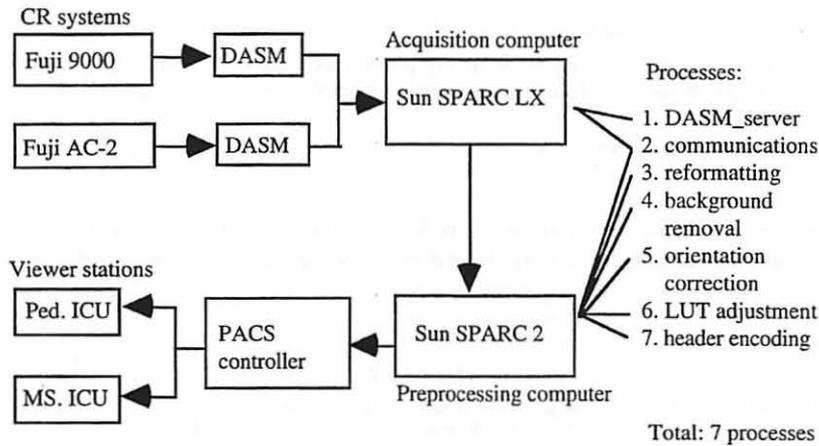


Fig. 2.1 The data flow and processing steps of CR in UCSF PACS and ICUs.

process modifies the default lookup table to enhance the visual quality of aligned images; and **header encoding** process converts the image header from Fuji proprietary format to the standardized ACR-NEMA 2.0 format. Future work will upgrade to DICOM 3.0.

These seven processes are independent demons running in the background of the two SUN SPARC computers. Processing of CR images is sequential through these steps and follows the FIFO model (first in first out). The conventional processing model of these demon procedure is illustrated in Figure 2.2.



Fig. 2.2 A diagram of the stream line procedure of preprocessing of CR images.

Every computational process takes a job from its input queue file, i.e., queue\_in, puts this job into its output queue file, queue\_out, after processing, and then removes that job from the queue\_in. For example, for process i, we denote its queue\_in as queue i and its queue\_out as queue i+1. Our experience indicated that this simple, single level processing structure suffers many drawbacks. First, it creates a bottleneck due to uneven computing speeds and workloads among the computational processes. For the slow computational processes, the jobs may be piled up in their queue\_ins. Meanwhile, faster processes are wasting the CPU cycles when there are no jobs to be done, waiting for new jobs. Second, the system often crashes by certain images with incorrect input parameters, software bugs, and system errors. For example, illegal patient ID may cause reformatting process crash which breaks preprocessing procedure for all following images.

### 3. PROCESS CONTROL THEORY

In this section, we present a control theory to optimize the acquisition and preprocessing time of CR images. The idea of this theory is based on the fact that all the seven demon processes have equal priority in accessing the multitasking

CPU. Suppose that there are  $m$  processes running in a multitasking computer and  $f_0$  is the clock speed, or process frequency,

of this computer. Let us further denote that the process frequency of the  $i$ th CR demon process is  $f_i$  and the number of computer cycles required in finishing one job is  $S_i$ . Every process frequency is assigned by the CPU and the sum of all process frequencies is roughly equal to the CPU frequency or clock speed, so we have the equation described by the following:

$$f_0 \approx f_1 + f_2 + \dots + f_m \quad (3.1)$$

The processing time required by the  $i$ th process is  $S_i/f_i$ . Assuming the demon's processes occupy most of the CPU time, the total processing time to process one image by  $m$  demon processes,  $T$ , can be approximated as follows:

$$T \approx S_1/f_1 + S_2/f_2 + \dots + S_m/f_m \quad (3.2)$$

Our purpose is to attain the minimum  $T$  by adjusting  $f_i$ , for  $i = 1, 2, \dots, m$ . Since  $f_0$  is a constant, Equations (3.1) - (3.2) now become an extreme value problem under certain conditions. With the solution of extreme value under conditions developed in mathematics, we find the  $i$ th frequency assigned by the CPU should be given by the following formula:

$$f_i = \sqrt{S_i} \cdot f_0 \cdot \left( \sum_{j=1}^m \sqrt{S_j} \right)^{-1} \quad (3.3)$$

and the minimum total processing time to process one CR image is given by the following:

$$T = \left( \sum_{i=1}^m \sqrt{S_i} \right)^2 / f_0 \quad (3.4)$$

Fig. 3.1 shows a curve describing the relationship between the processing frequency of the  $i$ th process and its cycles for one job according to Formula (3.3). The CPU assignment of every process frequency not only depends on its circle but also the other process cycles.

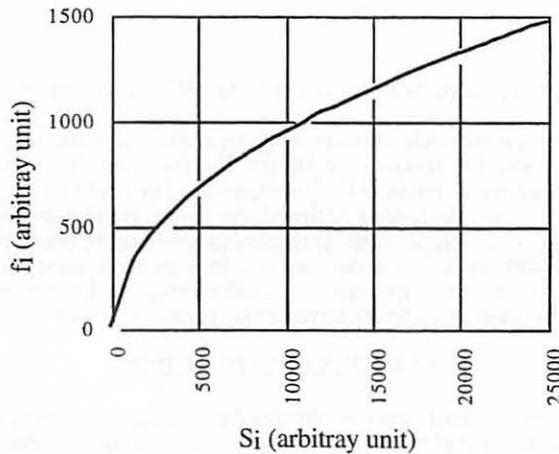


Fig. 3.1 The curve of the  $i$ th process frequency versus its cycle.

#### 4. FAULT TOLERANCE ALGORITHM

Two major causes of unreliability during medical image acquisition are as follows: one is image data with illegal parameters entered by the operators, or software bugs which cause processes crash. The other is system errors which may stop the operation of the CR scanner, although this kind of error happens infrequently. In image processing, a process crash in a job means that it can not finish this job, cannot dequeue this job from queue\_in and get the next job from queue\_in, regardless of how many times it has tried.

Some illegal parameters in the image header can be predicted while some cannot. To recognize the latter errors, an algorithm is developed to guarantee that an image that causes a process crash will not be processed again. Every medical image has its own identification information, such as the serial number and the scanning time. We use these parameters to recognize those images which cause processes to crash. Suppose  $D_i \{d_{1i}, d_{2i}, \dots, d_{ji}\}$  is a vector representing the identification of the  $i$ th medical image and  $d_{ji}$  is its  $j$ th branch vector. We define that  $d_{ji}$  does not correlate with  $d_{jk}$  if  $\langle d_{ji}, d_{jk} \rangle = 0$ , i.e.,  $d_{ji} \neq d_{jk}$ . Thus, if image  $i$  does not correlate with image  $k$ , then their identification vectors have the following relation:

$$\langle D_i, D_k \rangle = 0 \quad (4.1)$$

Otherwise, we say that image  $i$  correlates with image  $k$ .

Now, we apply the correlative recognition algorithm developed above to all the processes in order to identify the images which cause process crashing. After getting a job from a queue\_in, a process must compare the image identification parameters with the previous image according to Equation (4.1). If these two images correlate with each other, it means that this job has been processed more than one time but was not completed. Then, the process removes this job from its queue\_in and gets the next job in the queue. Implementing this algorithm in the multilevel processing structure which we will describe, will guarantee that all jobs are processed no more than one time in every process.

The detection of system errors can be done by periodic checking of the communication states of every process with a monitor process running in another site. If any error occurs in the system which effects the monitor server or parent processes, the central monitor immediately pages the system administrator through a special network, i.e., telephone line.

#### 5. IMPLEMENTATION OF CONTROL THEORY

From Figure 3.1, every process frequency assigned optimally by the CPU depends on not only its cycles but also all the other process cycles. Sometimes, the cycles of certain processes are changed due to images of different sizes. This requires the process frequencies to be assigned dynamically. Usually, it is very difficult to have exact and direct control on the process frequency through software. In the following, we present an implementation of the process control theory using the job control mechanism and the client/server model.

Let the number of jobs done by the  $i$ th process be  $n_i$  and the total time taken be  $\Delta T$ , the unit time spent for one job by the  $i$ th process, then is  $\Delta T/n_i$ . This should be approximately equal to  $S_i/f_i$ . That is:

$$\Delta T/n_i = S_i/f_i \quad (5.1)$$

Substituting Equation (3.3) into (5.1), we have:

$$n_i = \Delta T \cdot f_0 / \left[ \sqrt{S_i} \cdot \left( \sum_{j=1}^m \sqrt{S_j} \right) \right] \quad (5.2)$$

Since  $S_j$  ( $j = 1, 2, \dots, m$ ) can be pre-determined,  $f_0$  is a constant and  $\Delta T$  is same for all processes, we can attain the statistical minimum of the total processing time required by all processes to handle a CR image. This can be done by controlling the number of jobs in every process specified by Equation (5.2). Controlling the numbers of jobs for every process can be implemented by the client/server model shown in Figure 5.1.

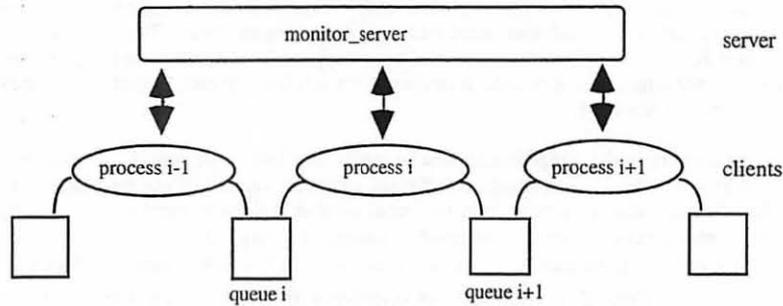


Fig. 5.1 Client/Server model of job control for CR acquisition and preprocessing.

In the client/server model shown in Fig. 5.1, a server process named **monitor\_server** controls the numbers of jobs of all client processes. All clients have to make requests to the monitor server after they get a job from their queue\_ins. The monitor server returns a message according to the parameter calculated from Equation (5.2) in order to notify the clients whether to process the job. The clients begin to perform functions only they get the permission from the monitor server; otherwise, they turn to the sleep state and repoll the server at regular intervals, e.g., 2-3 seconds.

#### 6. IMPLEMENTATION OF FAULT TOLERANCE

To implement the fault tolerance algorithm developed in Section 4 in all seven CR processes, we further divide every process into a parent process and a child process. As shown in Figure 6.1, (a) shows the original process, (b) the child and the parent processes, and (c) the functional diagram of a child process.

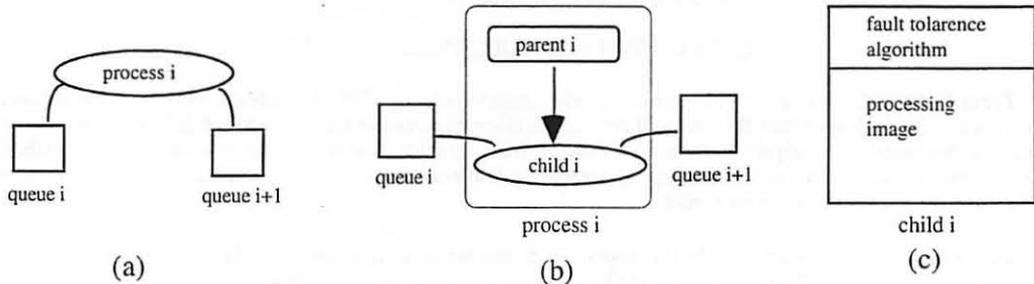


Fig. 6.1 Diagrams of child and parent processes and child process functions.

The major functions for parent processes are: (1) check their queue\_ins to detect if there is a job available; (2) generate a child process to perform desired functions; (3) monitor the processing time of the child process generated. If the child takes too long time to process a job, the parent will kill the child process. Because all these functions have nothing to do with the image data, the parent processes would not be effected by any illegal parameters or software bugs.

Every child process is composed of two parts as shown in Fig. 6.1(c). The first is the fault tolerance algorithm developed in Section 4, and the second is its image processing functions. The functions for a child are as follows: first, after being generated by its parent, it gets a job from its queue\_in and applies the fault tolerance algorithm to this image. If it fails the fault tolerance criterion, i.e., the image has been processed before but not completed, the child process removes this job from its queue\_in with an error message and then exits. Second, the child process reserves the image identification vector parameters in a special file in order to correlate with the next job. Third, it performs its assigned image processing function. Fourth, after completed the processing, the child moves the job into its queue\_out. Finally, it deletes the job from its queue\_in with the success message and then exits.

The merit of this child/parent architecture is that any error causing the crashing of a child process will not affect the parent process from continuing, such that the whole process still working.

7. EVENT-DRIVEN MULTILEVEL PROCESS STRUCTURE

To make the job control mechanism and the fault tolerance algorithm work together, an event-driven, multilevel processing structure is developed as shown in Figure 7.1. There are three process levels: monitor\_server, parent, and child. The major tasks and functions of these three levels have been described in Section 5 and 6. In what follows, we describe their additional functions within the event-driven software structure.

For the monitor server, apart from controlling the number of jobs of each process, it also monitors parent processes and the system environment. For example, it makes sure whether all parents are alive and whether there is enough memory space for parents to generate children. It receives events from child processes and transfers error messages to the central monitor in order to notify the system administrator.

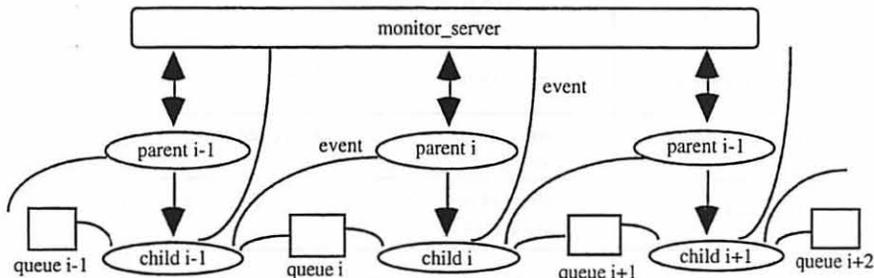


Fig. 7.1 The diagram of event-driven multilevel process structure.

It also accepts communication checking events sent from the central monitor. In other words, the monitor server is a server for all parents and children, but a client of central monitor. Figure 7.2 gives the communication relationship between the central monitor and the monitor server, as well as the parents.

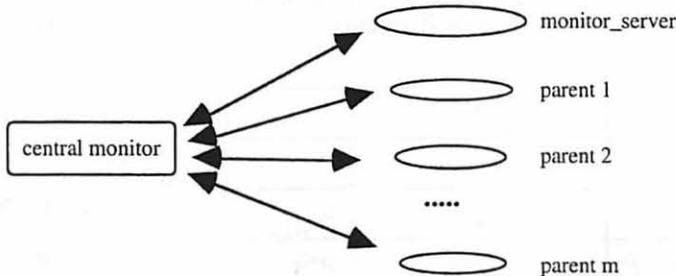


Fig. 7.2 The diagram of the communication relationship between the central monitor, the monitor server, and the parents.

For the parents, new functions are: receive trigger events sent from the preceding children, receive communication checking events sent from the central monitor. For children, the new functions are: send commands such as "job start", "job end," and "trigger" as well as any abnormal events to the monitor server and next parents. Table 7.1 summarizes the client/server relationships of various process types in the event-driven, multilevel processing structure.

Now, we describe how this event-driven, multilevel processing structure works. First, for the  $i$ th process, after receiving a trigger event sent from its preceding child, the parent first sends a request event to the monitor server which will decide whether this job can be done according to the control theory. It then generates a child process to perform the requested functions if it receives a permission event from the monitor server. Second, the generated child process gets a job from its queue\_in, applies the fault tolerance algorithm to this image to make sure it is a new job and then sends "job start" event to the monitor server. Third, the child process continues to perform image processing functions and sends the abnormal events to the monitor server if any error happens in this procedure. Fourth, having successfully completed the job, this child puts the processed job into its queue\_out, dequeues this job from its queue\_in with success message, sends the "job end" and triggered events to the monitor server and next parent ( $i+1$ ) respectively, and then exits. Otherwise, if it fails the fault tolerance algorithm, it dequeues this job from its queue\_in with an error message and sends the error event to the monitor server before exiting.

Table 7.1. Summary of Client/Server relationship in the event-driven, multilevel process structure.

Process Type	Client/Server relations
Central Monitor	server of monitor server, clients of parents
Monitor Server	server of parents and children, client of central monitor
Parent Processes	client of monitor server, server of central monitor and children
Child Processes	client of monitor server

## 8. PRELIMINARY RESULTS

We have implemented this event-driven multilevel process structure for monitoring and control of CR image acquisition and preprocessing in the hospital integrated PACS and ICUs at UCSF. All inter-process communications use TCP/IP protocols. To measure the performance of processing speed between the single level processing structure without control processes and multilevel process structure with control mechanisms, twenty CR images were used. Fig. 8.1 shows the curves of total processing time for each of the 20 images by using the single level structure and the multilevel structure.

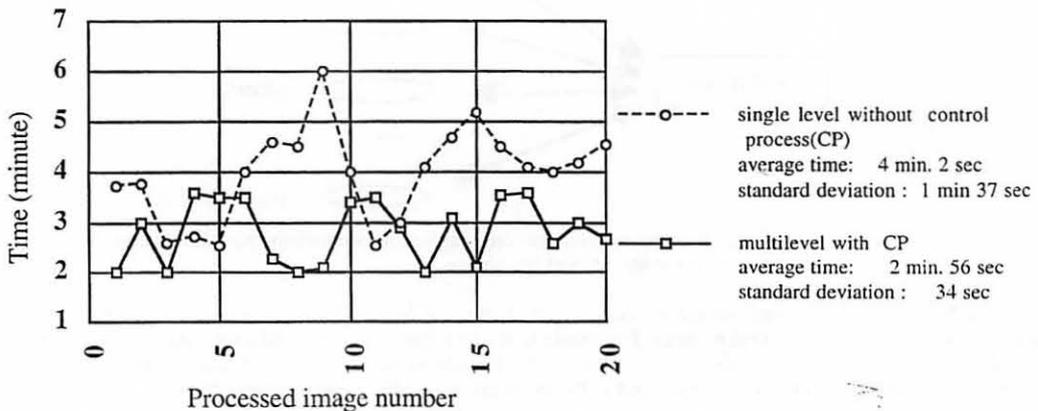


Fig. 8.1 The curves of total processing time for one image v.s. image number.

We see the average processing time for one CR image required by all processes with the single level process structure is 4 min. 2 sec., standard deviation is 1 min. 37 seconds. The average time for one CR image with event-driven, multilevel processing structure is 2 min. 56 sec., standard deviation is 34 sec.. Hence, the processing time for one image with the controlled, multilevel structure saves almost 25% compared with the traditional single level structure. We monitored the CR acquisition and preprocessing operations for over three months, 6514 CR images were acquired and processed using the new multilevel processing method with control processes. The total of 134 errors were detected automatically and corrected. No process crash was recorded and no image data was lost.

## 9. CONCLUSIONS

With the implementation of the event-driven, multilevel processing structure for real time monitoring and control of CR image acquisition and preprocessing, the processing speed is increased and the downtime is reduced. In addition, the reliability and processed image visual quality are improved greatly while the quality assurance time and user intervention are decreased. This event-driven, multilevel processing method can also be extended to other PACS image preprocessing applications.

## 10. ACKNOWLEDGMENT

This research was supported in part by National Natural Science Foundation of China.

## 11. REFERENCES

1. H. K. Huang, et al, "Multimedia in the radiology environment: current concept", J. Med. Imaging & Graphics, Vol.18, pp1-10, 1994.
2. T. L. Tokano eds, Computed Radiography, Springer-Verlag Press, New York, 1987.
3. M. F. McNitt-Gray, et al, "Image preprocessing for picture archiving and communication system", Investigative Radiology, Vol.27, pp529-535, 1992.

# Empower PACS with Content-Based Queries and 3-D Image Visualization

<sup>1,2</sup>Stephen T.C. Wong, <sup>2</sup>Kent Soo Hoo, Jr., <sup>1,2</sup>H.K. Huang

<sup>1</sup>Department of Radiology, Univ. of California, San Francisco, CA 94143-0628

<sup>2</sup>Bioengineering Graduate Program, Univ. of California, San Francisco/Berkeley, CA 94143-0640

## ABSTRACT

Current generation of picture archiving and communication systems (PACS) lacks the capabilities to permit content-based searches to be made on image data and to visualize and render three dimensional (3D) image data in a cost effective manner. The purpose of this research project is to investigate a framework that will combine the storage and communication components of PACS with the power of content-based image indexing and 3D visualization. This presentation will describe the integrated architecture and tools of our experimental system with examples taken from applications of neurological surgical planning and assessment of pediatric bone age.

## 1. INTRODUCTION

In the last decade, Picture Archiving and Communication System (PACS) has been the prevalent means for acquisition, storage, communication, and display of medical images [Huang 91]. A large scale PACS consists of many components: radiological imaging systems, textual data sources, data acquisition computers, data storage systems, and high resolution display stations interconnected by computer networks. The basic functions of a PACS fall under three major headings, data input, data management, and data output (see Figure 1), as in many information management systems.

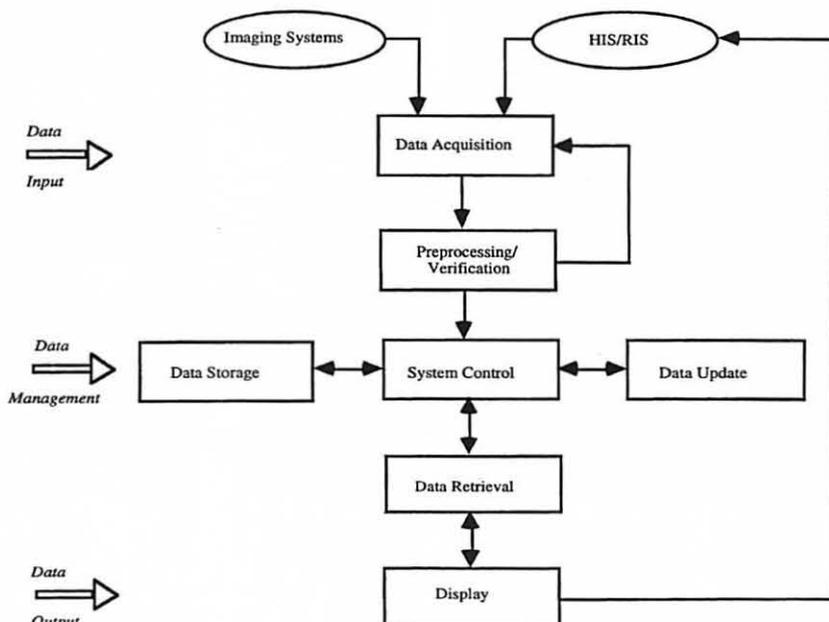


Figure 1. The basic functional components of PACS.

Referring to Figure 1, data input involves the automatic acquisition of image and data from various medical imaging devices and clinical information systems. Preprocessing functions, such as automatic removal of background in Computed Radiography (CR), are performed on the acquired data. The preprocessed and verified image and data are then stored into the centralized data storage archive. The multimedia data is viewed on the client high resolution display stations upon user requests. Sometimes, the user can also use the display station to edit or create new reports for updating patient case information in the PACS archive.

Figure 2 shows the hospital-integrated picture archiving and communication system (HI-PACS) implemented at UCSF which takes the client-server approach. The system contains a centralized, massive storage of multimedia clinical data acquired from imaging scanners and databases across the UC hospitals in the San Francisco Bay Area.

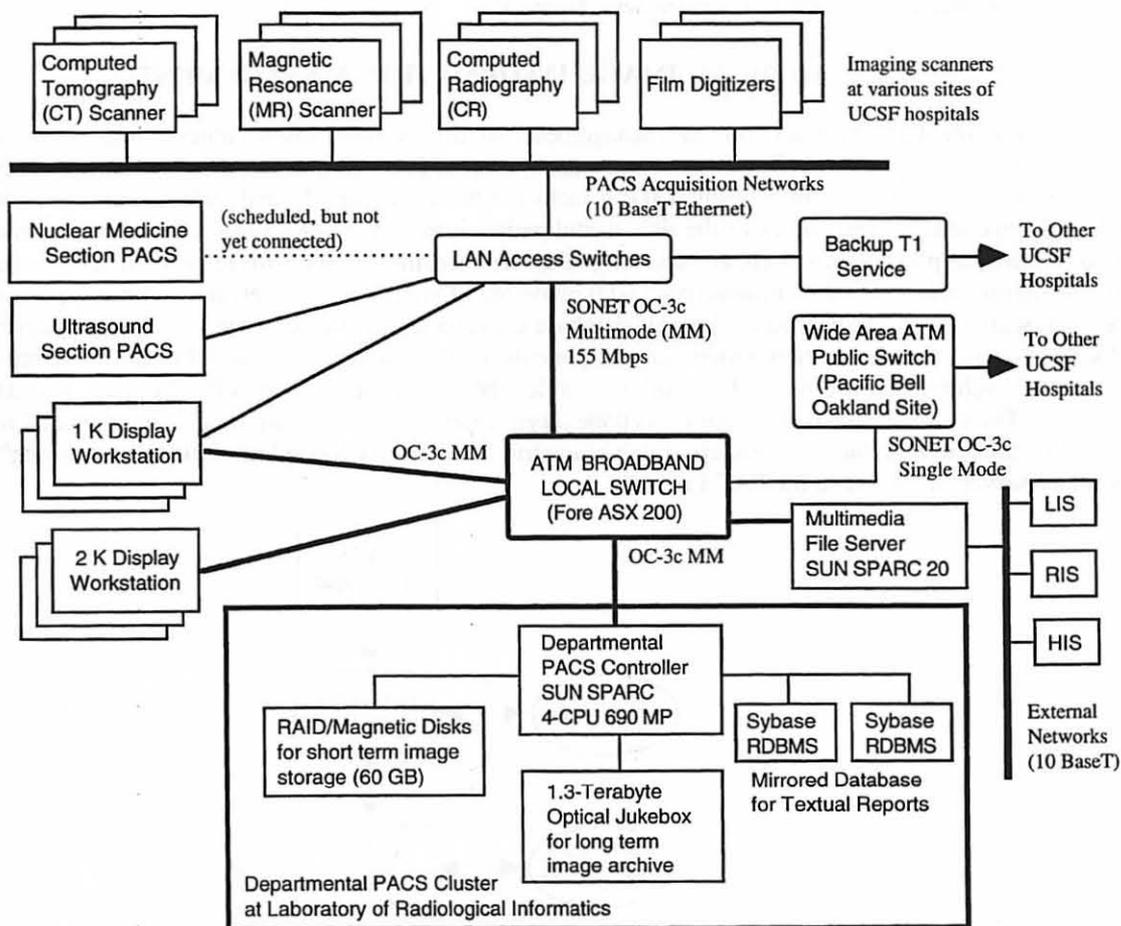


Figure 2. Schematic diagram of the hospital integrated PACS at UCSF.

The implementation of this HI-PACS emphasizes recognized and implemented standards, open systems connectivity, hierarchical memory management, database integration, and security [Wong 96]. The UCSF PACS connects not only various imaging scanners directly, but also smaller PACS of subspecialty sections and textual information sources, such as hospital information system (HIS), radiology information system (RIS), and library information system (LIS). The PACS fiber-optic backbone has been installed interconnecting four major facilities in UCSF, i.e., Moffitt and Long hospitals, ambulatory care clinic, and campus library, to provide a data warehouse integrating information from multiple imaging and data sources.

Wide area connections between UCSF and affiliated hospitals in the San Francisco Bay Area using asynchronous transmission mode (ATM - OC-3c, 155 Megabits per second (Mbps) for real-time teleradiology) are being implemented. Image communications between UCSF hospitals using frame relay technology (T1 link - 1.54 Mbps) for non real time telemanagement have been in operation for about two years. The data archive at the departmental PACS cluster can store up to 1.3 terabytes of image and text data. These PACS data can be retrieved and transmitted over the networks to high resolution, multi-monitor display stations located in intensive care units and radiology reading rooms for clinical review [Huang 96].

## 2. MEDICAL IMAGE INFORMATION MANAGEMENT

The intent of PACS is to provide management and fast review of vast volumes of image and text files in a digital radiology department. The file management system of PACS, however, can only handle query by a few keys, such as patient name or a hospital ID. It lacks the means to classify and index image data files by their information contents. Thus much of the rich, useful patient information stored in PACS has not been utilized for broader medical practice, research, and teaching. Furthermore, most of the current generation of display stations do not support real-time three dimensional (3-D) rendering of volumetric datasets due to limited processor power and hardware costs. To incorporate flexible database management and volumetric visualization capabilities into PACS, we need to add two new functional components, namely, data index and 3D data visualization, into the functional architecture of PACS. Data index provides the capabilities to search PACS image and data files by features. The extraction of image features is done *a priori* through semi-automatic segmentation and registration. Data visualization enables the rendering of volumetric data on PACS display workstations through powerful graphics servers available in the PACS networks.

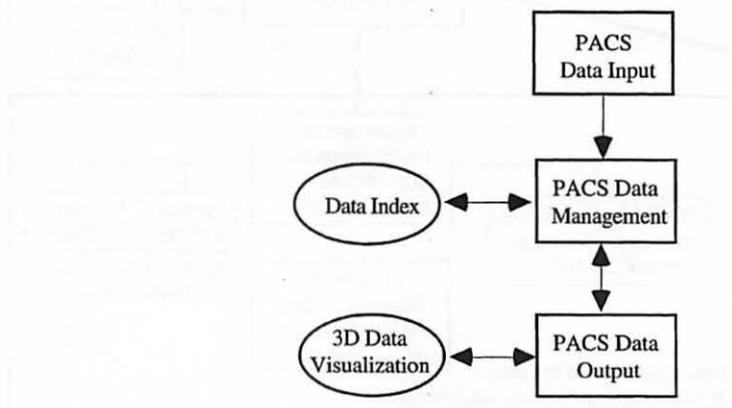


Figure 3. The addition of data index and 3D visualization functions in PACS.

To incorporate these new functions into PACS, we are implementing an image database server (IDBS) and a 3D visualization server in the UCSF PACS. The IDBS contains interactive tools that can be used to extract relevant features from PACS images and textual reports and to organize the extracted data, together with annotated image slices, into a common data model. The IDBS can also process various computer-aided medical imaging applications, such as fast indexing of medical records and activate rules for decision making. It consists of a set of sharable, multitasking components that interact with client workstations, PACS, and the image visualization server (IVS). The current IDBS is developed on a 4-processor SUN 690 MP SPARC computer (Sun Microsystems, Mountain View, CA) running Solaris Operating System (OS) 2.4. This database server contains three computational engines for providing database, knowledge base, and image processing services to remote client workstations on demands.

The image visualization server is being implemented on a high power SGI Onyx computer (Silicon Graphics Inc., Mountain View, CA) with up to 4 reality engines and intends to serve as a central node for computational intensive graphics operations. The connection of the Onyx computer will be ATM-based (currently is a 10 BaseT link) to reduce transmission time of large image data [ATM 93,93a].

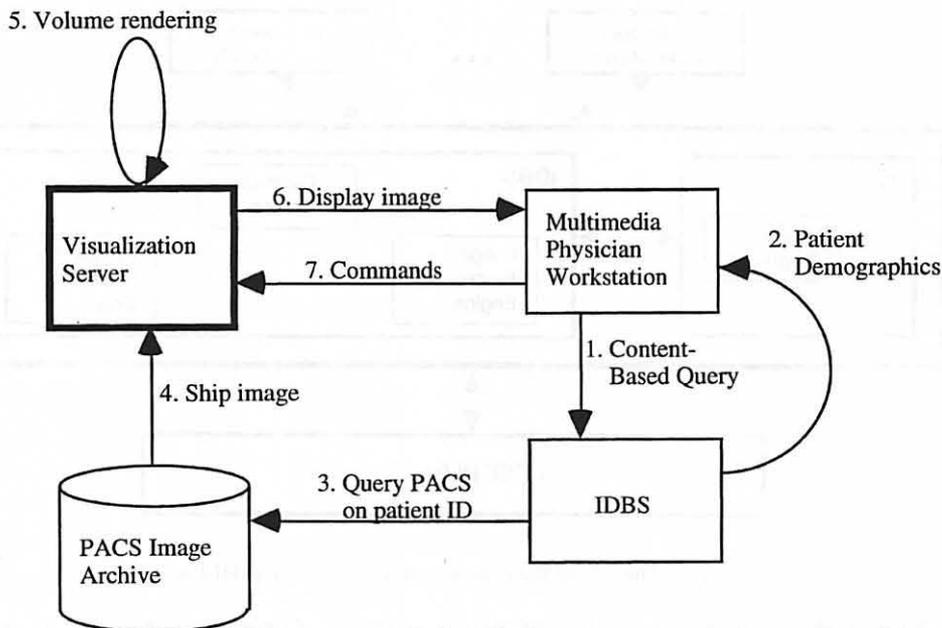


Figure 4: Application Scenario: Content-Based Retrieval and Volumetric Visualization of Clinical Images

Figure 4 considers an application scenario of indexing and visualizing PACS images. First, a neurosurgeon would query the IDBS for specific sets of patient images based on image content, such as tumor volume and type. The surgeon can also index the database server using disease symptoms and other keywords. The IDBS returns the retrieved patient demographics and pictures of interest to the physician. For raw image data not stored in the IDBS, the physician can retrieve them from the PACS's nearline optical jukebox using the patient name or hospital ID obtained from the IDBS. The volumetric image thus retrieved from PACS are represented in pixels. For better visualization of the tumor location and type, the image dataset is automatically routed to the IVS

where the images are texture mapped into the texel space [Foley 93] and then volume rendered. The 3-D rendered images are then sent to the physician's station for display. In addition, the surgeon can issue commands, such as rotate, translate, and reslice, to the IVS and the changes are updated in the local screen in real time.

### 3. CONTENT BASED IMAGE QUERIES

The IDBS is a centrally located core of application programs for accessing, processing, managing, and manipulating data from distributed information sources in the HI-PACS environment as if they were from a single database. This middleware database server, resided in between HI-PACS and client medical workstations, consists of three cooperating processes or engines, running in parallel. It also serves to filter all client requests through various, centrally controlled levels of security, such as IP (Internet Protocol) firewalls, access control, and image authentication. Figure 5 shows the key components in our medical image database framework.

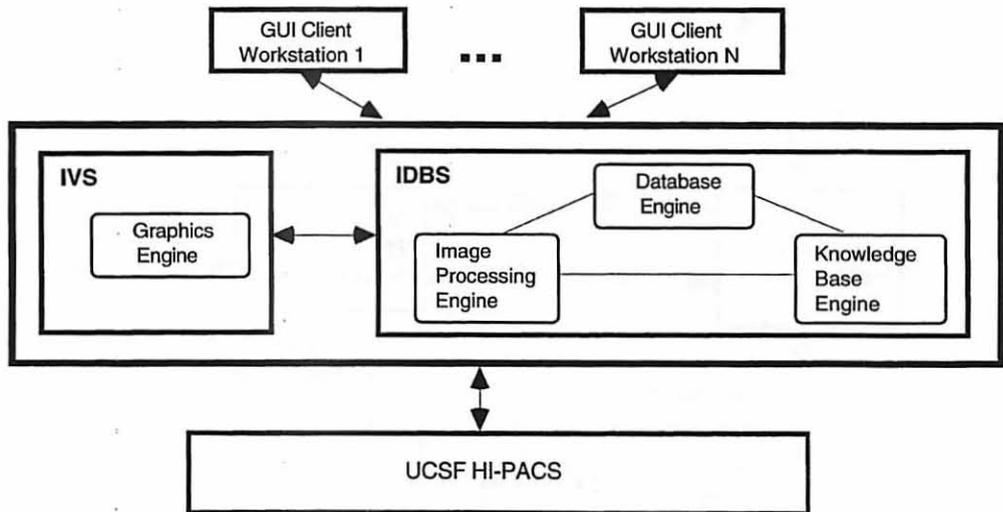


Figure 5. The new architecture with IDBS and VS over HI-PACS.

(1) **Database Engine** that contains a collection of objects and a set of operations that they can perform. Data models are devised to classify, organize, and represent classes of different images and text to support information retrieval by image content for various medical applications. The database engine treats system components and multimedia data of the entire database framework as a logical conglomeration of distributed, interacting objects of various levels of granularity.

For rapid prototyping purpose, we decided to use a commercially available database management system and then to develop customized software modules on top of the commercial system for special applications. The current database engine used is based on the object relational technology – Sybase database management system (Sybase Inc., Emeryville, CA) for the underlying persistent data storage. Future plan is to upgrade to

object-relational technology for more expressive data representation while maintaining the compatibility to existing relational systems.

- (2) **Image Processing Engine** that provides image processing services for various medical applications. These services include computationally intensive operations, such as image registration and feature extraction. The decision to locate heavy duty imaging application routines in the powerful image database server is to reduce the hardware and software requirements on the client workstations. The medical image processing and visualization software packages used include Volumetric Image Display and Analysis (VIDA) (University of Iowa College of Medicine, Iowa City, IA), Interactive Data Language (IDL) (Research Systems, Inc., Boulder Co) and several in-house developed programs [Wong 95a].
- (3) **Knowledge Base Engine** that consists of three kinds of knowledge base modules. Rule-based modules that encode general and specific medical rules for decision making, case-based modules that aid the navigation of database through similar or reference imaging cases, and model-based modules that provide interactive work-up models of different disease categories closely involving physicians making joint decisions. The knowledge base engine being developed is a mix of Prolog (Quintus Prolog, Mountain View, CA) and C programs.
- (4) **GUI Client Workstations** each provides a graphical user interface (GUI) for the local user to request applications of the image database server to perform information retrieval, manipulation, analysis, and update. Each client workstation also contains simple image processing capabilities, such as thresholding and filtering, but accesses the IDBS for more sophisticated database services. Our current application development tools include Gain Momentum, Motif, Web browsers, and several image processing library developer's kit.
- (5) **HI-PACS** consists of the PACS multimedia file storehouse and a federation of distributed, heterogeneous medical databases, imaging acquisition devices, and hospital information systems, providing the environment in which the above system components run.

#### 4. FEATURE EXTRACTION STEPS

Figure 6 of the following page illustrates the operational steps that the IDBS takes to extract image and textual features from the on line PACS data. The feature extraction is based on the *a priori* approach, rather than the dynamic and automatic feature extraction during user query as proposed in many non-medical image database applications [Guidivada 95, Flickner 95]. Image registration permits the combination of different types of functional (such as PET and SPECT images) and structural information (such as MRI images), setting the stage for feature extraction. We use the automatic Ratio Image Uniformity algorithms developed by Woods for registration of neuroimages including PET, MRI, MRS, and MEG [Wood 93, Wong 95a]. The correlated image datasets are encoded into the targeted data model to serve as definitive indexing in image query. For example, registering the functional images of PET with the MRI images of the same patient, allows the intrinsically better spatial resolution of MR images to be used in quantitatively analyzing functional information (metabolic count of glucose consumption) of captured PET scans.

The segmentation and extraction of medical images are done interactively using the Volumetric Image Display and Analysis (VIDA) software [Hoffman 94]. We divide the image features into primitive and logical. Primitive features are directly obtained from the medical images, such as volume, shape, and texture of certain organs in CT images or metabolic activities of brain tissue in PET scans. Regions of interest are semi-automatically outlined in the MRI images to obtain anatomic volumes and cross-registration with functional images

allows quantitative analysis of functionality in the corresponding anatomic region. Logical features are abstract representations of images at various levels of detail and deeper domain semantics. For example, whether the volume of an anatomic structure is normal or whether certain brain tissue is hypometabolic in reference to certain established data. These logical features are synthesized from primitive ones and additional domain knowledge. All features extracted are entered into appropriate attributes of the objects defined in a data model to facilitate subsequent information query by content. In epilepsy, we have extracted MRI anatomical volume, PET glucose uptake count, MRS spectra, and MEG dipole polarization for the amygdala and hippocampus.

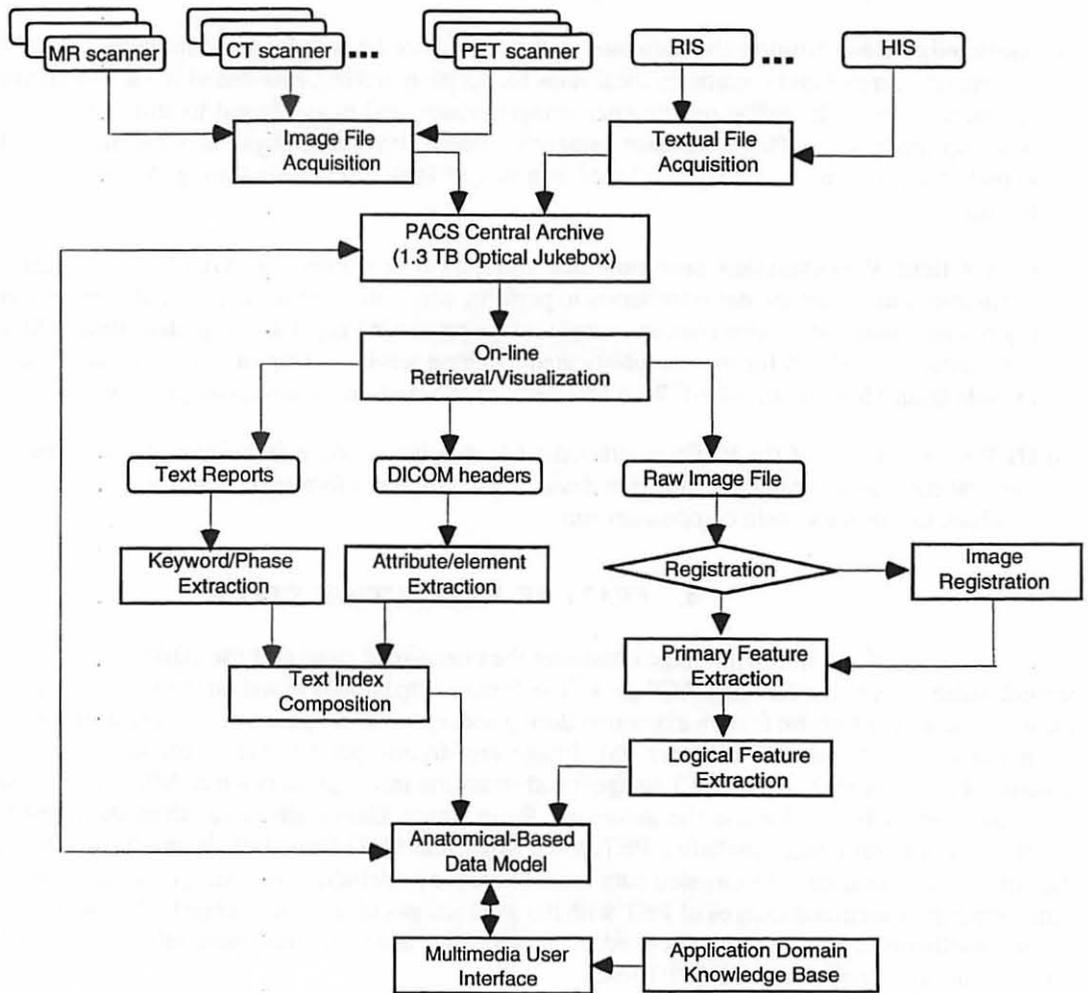


Figure 6. The operational flow of extracting image and text features into an anatomical based data model for subsequent content-based image indexing in IDBS. Specific knowledge or heuristics is triggered to aid the query and navigation of the medical image database.

The extraction and composition of textual data from the diagnostic reports and DICOM (Digital Imaging and Communications in Medicine) header fields can be automatic. Certain keywords or phrases are automatically extracted from the physician textual reports for indexing purposes. UCSF HI-PACS image files all have the DICOM headers which contain patient and imaging exam information. The DICOM header is organized into sequential data element tags, each consisting of a group number and element number. For example, the value of Patient's Name is located in group 0010, element 0010 (see Table 1). This value is automatically extracted and entered into the IDBS column Patient's Name.

Group	Element	Name
0010	0010	Patient's Name
0010	0020	Patient ID
0018	0015	Body Part Examined

Table 1: Data Element Tags from the DICOM header

A multimedia physician workstation prototype was developed using an object-oriented multimedia GUI builder. Thus, user friendly interfaces for retrieving, visualizing, and analyzing PACS images, can be rapidly developed for specific domain applications. We have developed databases and interfaces for two clinical applications, namely, age assessment of pediatric bone images and presurgical planning in epilepsy.

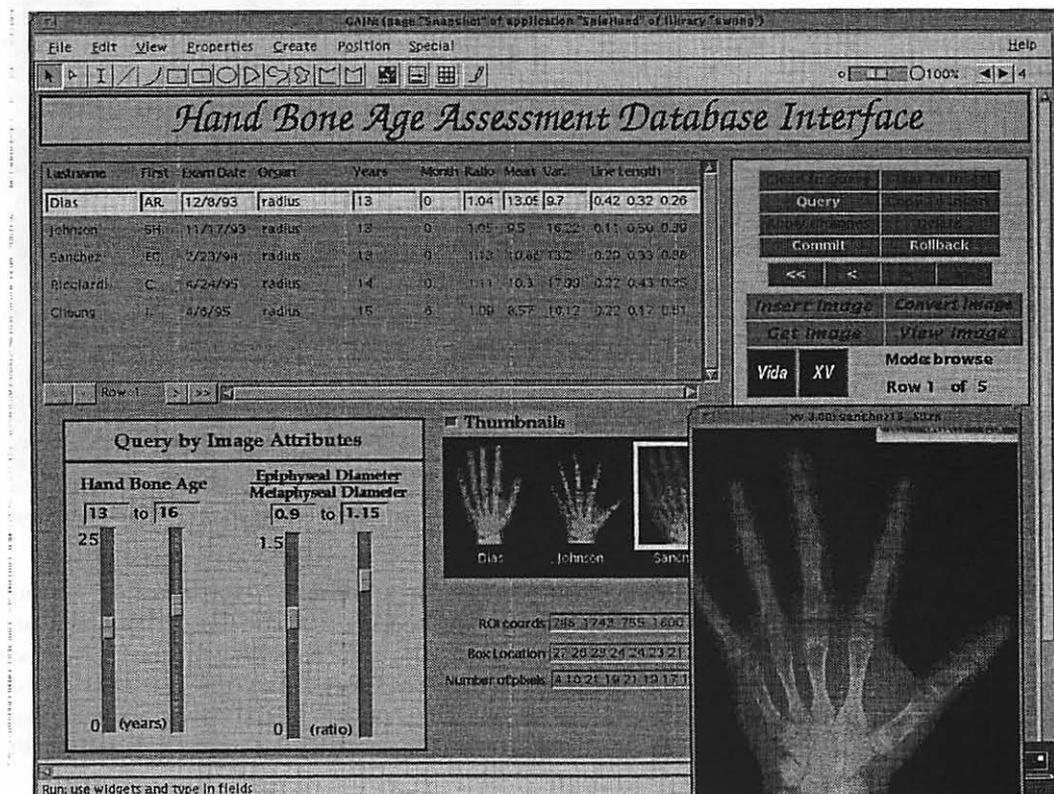


Figure 7. Content-based retrieval for hand bone imaging.

Based on a radiological examination of the skeletal development of a left hand wrists, the bone age is assessed and compared with the chronological age of the patient. A discrepancy indicates abnormalities in skeletal development. Query of the IDBS for pediatric hand bone images can be by image content, e.g., radius bone age or ratio of epiphyseal and metaphyseal diameters, by patient attributes, e.g., name, age, and exam\_date, or a combination of these features. Programs for extracting features of hand bone images were discussed in [Pietka 91,93]. The sliders in the “Query by Image Attributes” window can be used to specify the range of the image attributes for data retrieval. The IDBS returns with a list of five patients and representative thumbnail images satisfying the combined image and patient attribute constraints. The user can click on any thumbnail image to retrieve, visualize, and analyze the original digitized hand radiographs.

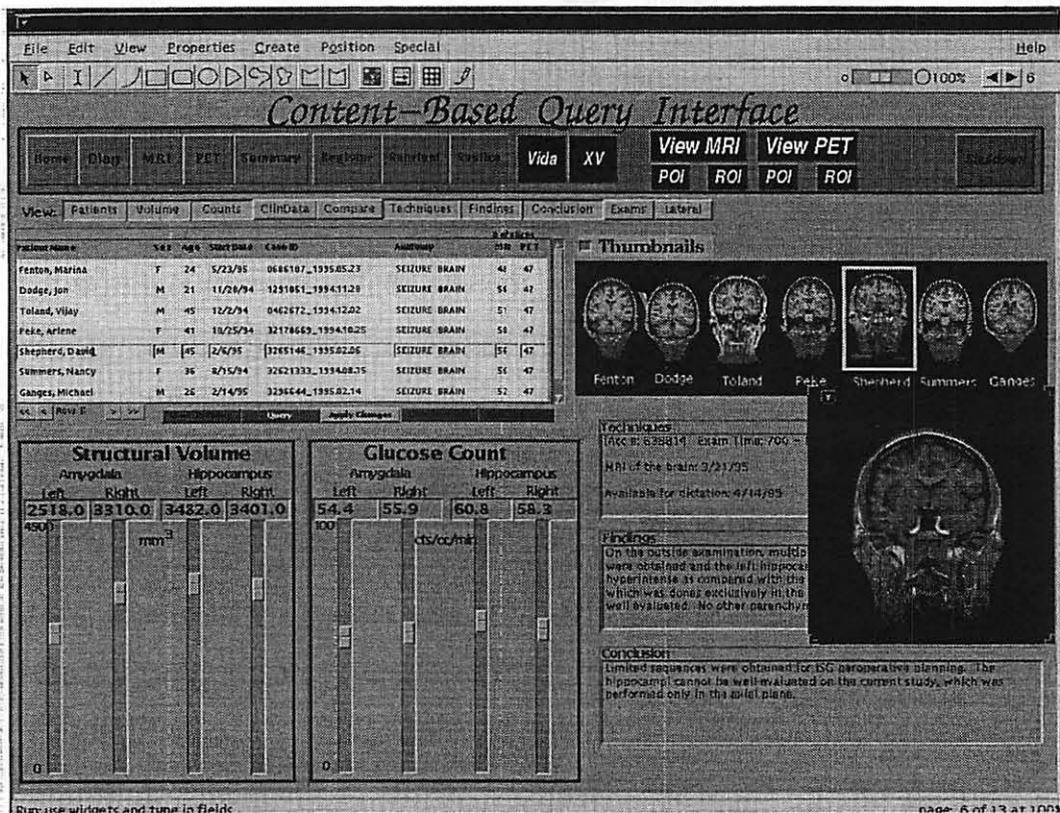


Figure 8. Content-based retrieval of MRI images based on ranges of structure/function attributes.

In Figure 8, we illustrate the use of content-based information query for assisting the presurgical evaluation of complex partial seizures. In this case, the user first specifies the structural, functional, and textual attributes of the MRI studies of interest. The IDBS returns a list of patients satisfying the query constraints and a set of representative pictures of interest of the image sets in thumbnail form. The user then clicks on one of the thumbnail images to zoom it up to the full size or to retrieve the complete 3-D MRI dataset for further study. After studying the retrieved images, the user can update the database with new pictures of interest, regions of interest, and image attributes, and textual reports.

## 5. 3-D IMAGE DATA VISUALIZATION

Although the digital imaging modalities such as MRI, PET, and CT capture 3-D structural and functional images, they are reviewed by radiologists as 2-D slices. The rendering of large 3-D datasets requires huge amounts of computation which has until now limited the use 3-D visualization in clinical environments. Using the specialized graphics hardware of SGIs' Onyx Reality Engine architecture as the IVS and the high speed broadband communication network, such as 155 Mbps ATM, it is possible for using a low end, economic workstation to visualize and manipulate 3-D datasets in real-time.

The basic ideal is to map a stack of 2-D images slices, e.g., 100 MRI images slices of the whole brain, into a 3-D textual mapping space and then volumetrically render the images; using the protocol similar to the one discussed in Figure 4. In the texel space, many efficient and sophisticated manipulation tools can be applied to enhance the visualization of medical images. Parts of the brain can be made transparent or removed to view internal structures while arbitrary cutting planes can be composed in real-time. Multiple 3-D image sets can be viewed concurrently, registered, assigned separate texture lookup tables, and then merged together. For example, a PET study and MRI study so merged allow the determination of geometric concordance between abnormal functional and structural features. In the near future, we will remotely display the 3-D image on client multimedia workstations. The client image manipulation commands will be sent to the IVS, instead of sending the entire new 3-D image back and forth between the visualization server and display client. This will reduce response time and network traffic.

Figure 9 shows an example of volumetric brain image mapped from the pixel space to the texel space using OpenGL programming language. The highlighted box indicates the frontal lobes of the brain. This example was performed using the Onyx IVS and an SGI Indy desktop as a client display station within the UCSF PACS networks.

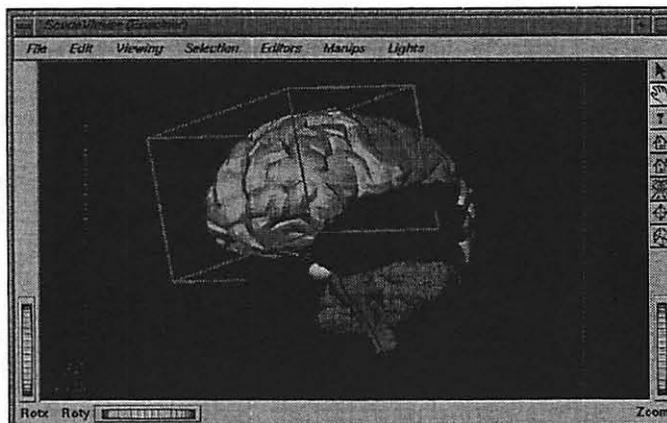


Figure 9. A volumetric brain image rendered in OpenGL with the frontal lobe highlighted by the rectangle b (Courtesy of Dr. E. Grant, UCSF and Dr. P. Cahoon, Univ. British Columbia).

## SUMMARY

In this presentation, we have described the value-added PACS architecture for supporting content-based image retrieval in limited applications and for visualization of large volumetric brain image data. This architecture

is based on distributed client/server concept. It consists of three major servers: the PACS server for traditional image display and archive functions, the image database server for indexing PACS patient cases by features and keywords, and the image visualization server which acts as a graphical node of the PACS for 3D rendering and display on inexpensive PACS workstation. Work is also in progress to extend this concept for World Wide Web access [Wong 96a].

## REFERENCES

- [ATM 93a] *ATM User-Network Interface Specification, Ver 3.0*, September, ATM Forum, 1993.
- [ATM 93b] *The ATM Forum BISDN Inter Carrier Interface (B-ICI) Specification, Ver 1.0*, ATM Forum, August, 1993
- [Flickner 95] M Flickner, H Sawhney, W Niblack, J Ashley, Q Huang, B Dom, M Gorkani, J Hafner, D Lee, D Petkovic, D Steele, and P Yanker. Query by image and video content: The QBIC system, *IEEE Computer*, September, 1995;28(9):23-32.
- [Foley 93] JD Foley, A van Dam, SK Feiner, JF Hughes. *Computer graphics: Principles and practice*, 2nd Edition, Addison Wesley, NY, NY, 1993.
- [Guidivada 95] V Guidivada and V Raghavan. Content-Based Image Retrieval Systems. *IEEE Computer*, Sept., 1995:18-22.
- [Hoffman 94] E. Hoffman, *VIDA (volumetric image display and analysis) operational manual*, Dept. of Radiology, University of Iowa College of Medicine, 1994.
- [Huang 91] HK Huang, O Ratib, AR Bakker, G Witte, eds. *NATO ASI Series F, PACS in Medicine, Vol. 74*. New York: Springer-Verlag, 1991:213-230.
- [Huang 96] Current Status of the UCSF Second Generation PACS. *SPIE Medical Imaging 96: PACS*, Newport Beach, CA, 1996.
- [Pietka 91] E Pietka, F McNitt-Gray, ML Kuo, HK Huang, Computer-Assisted Phalangeal Analysis in Skeletal Age Assessment, *IEEE Trans. on Medical Imaging*, Vol. 10, No. 4, 1991:616-620.
- [Pietka 93] E Pietka, L Kaabi, ML Kuo, HK Huang, Feature Extraction in Carpal-Bone Analysis, *IEEE Trans. on Medical Imaging*, Vol. 12, No.1, 1993:44-49.
- [Woods 93] RP Woods, JC Mazziotta, and SR Cherry. MRI-PET registration with automated algorithm. *J Computer Assisted Tomography* 1993; 17(4):536-546
- [Wong 95a] STC Wong, R Knowlton, M Chew, HK Huang. Integrating Multidimensional Imaging, Multimodality Registration, and Multimedia Database for Epilepsy Diagnosis. *SPIE Medical Imaging 95*, 1995.
- [Wong 96] STC Wong, HK Huang. A hospital-integrated image base management system. *IEEE Trans. Systems, Man, and Cyn.*, July, 1996, in press.
- [Wong 96a] STC Wong, P Whaley, CS Cheong, K Soo Hoo, J Wang, HK Huang. Interactive Query and Visualization of Medical Images on the World Wide Web, *SPIE Medical Imaging 96: Display*, Newport Beach, CA, 1996.

# **High Speed Image Transmission Via the Advanced Communication Technology Satellite (ACTS)**

Todd M. Bazzill, H.K. Huang  
University of California, San Francisco, Department of Radiology,  
Laboratory for Radiological Informatics  
San Francisco, California. 94143-0628

George R. Thoma, L. Rodney Long, Michael J. Gill,  
Lister Hill National Center for Biomedical Communications,  
National Library of Medicine, Bethesda, Maryland. 20894

## **Abstract**

We are developing a wide area test bed network using the Advanced Communication Technology Satellite (ACTS) from NASA for high speed medical image transmission. The two test sites are the University of California, San Francisco, and the National Library of Medicine. The first phase of the test bed runs over a T1 link (1.544 Mbits/sec) using a Very Small Aperture Terminal (VSAT). The second phase involves the High Data Rate Terminal (HDRT) via an ATM OC 3C (155 Mbits/sec) connection. This paper describes the experimental set up and some preliminary results from phase 1.

## **Background**

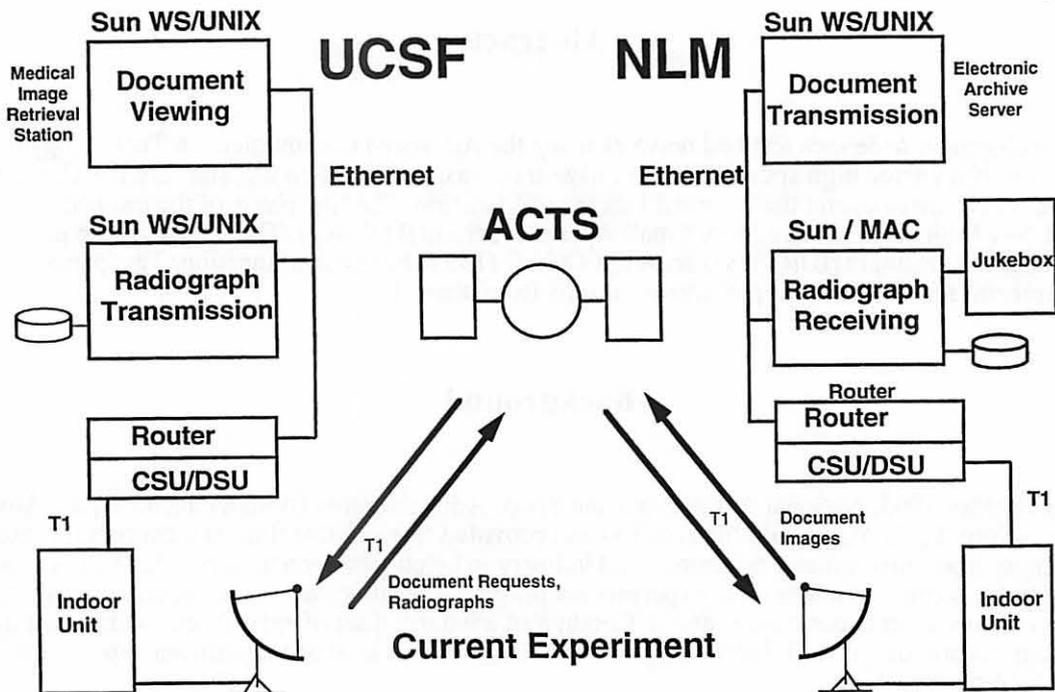
In October 1993, National Aeronautics and Space Administration (NASA) launched the Advanced Communications Technology Satellite (ACTS) and provided free satellite time and support for interested experimenters from government, academic, and industry in helping to open up a new Ka-band channel for high bandwidth communications. The experimental program objectives are " 1. Conduct a complete set of technology verification experiments, and 2. Conduct of a balanced set of experiments which evaluate the potential application of the ACTS technologies, with the goal of at least one significant experiment in each of the application areas. "

In 1994, the Laboratory for Radiological Informatics (LRI) at University of California, San Francisco (UCSF) was selected by the High Performance Computing and Communication (HPCC) Program to conduct a high performance testbed network for telemanagement of neuro-imaging under the auspices of the National Library of Medicine. The LRI's HPCC project is also supplemented by a California Research and Education Network program (CalREN) sponsored by Pacific Bell in which the 155 Mbits/sec OC3C Asynchronous Transfer Mode (ATM) technology network connection is made available for LRI to conduct the HPCC project in the San Francisco Bay area. The combined program between HPCC and CalREN makes it possible for the first time to establish a high speed test bed network in a clinical environment. However, this test bed network is limited to a metropolitan area because the connection does not extend outside of San Francisco Bay area. In order to fully test the capability of this

high speed test bed network, an extension of the network from the Bay area to national geography is necessary. The announcement from NASA on ACTS provides us with an opportunity to extend the test bed network to a wider area using satellite communication without the expense of developing terrestrial communication links. A proposal to NASA was submitted by Lister Hill National Center for Biomedical Communications, National Library of Medicine (NLM) and UCSF. In the proposal, high speed transmission using the ACTS for biomedical images is the goal. We are among several sites selected for the biomedical image communications experiment by NASA in 1995.

### Experimental Conditions

The experiment is conceived as having two phases, the first phase being run over a T1 link (1.544 Mbits/sec) using a Very Small Aperture Terminal (VSAT *Figure 1*).



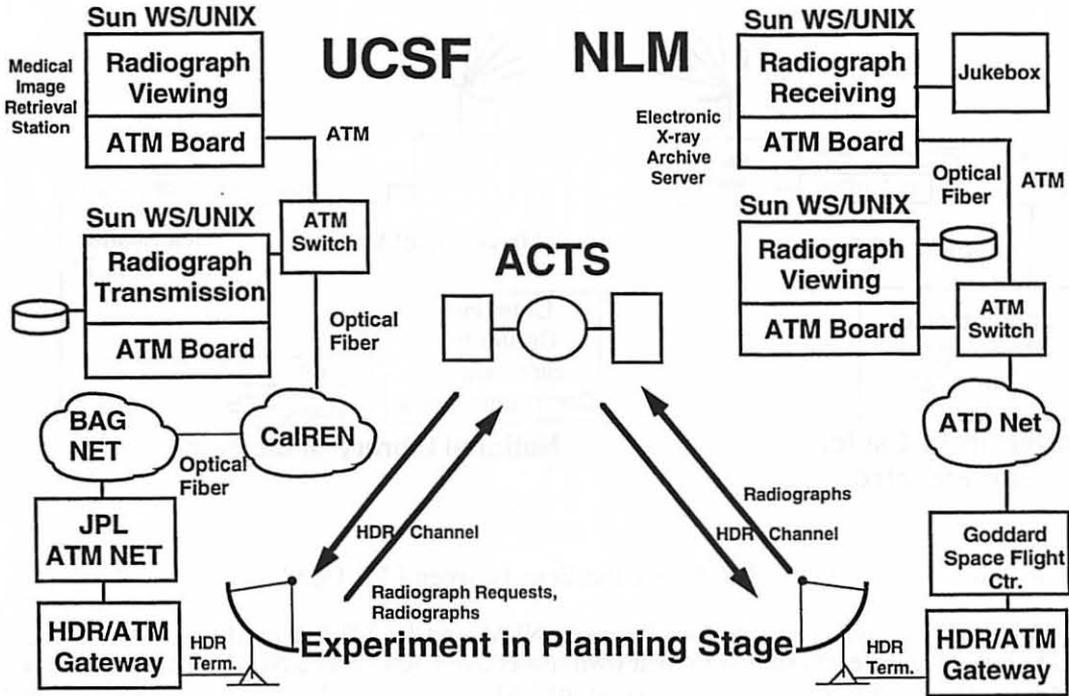
(Figure 1: Schematic showing the Current ACTS T-1 Experiment)

Figure 1 shows the current configuration of the UCSF - NLM model. Each has installed a 1.2 meter satellite dish and earth station which consists of the various electronics responsible for power, timing, and data flow to and from the satellite. The output of the earth station is similar to the output of a standard T1 land line and is connected to a CSU/DSU (carrier service unit / data service unit) and then interfaced to a router to connect to the LAN via the Ethernet protocol. Once the connection between the earth station and the LAN's have been made, document and image transfers can move seamlessly through the VSAT.

The second experiment is envisioned as duplicating the same procedures as in phase 1 except using a High Data Rate Terminal (HDRT) which communicates to the VSAT via a OC3C link. This HDRT connection

has a throughput of 155 Mbits/sec from which it will be connected to our current ATM (Asynchronous Transfer Mode) network.

Results from the T1 phase are presented in this paper. The ATM phase is still in the planning stage, *Figure 2* represents the architecture for the experiment.



(Figure 2: Schematic showing the Planned ACTS ATM Experiment)

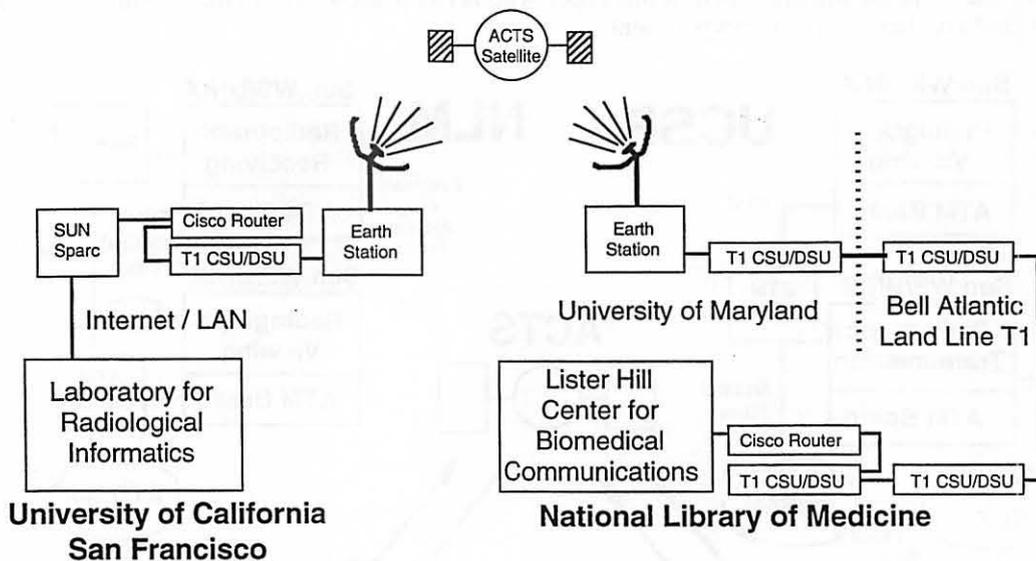
To communicate at a high data rate via the HDRT connection, each earth station and satellite dish will be replaced with a multi-meter dish and a trailer to house the electronics and Light Guides. NLM will have a direct connection from its gateway to a receiving station at Goddard Space Flight Center, Greenbelt, Maryland via the ATD Net (Advanced Technology Demonstration Network.) LRI will interface the HDRT through JPL (Jet Propulsion Laboratory) located in Pasadena, CA via BAG NET (Bay Area Gigabit Network). From there, the ACTS ATM net will be connected to the UCSF LAN/WAN ATM networks via a HDRT/ATM gateway device. This connection will allow NLM access to the UCSF ATM networks which spans several buildings and hospitals in the San Francisco Bay area [1].

## Materials and Methods

### Setting up the satellite communication

The Ka-band (2.5 GHz. bandwidth) of the ACTS is used and links UCSF with a NASA earth station at University of Maryland at College Park, MD. The University of Maryland is connected to NLM through a dedicated T1 land line provided by Bell Atlantic (*Figure 3*.)

# ACTS Physical LAN Connections

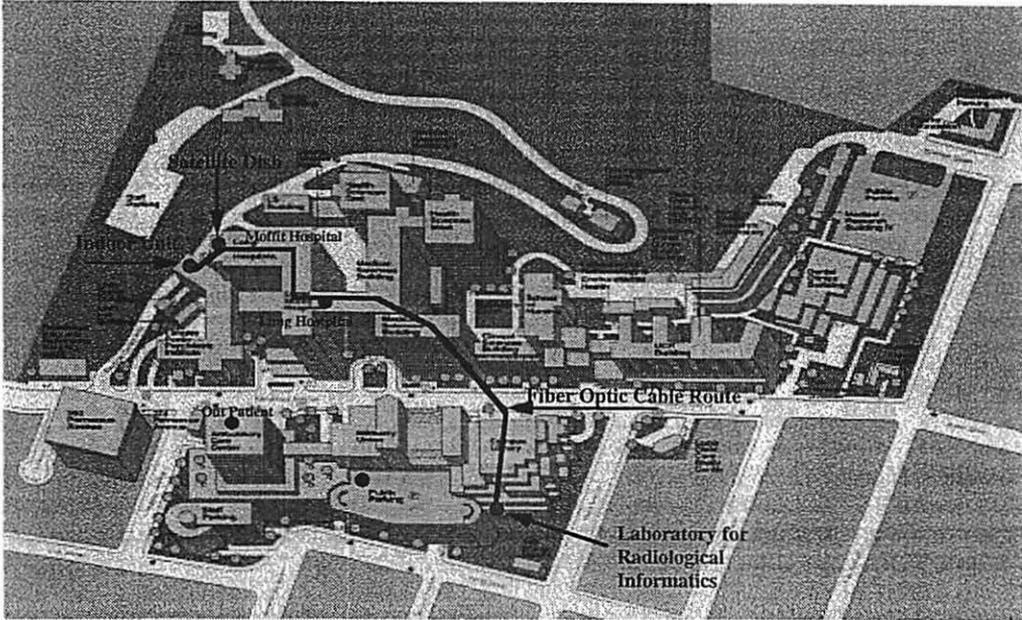


(Figure 3: ACTS Physical LAN Connections between LRI, UCSF and CEB, NLM)

Figure 3 shows the hardware connections between NLM and UCSF. Each site must use a router and CSU/DSU to interface the earth station to their own respective LAN. Since NLM is using a local land line T1 to connect to the earth station, an additional set of CSU/DSU's must be used to properly communicate through the local carrier.

In May of 1995 a 1.2 meter satellite dish and earth station were installed by NASA at UCSF. NASA dispatched a site visit team of engineers to determine if the location for the installation was feasible, and the time needed to install the system based on three criteria. The first and most important criterion was to investigate if the satellite could be located from the planned site for the dish (outdoor unit). A hand held GPS (global positioning satellite) unit was positioned at the dish site to establish a link to a minimum of three other satellites from a maximum of five. Once three satellites locked on to the GPS unit, the coordinates of the dish location were displayed. From this information, the "look" angle was calculated to determine if the propagation path to the satellite from the proposed antenna installation site was clear of obstructions. The second criterion was to verify that there are acceptable ways to route the cables from the dish location to the indoor earth station. We planned to install the dish on a overhang of a neighboring building in which a previous satellite dish was located with existing cable route to the location of the indoor earth station. By choosing this route we would be able to follow the existing cable path without any obstructions or hindrances and therefore would save the construction cost. Third, the NASA team wanted to verify that the proper power, grounding for both indoor and outdoor units, and anchoring of the dish were available. NASA determined that it would take approximately 2-3 days to install the system. The installation took a week because there were difficulties locating the "bird" (satellite) after the dish was mounted and powered up. Minor adjustments to the "look" angle corrected the problem.

The ground station (indoor unit) is directly connected to the Laboratory for Radiological Informatics HI-PACS (Hospital Integrated Picture Archiving and Communication System) infrastructure using the standard Ethernet protocol through a fiber optic network.



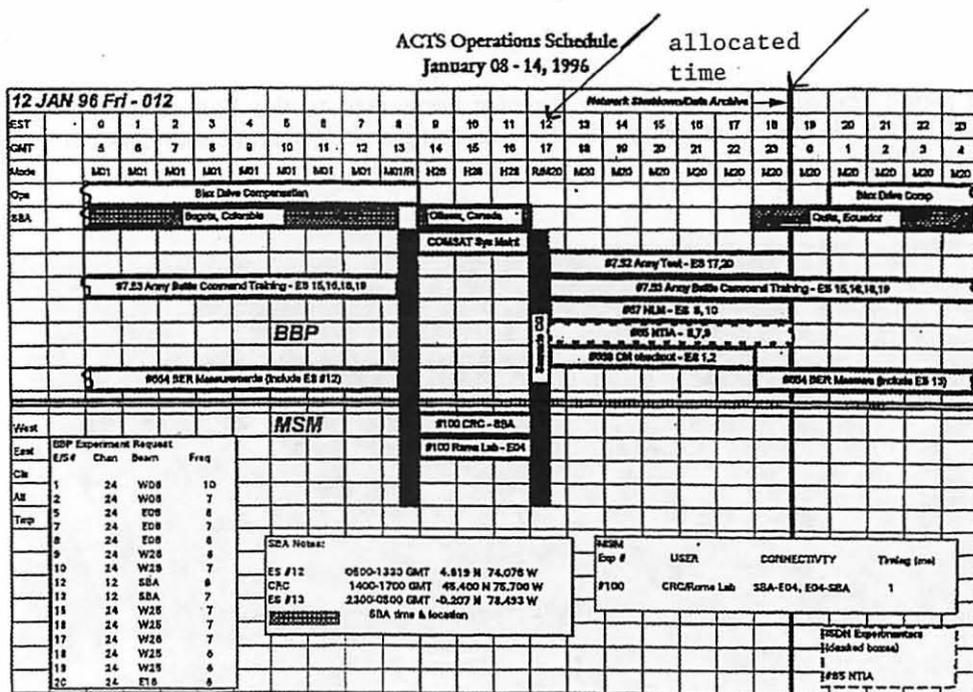
(Figure 4: Location of the ACTS Dish, Earth Station and their connections to the LAN at LRI, UCSF)

Figure 4\* shows an over head view of UCSF's Parnassus Campus. As shown, the satellite dish is located in rear of Moffitt Hospital and the indoor unit is on the ground floor of the adjacent building (Langley Porter). Using LRI's extensive fiber optic infrastructure which spans multiple buildings and levels on the UC campus, we were able to provide a link between the indoor unit and LRI's PACS computer room which is located across the street on the ground floor of the Campus Library. From here we are able to route the textual and image data anywhere within UC and our WAN environments.

Images transmitted through the ACTS can be displayed at workstations throughout the Radiology department. The satellite link allows NLM remote access to the HI-PACS at UCSF. UCSF has remote access to NLM databases, such as digital spinal X-ray.

---

\* Map courtesy of UCSF's WWW site.



(Figure 5: ACTS Operations Schedule, arrow shows the time period on a given day (Jan. 12, 1996) allocated by NASA for this experiment)

Figure 5 shows a typical testing time allocation by NASA. The ACTS Operations Schedule is faxed to each Earth Station manager each week. The table shows the usage time allocation for each site and the connections. For example: the week of January 8, 1996 UCSF station #10 and NLM station # 8 have a time slot of 12:30 P.M. to 6:59 P.M. EST. Also, in the lower left corner the table shows each Earth Station # and the respective number of the allocated 64Kb channels, beams used, and the frequency of transmission.

## Testing Procedures

There are two options to establish a link with NLM via the VSAT network : either dial in the correct access code including the number of 64Kb channels desired (up to 24), and the earth station to be connected directly at the indoor earth station; or, call NASA Ground Control and request for a link to be established with the same parameters. Once a link has been established between the VSAT and each earth station, the connection is transparent to the users. There are no delays or special codes to use to transfer information from one site to the other. At this point NLM and UCSF can send images and textual data to each other using various algorithms and parameters to track their performances.

## Multisocket TCP/IP

Conventional file transfer uses a single socket pair to establish the transmission between the send and the receive node. During the transmission, one principal server processor waits for a client to send for a request before a logical connection is established. Once established, data may be sent between the client and the server bidirectionally through the actual physical connection between this pair. The limiting factors in single socket transmission are the idling time of the server waiting for requests and the fixed transmission rate of the physical connection between the pair. The multisocket concept of transmission circumvents these bottle neck problems by first dividing the data to be transmitted into equal segments and

then establishes multi server and client processes. Multi servers and clients set up can take advantages of hardware resources with multiprocessors, better use of network routing, and exploitation of latencies resulting from TCP data acknowledgments [2]. The result is a faster data transmission rate between two nodes. The multisoocket set up is especially beneficial for large data sets like medical images.

At NLM, multisoocket TCP/IP was implemented to test the improvement of ACTS transmission rate between NLM and UCSF. So far in the initial experiment phase, we estimated that 80% of the maximum TCP throughput on bulk file transmissions from UCSF to NLM can be achieved by using the NLM multisoocket algorithm. This was 2.5 times faster than the conventional FTP. A similar rate was achieved for transmission from NLM to UCSF. Even though the basic capability of the multisoocket algorithm to achieve high throughput on this link has been validated, additional work needs to be done, for example: Testing the stability and performance of the algorithm and link fore repeated back-to-back transfers of large files; this mode of operation is needed to do mass transfers of image libraries efficiently.

### **Experimental Projects**

The following three projects are planned to test the ACTS T-1 transmission rates and reliability:

1. By using the Multisoocket File Transfer Program, developed by Communications Engineering Branch (CEB), test a prototype system to rapidly and accurately transfer a large volume of uncompressed digital x-ray image files. This transfer of images is to test the performance of a two way communications with massive data sets through the use of multiple sockets. This test will allow us to measure performances on continuous data transfers and to modify or tune the transmission parameters to consistently provide faster transmission times.

2. Test a medical information system that provides remote access to a national medical database at NLM in real time. The application to be run is the MIRS (Medical Information Retrieval System), developed by CEB. When in operation, this system will allow users throughout the nation to have immediate access to very large repositories of text and pictorial information. These repositories are expected to provide access to current medical literature, plus medical image data. Examples image data which may be available include digitized musculoskeletal x-ray images, MRI, CT, and digital photographic cross-sections of the human body. This will allow us to test one way communication of images from a single source (NLM) to multiple sites (laboratories at UCSF).

3. Test access to clinical radiology image databases to measure the performance of accessing a database where image data is acquired continuously 24 hours per day, seven days per week. NLM will install a client process on either a MAC or Sun workstation which will allow the use of the Radiology Workshop program to query in real time of UCSF's PACS image and textual databases [3]. NLM will be focusing on case histories relating to epilepsy and myelination. Radiologists and other interested medical workers will visit NLM to use and evaluate the system. In each case, NIH researchers will be identified who are interested in viewing the case histories by using the Radiology Workshop client in the CEB lab. Researchers are expected to access both text and image data, principally MRI case studies. Feedback and a critique of the features of the system will be solicited from the users. Data on the type of queries made by the users and the type of queries desired by the users will be collected.

## Results

Preliminary results showed that we can achieve raw data transfer rates of 160 KB/sec which is accounting for 80% of the theoretical T1 throughput of the ACTS. We believe that further fine tuning may allow us to further increase this efficiency. This raw data rate includes time to transmit application data plus all packet overhead, but no file reading or writing time. This is the rate we use to evaluate how much of the T1 line capacity we are achieving. (The corresponding application data rate, which includes time to transmit application data, packet overhead, and file I/O time, approaches 110KB/sec.) Currently, we are using Berkeley socket programming techniques which allows us to customize the transmission parameters.

We have noted that as of now transmissions from NLM to UCSF required many retransmissions. In the initial experimentation, we observed a similar phenomenon for transmissions from UCSF to NLM. This was traced to particular parameter settings on the router at the receiving (NLM) end. By resetting these parameters, we were able to virtually eliminate the retransmissions. However, as of to date, we have not been able to achieve similar results by resetting parameters in the UCSF router when transmitting from NLM to UCSF. This retransmission problem remains an open area of investigation. The different behavior of the multsocket algorithm and of FTP; depending on the transmission direction, may be related to these retransmissions and requires further investigation.

One draw back of using the ACTS system in a large scale test is the massive instrumentation and maintenance involved in setting up and maintaining the communication links. As a result, the ground supported equipment is quite expensive. Another problem we encountered is the reliability of the ground station, it requires tuning from time to time.

## Conclusions

The ACTS T-1 transmission offers the opportunity to exploit a test environment to gain experience in operating the high-bandwidth, wireless wide area network applications of the future. It further offers a chance to gain experience in the direct use of satellite channels for delivery of medical information with large files. The use of satellite communications may become increasingly useful in the delivery of time-critical, high bandwidth medical image information to remote areas or simultaneously to many areas.

## Acknowledgment

We gratefully acknowledge the support of the NASA ACTS Project and Experiments Office, especially Ron Schertler, Chief of the ACTS Experiments Office, and our liaison, Barry Fairbanks of the Analex Corporation; UMCP support provided by Dr. Timothy Kirkwood and his staff at the Center for Satellite and Hybrid Communications Networks; and the invaluable networking consultation and support provided by Dr. Hans Kruse of the Ohio State University.

## References

- [1] National Aeronautics and Space Administration (NASA) Project "VAMA: VSAT Access to Medical Archives" Internal Report, April 1995 Communications Engineering Branch, NLM.
- [2] L. Rodney Long, Lewis E. Berman, Leif Neve, Gautaam Roy, and George R. Thoma "An Application-Level Technique for Faster Transmission of Large Images on the Internet" SPIE Multimedia Computing and Networking Feb 6-8, 1995 San Jose, Cal. Vol: 2417 Pg 116-129
- [3] Ramaswamy MR, Wong AWK, Lee JK, Huang HK. "Accessing Picture Archiving and Communication System Test and Image Information Through Personal Computers" AJR, 1994; 163: 1239-1243

## **Current Status of the UCSF Second Generation PACS**

H.K. Huang, Ronald L. Arenson, Albert W.K. Wong, Todd Bazzill, Andrew S.L. Lou, Katherine Andriole, Jun Wang, Jianguo Zhang, Stephen Wong

Laboratory for Radiological Informatics,  
Department of Radiology, University of California, San Francisco, CA 94143-0628

### **ABSTRACT**

This paper describes the current status of the second generation PACS at UCSF commenced in October 1992. The UCSF PACS is designed in-house as a hospital-integrated PACS based on an open architecture concept using industrial standards including UNIX operating system, C programming language, X-Window user interface, TCP/IP communication protocol, DICOM 3.0 image standard and HL7 health data format. Other manufacturer's PACS components which conform with these standards can be easily integrated into the system. Relevant data from HIS and RIS is automatically incorporated into the PACS using HL7 data format and TCP/IP communication protocol. The UCSF system also takes advantage of state-of-the-art communication, storage, and software technologies in ATM, multiple storage media, automatic programming, multilevel processes for a better cost-performance system. The primary PACS network is the 155 Mbits/sec OC3 ATM with the Ethernet as the back-up. The UCSF PACS also connects Mt. Zion Hospital and San Francisco VA Medical Center in the San Francisco Bay area via an ATM wide area network with a T1 line as the back-up.

Currently, five MR and five CT scanners from multiple sites, two computed radiography systems, two film digitizers, one US PACS module, the hospital HIS and the department RIS have been connected to the PACS network. The image data is managed by a mirrored database (Sybase). The PACS controller, with its 1.3 terabyte optical disk library, acquires 2.5 gigabytes digital data daily. Four 2K, five 1,600-line multiple monitor display workstations are on line in neuroradiology, pediatric radiology and intensive care units for clinical use. In addition, the PACS supports over 100 Macintosh users in the department and selected hospital sites for both images and textual retrieval through a client/server mechanism. We are also developing a computation and visualization node in the PACS network for advancing radiology research.

### **KEYWORDS**

PACS, ATM, display workstation, networking, archival

### **INTRODUCTION**

The University of California at San Francisco (UCSF) is a health sciences campus with two medical centers, the main campus located at Parnassus Avenue and the Mt. Zion Medical Center located at Divisadero Street. The two medical centers are two kilometers apart. Table 1 shows some statistics of the two campuses. In addition, UCSF also affiliates with the San Francisco VA Medical Center. The second generation PACS at UCSF was designed for connection with all three medical centers.

**Table 1      The University of California at San Francisco**

<u>Parnassus Campus</u>	<u>Mt. Zion Campus</u>
560 Beds/76% occupancy	230 Beds/61% occupancy
22,200 admissions	6,800 admissions
7.1 day average length of stay	7.9 day average length of stay
289,000 patient visits	29,500 patient visits
23,000 emergency room visits	17,500 emergency room visits

In a previous report, we described the second generation PACS design concept [1]. This paper describes the current implementation status. The design concept of the second generation PACS at UCSF is based on standardization and open architecture. Table 2 summarizes some industry standards used in the system.

**Table 2      Certain Standards Used in PACS**

• Image format	Dicom 3.0 (ACR/NEMA 2.0) <sup>2</sup>
• Data format	Health Care Data - HL7 (Health Level 7) <sup>3</sup>
• Computer Operating System and Language	UNIX Operating System C Programming Language X-Window User Interface
• Communication Protocol	TCP/IP
• Database Query	SQL Structured Query Language

Open architecture design means that all application software programs are portable. Any change in hardware platform or software operating system (OS) requires only the development of an adapter (software layer) between the existing application software and the new platform or the operating system, and the application software remains intact. This design philosophy minimizes software development cost which remains a major obstacle in PACS implementation.

## **2. INFRASTRUCTURE DESIGN**

This section describes the intelligent image archive and distribution in the infrastructure design of the second generation hospital-integrated PACS. The infrastructure is based on a composite staging mechanism utilizing multiple storage media, hospital information system (HIS) and radiology information system (RIS), and the client server concept. Figure 1 shows a simplified PACS schematic at UCSF.

Two major aspects were considered in the implementation of the second-generation UCSF PACS: data integrity, which promises no loss of images once the PACS receives the images from the radiologic imaging systems; and system efficiency, which minimizes access time for images at the display stations. The following describes some major components.

### **2.1 Storage Management**

#### *Local Storage Management for Image Acquisition*

To ensure data integrity from image acquisition devices, the UCSF PACS always retains two copies of an individual image on separate storage devices until a successful archive of the image in the long-term optical disk library has been made. This back up scheme is achieved via the PACS inter-component communication in four levels:

- (1) Level 1 - At the radiologic imaging device

Images are not deleted from disk storage (Fig. 2A) of the imaging devices unless either the technologist verifies via the PACS terminal or the PACS controller acknowledges the successful archiving of individual images. Should any failure of the acquisition or the archive process occur, images can be re-sent from these imaging devices to the PACS.

(2) Level 2 - At the acquisition computer

Images acquired in the acquisition computer remain in its local magnetic disks (Fig. 2B) until the archive system acknowledges back to the acquisition computer that a successful archive has been accomplished. These images are then deleted from their residing magnetic disks (Fig. 2 A, B) so that storage space from these disks can be reclaimed.

(3) Level 3 and 4 - At the Archive System

Images arriving in the archive server from various acquisition computers are not deleted before their successful archive to the optical storage (Fig. 2D). On the other hand, all archived images are stacked in the archive server's cache magnetic disk (Fig. 2C) and will be deleted based on their aging criteria (e.g., number of days since an examination was performed, discharge or transfer of a patient, etc.). Figure 2 shows this four level archive scheme ensuring the data integrity.

*Local Storage Management for Image Display*

The storage management system for image display features three levels of user-accessible storage media: (a) RAID (redundant Array of Inexpensive Disks, Fig. 2E) in the display station for immediate access for current images; (b) magnetic disks in the archive server for fast retrieval of cached images (Fig. 2C); and (c) erasable magneto-optical disks and WORM disks (Fig. 2D) in the optical disk library for retrieval of any historical images. All high-resolution (2,048 x 2,500 pixels) and ICU ( 1,280 x 1,600 pixels) display stations in the PACS are configured with RAID of different capacities from 5 to 30 gigabytes. With this configuration, both 2K and 1K images can be displayed in less than two seconds.

## **2.2 Intelligent Archive Server**

*HIS/RIS/PACS interfacing*

Interfacing the HIS allows the storage management system to receive patient admission, discharge, and transfer (ADT) messages. Interfacing the RIS, on the other hand, allows the storage management system to receive information such as a patient arrival, examination scheduling, examination cancellation, examination completion, reports, etc. These events trigger the storage management system to perform image routing, stacking, aging, pre-fetching, studies grouping, and platter management mechanisms described in references 1 and 4. Exchange of messages among these heterogeneous computer systems is conducted using the Health Level Seven (HL7) standard data format with TCP/IP protocols on a client/server basis.

*Folder Manager*

The storage management system is characterized by its on-line patient folder management [4]. When the first radiological examination is scheduled, a patient folder is created in the PACS controller for the given patient. During the patient's hospital stay or ambulatory examination, this folder remains in the display station(s) for immediate access until the patient is discharged, transferred, or other aging criterion (e.g., two days after an out-patient visit) is met. The patient's ADT (admission, discharge, and transfer) information is obtained directly from the HIS and RIS. Images and associated data from any new examinations of the patient are continually added to the existing folder so that no redundant pre-fetching procedures will be performed. By applying the folder manager concept, the pre-fetching mechanism is only performed once per patient hospital stay or visit.

### *Integration with other manufacturer's PACS components*

Other manufacturer's PACS components using the Dicom 3.0 standard can be integrated into the UCSF PACS as well. An example is the Aegis ultrasound PACS (Acuson, Mountain View, CA) [5]. In this case, PACS treats the Aegis as a PACS acquisition device and organizes US images the same way as CT and MR images in the patient's image folder. Figure 3 shows the schematic of the connection.

## **2.3 Networking**

There are two digital networks and one video network in the UCSF PACS. The fiber optic video broadband network is used to connect all CT and MR scanners for real time patient monitoring during a scan. The operating principles are described in reference 6. Figure 4 depicts the schematic of the video system which connects eight CT and MR scanners and distributes their real time images to eight monitoring systems, one of which is shown in Figure 5.

The PACS digital networks are used in both local area networks (LAN) and wide area networks (WAN). In both LAN and WAN, asynchronous transfer mode (ATM OC3, 155 mbits/sec) technology is used as the primary network [7]. Ethernet and T-1 are used as the back-up networks for LAN and WAN, respectively. In case the primary network goes down, the back-up network is activated automatically by the PACS controller (described in Section 2.5). Figure 6 shows the logical WAN and LAN network connection.

## **2.4 Image Display**

The development of image display workstations is based on three implementation methods: using existing in-house workstations, working with manufacturers to develop new workstations, and distributing images and patient textual data to existing low-end desk top Macintosh computers. In the first type, we modified the two-monitor 2K display workstations developed by our group earlier by adding the HIS/RIS interface and some extra display functions [8]. An example is the Montage function which allows the assembly of images from different examinations into one file [9]. These workstations are used in the neuroradiology and pediatric radiology sections. Second, we worked with ISG (Toronto, Canada) to develop a two-monitor 1,600-line display station for intensive care unit ICU applications [10]. Figure 7 shows the schematic.

The third method is to develop a file server to distribute integrated PACS images and textual data to Macintosh desktop computers for individual physicians to review, teach, and perform research. The design concept, implementation, and user feedback are described in reference 11. This physician desktop workstation with full access to PACS and related data is a unique feature in the UCSF second generation PACS not available in any other commercial PAC systems.

## **2.5 PACS Controller**

The PACS controller is an intelligent machine that controls the flow of data within the entire PACS from acquisition (data input) to archive (long-term and short-term storage) and display (data output). The primary functions of the PACS controller include: (1) accepting images from acquisition nodes; (2) accepting HIS/RIS data; (3) updating global PACS database; (4) archiving images to optical disks; (5) routing images and HIS/RIS data to display workstations; (6) handling retrieval requests from display stations and (7) monitoring the network performance.

The computer system is based on the SUN 690 with four central processing units (CPUs), which allow multiple processes running simultaneously with minimal shared CPU time. Two standard network interfaces, the Ethernet and the 155-bit/sec bandwidth asynchronous transfer mode (ATM) networks, are used for receiving and distributing images. The PACS controller monitors the two networks and switches the ATM to Ethernet should the former fail. The optical disk library (Cranel, OH) attached to the PACS controller supports both erasable magneto-optical disks and WORM disks. A composite staging mechanism is implemented in the PACS controller to manage images stored in its

multiple storage media: magnetic disks (immediate-access temporary storage), erasable magneto-optical disks (longer-term archive data cache), and WORM disks (permanent storage). A mirrored database (Sybase, Emeryville, CA) is used for image management, the detail of which is described in reference 12. Table 3 summarizes the infrastructure characteristics of the UCSF PACS.

**Table 3 Infrastructure Design of the UCSF PACS**

<u>Function</u>	<u>Status</u>
Open architecture	yes
Connectivity	yes
Standards	ACR/NEMA, DICOM HL7 TCP/IP
Interface to HIS/RIS	HL7 and TCP/IP
Interface to other vendor PACS modules	DICOM
Mirrored database	yes
Auto routing	yes
Image pre-fetching	in-progress
Image sequencing	in-progress

### 3. IMPLEMENTATION

This section describes the current implementation status of the UCSF PACS.

#### 3.1 Networking

Table 4 describes the characteristics of the UCSF PACS network implementation as of today.

**Table 4 Characteristics of UCSF PACS Network**

<u>Function</u>	<u>Description</u>
Media	The network backbone consists of primarily fiber optics. CAT 5 UTP is used for workstation connections.
Router/ Gateway/Hub	Routers and gateways are used to divide subnets and ensure security among the networks. The design is a distributed hub configuration with multiple subnets. For example, Genesis network is used to connect the digital modalities i.e. CT, MR to the PACS controller. The departmental network is used to distribute images to Macintosh's users, etc.
Communications Technologies	ATM WAN and LAN via fiber are used for acquisition and distribution. Ethernet and T-1 are used for back-up.
Image Acquisition	PACS external network is used for transferring image data from the imaging modalities to the acquisition computers. The internal network with firewall protection is used to transfer data from the acquisition computers to

the PACS controller. The external network is also used as a back-up network for the display workstations.

Image distribution	Distribution to the 1K stations is done via ATM, a dedicated 155 Mb/s fiber optic connection directly to the PACS controller. Distribution to the 2K systems is done through an ATM to Ethernet switch which has dedicated 10 mb/s connections to each workstation.
Connection to the outside world	The network has T1 and ATM WAN connections to two affiliated medical centers, direct connections to the internet via the departmental network, and to over 100 Macintosh computers in the department.

### **3.2 Image Acquisition Component**

#### *CT and MR*

The UCSF CT/MR image acquisition systems consist of multiple vendors' equipment including GE, Siemens, and Imatron. The network is connected to 10 CT/MR scanners (CT: 3 GE spiral, 1 GE 9800 Quick, 1 Imatron; MR: 4 GE Signa 5, 1 Siemens Vision). The GE 9800 CT scanner uses GE IDNET-II configuration for interfacing and image acquisition processes. GE spiral CT scanners and Signa 5X MR scanners are based on DICOM standard communication protocol (upper level for TCP/IP). The Siemens and Imatron use their own proprietary interface protocols.

The PACS acquisition computers have a patient ID verification algorithm to correct typographical errors by imaging modality technicians, a mechanism to automatically recover image acquisition processes, and a central paging scheme to automatically page service engineers for system fatal errors [13]. The PACS uses both ACR/NEMA 2.0 as well as DICOM 3.0 header information format [12].

#### *Computed Radiography (CR)*

The UCSF PACS connects to one FCR AC2 and one FCR 9000 system. Both systems use ST-V (standard) and HR-V (high resolution) photostimulable phosphor plates. The digital interface to PACS for both systems uses DMS Bus with a RS-485 cable (which is a combination of RS-232 serial bus for messages and textual information and an RS-422 parallel bus for image data) connection to the data acquisition system manager (DASM). The DASM is basically a ring-buffered SCSI disk which transmits both textual and image data from the CR to the SUN acquisition computer over SCSI cable. All ICU portable, pediatrics, and newborn radiographic examinations use CR. We are acquiring between 2.0 to 2.5 Gigabytes/day image data including CT, MR, and CR but excluding US.

Two important features developed for automatic CR image acquisition are automatic background recognition and removal, and multi-layer adaptive process control. The former allows the automatic background removal from CR which allows a high percentage of correct automatic rotation and look-up table. The multi-layer adaptive process control guarantees no loss of CR images from the CR to the acquisition computer.

### **3.3 The Display Workstation**

The UCSF display workstations are comprised of 2K display stations, 1K display stations and Macintosh stations. The description of the 2K stations is given in reference 8. Additional features developed recently are a local montage feature for users to select images from different files for display. Another feature with 1 on 1 for the current image and 4 on 1 for historical images was specially designed for pediatric CR image viewing. The study list automatic update feature alerts users that the active (selected) patient has a new image file just arrived. Another new feature is a workstation usage

statistic software package to track system usage. This package allows refinement of workstation software based on users' working habits.

The 1K system was developed in collaboration with ISG Technologies, Inc. (Ontario, Canada) based on off-the-shelf components. This 1K workstation can support either 2, 3, or 4 1,600 line monitors, and has a user friendly interface. Figure 8 shows the schematic of the ICU distributed workstation design in which the 1K workstations are used.

The PACS also supports over 100 physicians' desktop Macintosh users to retrieve images and related PACS data for research, teaching, and case review. From these Macintosh computers, radiologists can select other printing resources in the department for hardcopy output. Table 5 summarizes the three types of image display stations.

**Table 5 Description of Three Types of Display Workstations**

<u>Types of Workstation</u>	<u>Description (No. of Stations)</u>
2K	SUN 4/470, Megascan display boards with two monitors.(4)
1K	SPARC 20, Turbo GX+ boards supporting 2, 3 or 4 monitors (5)
Physician desktop	Macintosh (over 100)

#### 4. DISCUSSION

The UCSF second generation hospital-integrated PACS was designed in-house based on the SUN workstation with the ACR/NEMA, DICOM, and HL7 standards. Relevant data from HIS and RIS is automatically incorporated in the PACS. Other manufacturer's PACS components which conform with these standards can be easily integrated into the system. The UCSF system also takes advantage of newer communication, storage, and software technologies in ATM, multiple storage media, and automatic programming for a better cost-performance system.

The new concepts of auto-routing, pre-fetching, and auto-sequencing have influenced our design in the UCSF PACS controller. With the client-server concept, we are able to implement the Macintosh server to distribute PACS images and patient textual data to every Macintosh user in the department. This feature is unique in our system.

The UCSF hospital-integrated PACS has been developed to such a stage that the infrastructure will not become obsolete and it can support any new PACS components as well as evolve continuously for higher reliability, better performance, and lower cost.

In addition to using PACS as an image management tool, we are also developing research programs to advance PACS technology as well as projects requiring large databases from PACS. In particular, the following six on going projects require large database support, and a computation and visualization node in PACS.

1. Brain: Non invasive epilepsy surgical planning using multimodality brain image,
2. Brain Temporal analysis of multiple sclerosis drug therapy progress (5 MRI sequences),
3. Brain: Quantitative assessment and staging of myelination disorder for children (MRI, MTC, Ultrasound),
4. Chest: Timeline quantitation of lung nodule volumes over treatment period (spiral CT),

5. Bone: Bone age assessment with digital hand radiographs (CR and digitized film), and
6. Prostate gland: Registration of prostate gland and tumor CT/MR images for conformal therapy planning.

The basic architecture of a computation and visualization node in a PACS environment is shown in Figure 9. The visualization engine used in this node is a SGI Onyx with two Reality Engines (Silicon Graphics, Mountain View, CA).

### ACKNOWLEDGEMENT

This work was partially supported by NLM Contract N01-LM-4-3508; IBM Almaden Research Center; CalREN grant ATMN-007, Pacific Bell; Lumisys; and Abe Sekkei, Japan.

### REFERENCES

1. H.K. Huang, R.L. Arenson, S.L. Lou, W.K. Wong, K.P. Andriole, et al. "Second generation picture archiving and communication systems" *Proc. SPIE Medical Imaging*, vol. 2165 pp. 527-535, 1994.
2. "ACR/NEMA Digital Imaging and Communication Standard Committee." *Digital Imaging and Communications ACR-NEMA 300-1988*. Washington, D.C.: National Electrical Manufacturers Association, 1989.
3. "Health Level Seven (HL7)." *An Application Protocol for Electronic Data Exchange in Health Care Environments Version 2.1*. Health Level Seven, Inc. Ann Arbor, Michigan, 1991
4. R.L. Arenson, D.E. Avrin, A.K.W. Wong, R.G. Gould, H.K. Huang. "Second generation folder manager for PACS." *Proc. Comp. Appl. to Assit Radiology*, vol. xxx, pp. 601-605, 1994.
5. M. Moskowitz, R.G. Gould, H.K. Huang, et. al. "Initial clinical experience with ultrasound PACS." *Proc. SPIE Medical Imaging*, vol. 2435, pp. 246-251, 1995.
6. H.K. Huang, R.H. Tecotsky, T. Bazzill. "A fiber-optic broadband CT/MR video communication system." *J. Digital Imaging*, vol. 5, no. 1, pp. 22-25, 1992.
7. H.K. Huang, R.L. Arenson, W.P. Dillon, A.S.L. Lou, T. Bazzill, A.W.K. Wong. "Asynchronous transfer mode (ATM) technology for radiologic communication. *Am. J. Roentgenology*, vol. 164, pp. 1533-1536, 1995.
8. H.K. Huang, R.K. Taira, S.L. Lou, W.K. Wong, et. al. Implementation of a large scale picture archiving and communication system. *J. Comp. Med. Imaging and Graphics*, vol. 17, no. 1, pp. 1-11, 1993.
9. M.A. Linnolt, S.L. Lou. "Improved reading efficiency by using Montage in a clinical PACS display." *Radiology*, vol. 197(P), pp. 257, 1995.
10. H.K. Huang, A.W. Wong, T.M. Bazzill, K.P. Anriole, J.K. Lee, J.G. Zhang. "Asynchronous transfer mode - distributed PACS server for intensive care unit applications." *Radiology*, vol. 197(P), pp. 247, 1995
11. M.R. Ramaswamy, A.W.K. Wong, J.K. Lee, H.K. Huang. "Accessing a PACS' text and image information through personal computers. *Am. J. Roetgenology*, vol. 163, pp. 1239-1243, 1994.
12. A.W.K. Wong, H.K. Huang, R.L. Arenson, J.K. Lee. "Multimedia archive system for radiologic images." *Radiographics*, vol. 14, pp. 1119-1126, 1994.
13. S.L. Lou, J. Wang, M. Moskowitz, T. Bazzill, H.K. Huang. "Methods of automatically acquiring images from digital medical systems." *J. Comp. Med. Imaging & Graphics*, (In press).

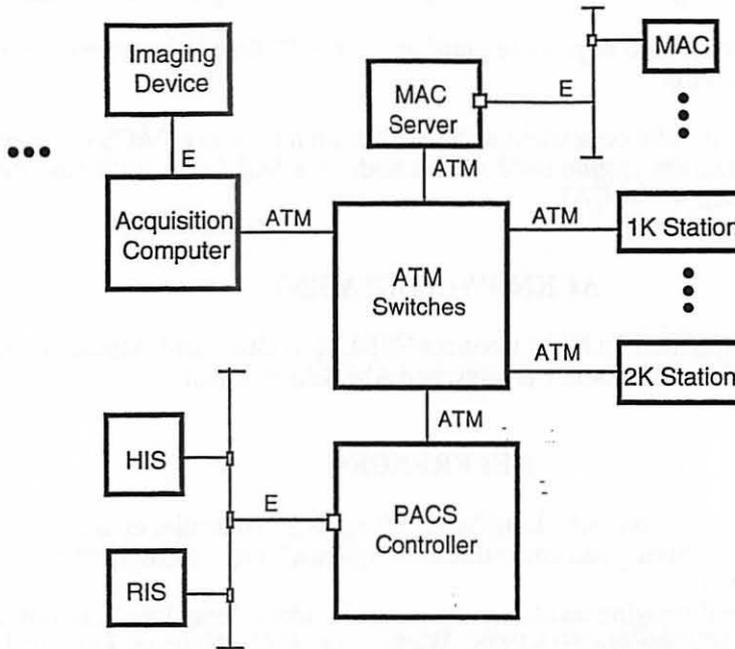


Figure 1. Simplified UCSF Second Generation PACS Schematic

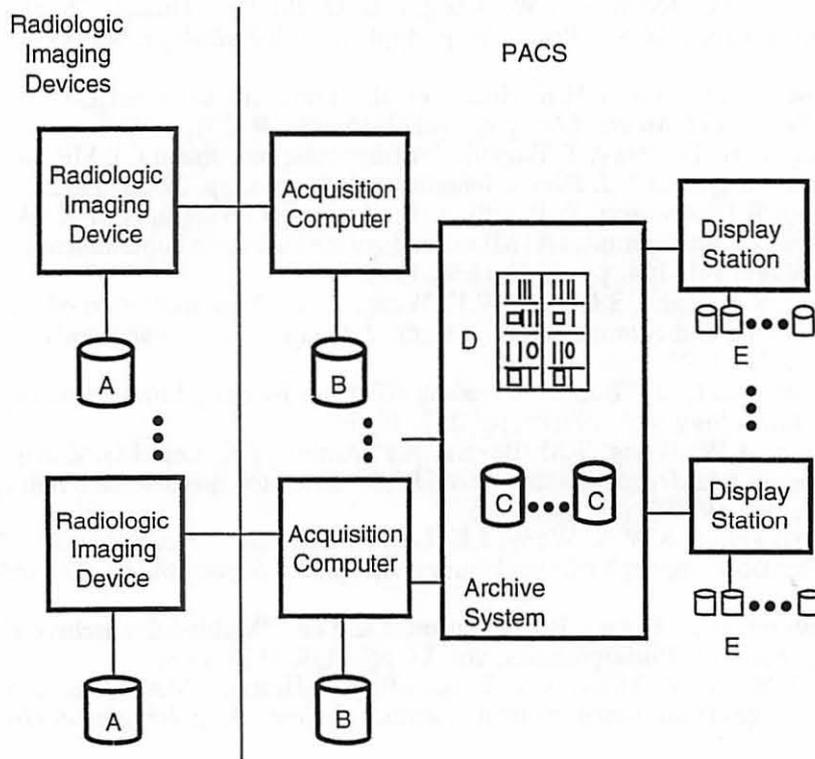


Figure 2. Four level image storage scheme to assure no images will be lost

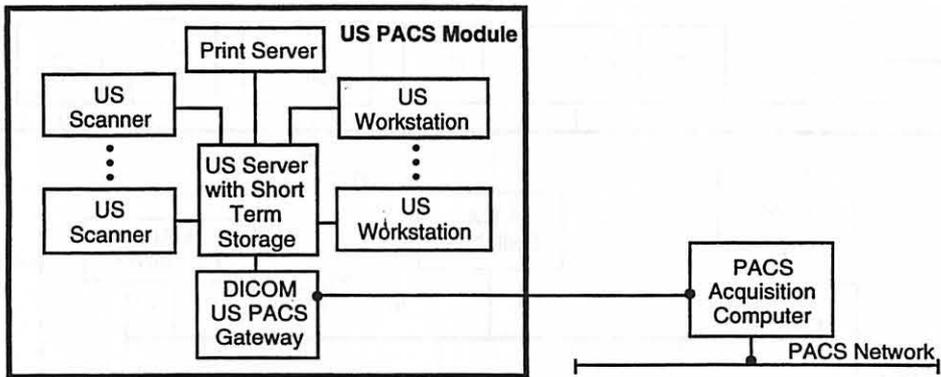


Figure 3. Schematic of the US connection to PACS with the DICOM Gateway

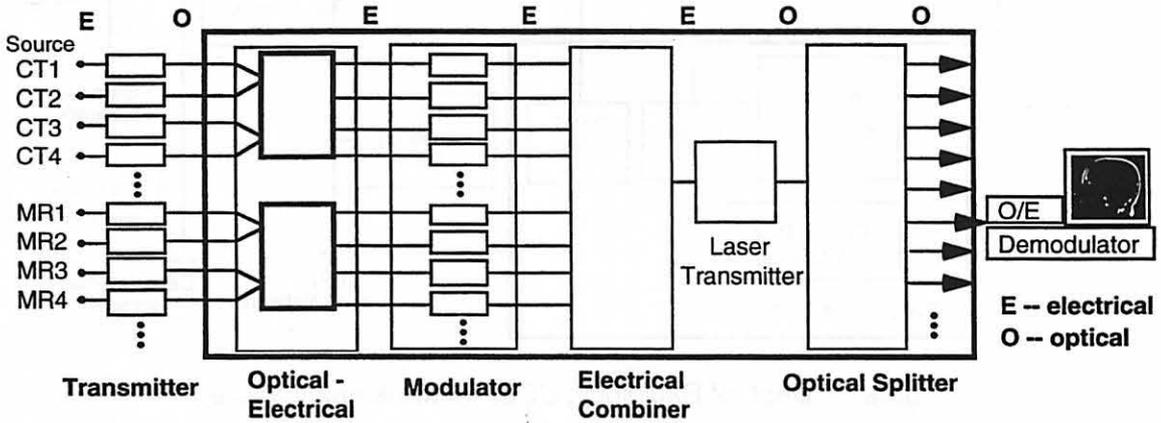


Figure 4. Schematic of the fiber optic broadband real-time video system for monitoring patient real-time CT/MR scans

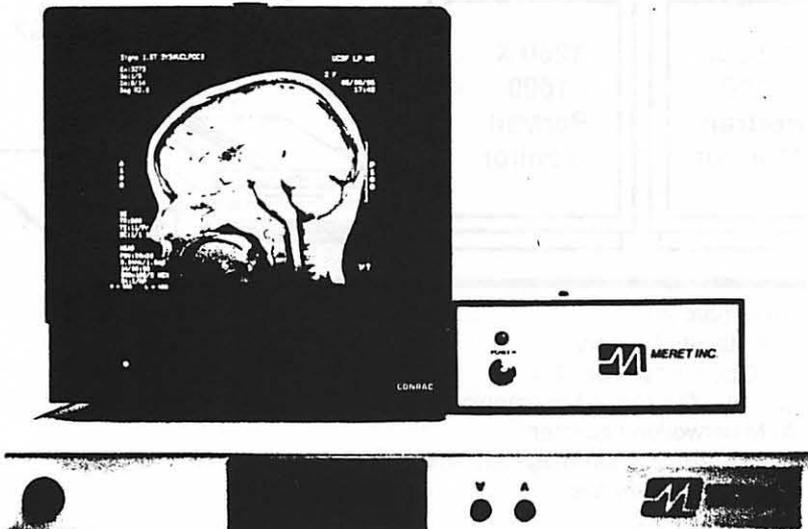


Figure 5. A video monitoring display station showing a real time image from one of the eight CT/MR scanners (MRLP) connected to the fiber optic broadband video system

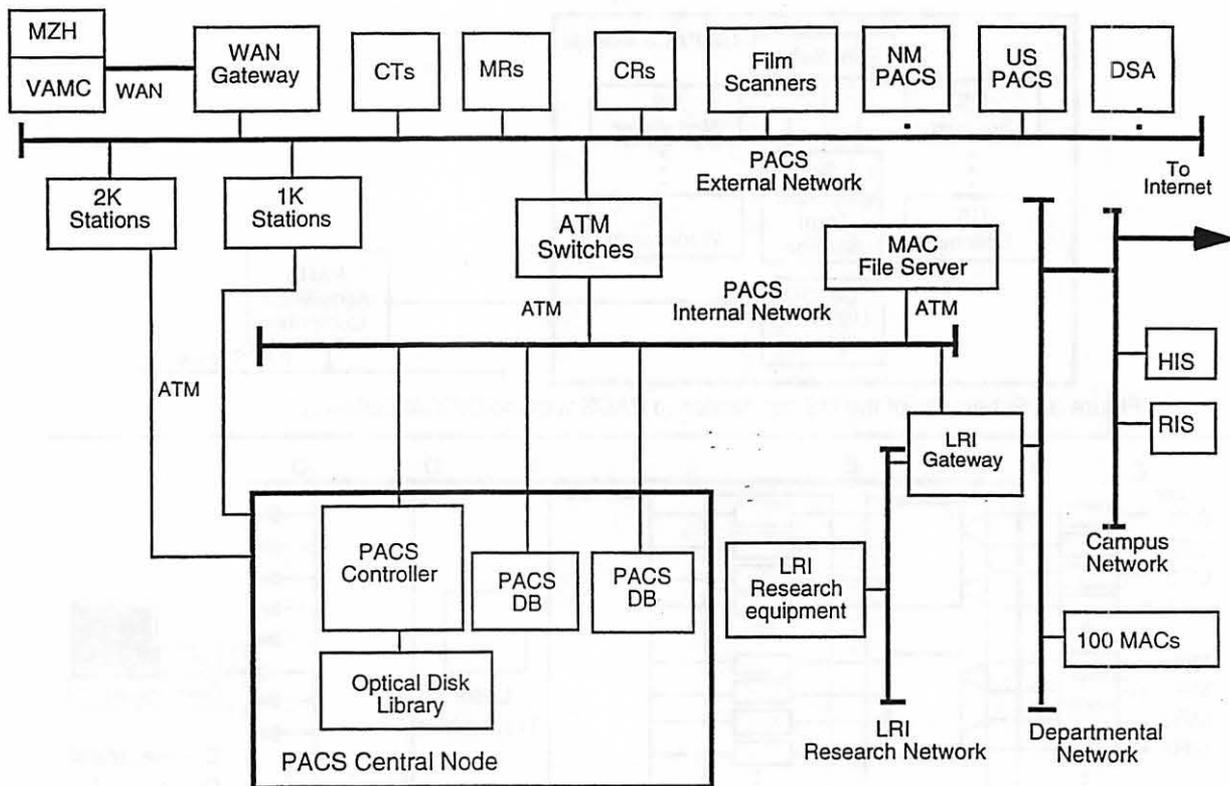


Figure 6. Dept. of Radiology, UCSF network architecture

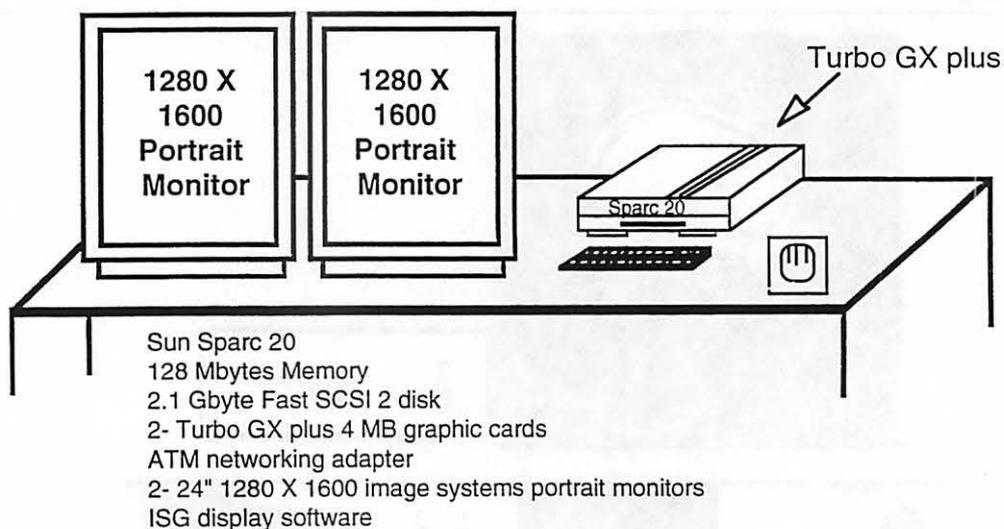


Figure 7. Schematic of the 1K workstation

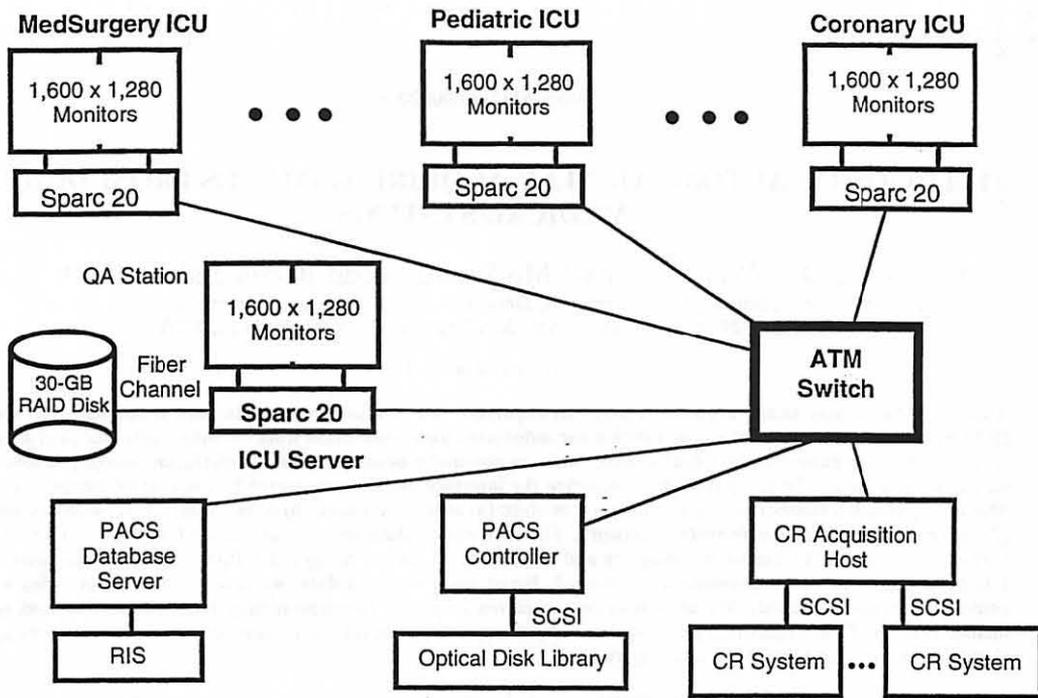


Figure 8. Schematic of the ICU distributed workstation design with multiple 1K workstations

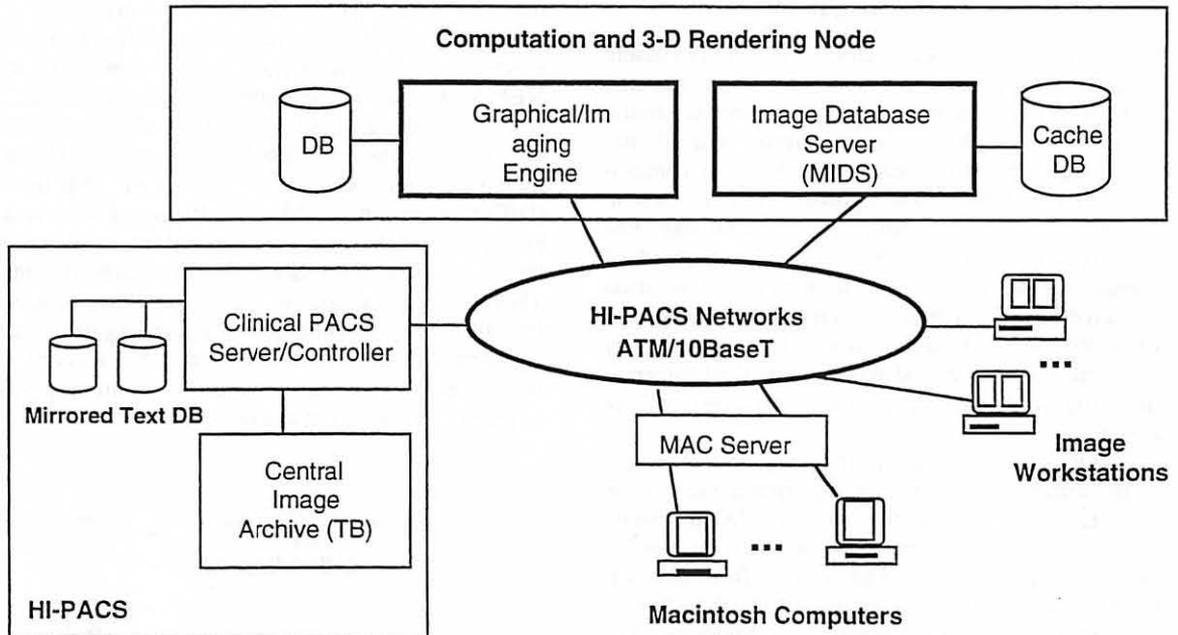


Figure 9 Simplified HI-PACS network with a Computation and 3-D Rendering Node



0895-6111(95)00020-8

## METHODS OF AUTOMATICALLY ACQUIRING IMAGES FROM DIGITAL MEDICAL SYSTEMS

S. L. Lou,\* Jun Wang, Michael Moskowitz, Todd Bazzill and H. K. Huang  
 Laboratory of Radiological Informatics, Department of Radiology, University of California,  
 Box 0628, 505 Parnassus Ave, San Francisco, CA 94143-0628, USA

(Received 13 January 1995)

**Abstract**—Automated image acquisition plays an important role in a picture archiving and communication system (PACS). However, there is no single solution for automated data acquisition from existing digital medical imaging systems. We have gained a great deal of experience on automatic acquiring data by interfacing imaging scanners of major manufacturers. In this paper, we categorize the interface methods supported by the current image scanners. This categorization consists of five architectural models: (a) sequential chain; (b) direct interface; (c) memory access; (d) shared disk; and (e) interconnected network. The cost, rate of data transfer, and ease of implementation of each model are discussed. To ensure the integrity and availability of patient images in a PACS system, automated fault tolerance design in image acquisition is required. Based upon our field data, we report common scenarios which cause the acquisition to fail. We also describe techniques employed to automatically restart the operations which include recovery from acquisition processes' errors and traps, image acquisition computer down-time occurrence, and shutdown occurrence of medical imaging system.

**Key Words:** Picture archiving and communication system, Image acquisition, Magnetic resonance imaging (MRI), Computed tomography (CT), Ultrasound, Computed radiography (CR), Film scanner

### 1. INTRODUCTION

With the advance of technology, an increasing amount of patient information (such as images and text) is being generated and processed in digital format by health care institutions. The digital technology provides advanced methods of communication and creates new approaches to education, research, practice of medicine, and publishing. For these reasons, there is a demand within the medical community to have access to patient information scattered in various digital systems. The development of a PACS meets this demand by automatically acquiring, archiving, and retrieving patient information. Figure 1 shows the basic components of the PACS.

In PACS, automatically acquiring patient images from different digital imaging modalities is an essential component of the complete PACS system. However, this component remains the most difficult task, because these modalities can be from different manufacturers and they may be distributed over geographically separate hospitals or buildings. The lack of open systems' interoperability among different manufacturers' imaging systems remains the

major obstacle. To circumvent this problem, the American College of Radiology and the National Electrical Manufacturers Association (ACR-NEMA) formed a committee to address this issue by establishing a standard. The standard has been evolved since the early 1980s to become the Digital Imaging and Communications in Medicine (DICOM) version 3.0. Most health care institutions, however, still have a large inventory of imaging systems based on "closed" architectural designs. These systems cannot be readily modified to meet the DICOM specifications, and upgrading or replacing them would be costly. To retain the services of the existing imaging systems for automatically acquiring images is a challenging issue.

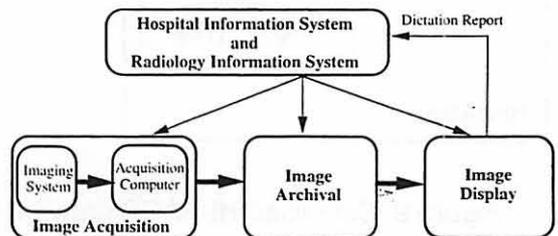


Fig. 1. A brief diagram shows the major subsystems in the picture archiving and communication system.

\*To whom correspondence should be addressed. Tel.: 415-476-5055.

In our effort of developing PACS in two medical centers since 1983, we have successfully integrated many imaging systems from major manufacturers into our PACS (1, 2). These include five laser film digitizers, four CR systems, nine MR scanners, 10 CT scanners, and one ultrasound PACS. With the experience gained from this developmental effort, we present this paper to focus on two subjects: (1) automated image acquisition interface methods; and (2) automated acquisition fault tolerance methods. The first subject discusses available interface methods which make automated image acquisition possible. The second subject presents methods to automatically recover image acquisition processes from potential faults and to assure that the operation of the image acquisition works around the clock.

## 2. AUTOMATED ACQUISITION INTERFACE METHODS

Generally speaking, the PACS image acquisition system consists of three major components: (1) a medical imaging system; (2) a computer system that acquires images from the imaging system (i.e. an acquisition computer); and (3) an interface mechanism (hardware and software) between the imaging system and acquisition computer. Specifically, the interface mechanism can be further categorized into five models: (a) sequential chain; (b) direct interface; (c) memory access; (d) shared disk; and (e) interconnected network. In the following, we illustrate the interface mechanism of each model by using manufacturers' imaging systems as examples. Each model's costs, data transfer rate, and ease of implementation are discussed.

### 2.1. Sequential chain model

The sequential chain model is defined as when the PACS acquisition computer links a medical imaging system through a chain of interface devices provided by the manufacturers. An example of the sequential chain model is the IDNET-1 (Integrated Diagnostics Network, version 1.0) solution provided by General Electric Medical Systems (GEMS) (General Electric Company, Milwaukee, Wisconsin) to acquire CT images from the GE-9800 CT scanners in late 1980s. Figure 2 shows the schematic diagram of IDNET-1 (dotted box). The parallel peripheral interface (PPI) board, residing in the scanner system, functions as a virtual magnetic disk driver that reads the data from the disk of the scanner to one of the network interface equipment (NIE) units, which is located next to the scanner system. At the

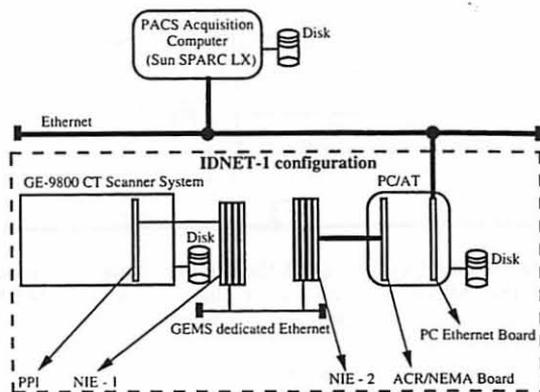


Fig. 2. The configuration of the IDNET-1 an example of the sequential chain model. Image data propagates through the PPI, the NIE-1, the GEMS dedicated Ethernet, the NIE-2, the PC/AT, and the PACS acquisition computer.

NIE-1 node (see Fig. 2), the image data and associated text data are encoded into the ACR-NEMA format and transmitted to the NIE-2 using a GE proprietary data transfer protocol and a dedicated GEMS network which is Ethernet-based.

The NIE-2 is the second node in the GEMS network and has a standard ACR-NEMA output (50-pin connector) to an ACR-NEMA interface board which is resided in a PC/AT. The PC/AT transmits the image data to the acquisition computer through the Ethernet with a PC Ethernet board.

Since this configuration requires several interface units, the cost of the interface is high (approximately \$40,000) and the connectivity is complicated. It requires several man-months to implement this configuration. Further, it is difficult to measure the elapsed time spent in each interface unit for determining the data transfer performance. The observation, however, is that the data transfer between the PC/AT and the PACS acquisition computer is the most time-consuming. We have measured the data transfer rate between a PC/AT and a Sun minicomputer 3/260 (Sun Microsystems, Mountain View, California) at about  $50 \text{ kb s}^{-1}$  (3). The sequential chain model has the disadvantages of being costly, complex, and having a low transfer rate. However, it was the only solution available to automatically acquire the images from the GE-9800 CT scanners to the PACS until GEMS introduced the IDNET-2. The interface configuration of IDNET-2 is simpler than that of IDNET-1; and it is categorized into the shared disk model discussed in Section 2.4.

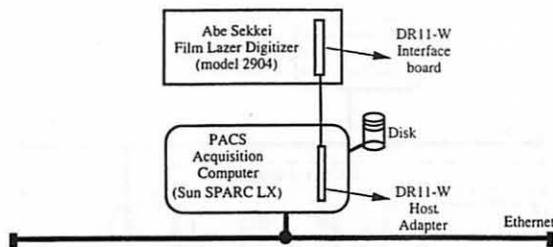


Fig. 3. The diagram shows the DR11-W interfaces used in the Abe-Sekkei film laser digitizer to the PACS acquisition computer.

### 2.2. Direct interface model

The direct interface model is composed of a PACS acquisition computer connected to a medical imaging system through a standard electronic interface board. An example for this model is the DR11-W interface (or SCSI interface) of the Abe-Sekkei (AS) film laser digitizer (Abe-Sekkei Inc., Tokyo, Japan). Figure 3 shows the interface configurations of the imaging system. In Fig. 3, the image data from the scanner are buffered in the DR11-W interface board and transmitted to the PACS acquisition computer. The buffer size in this application is 32 kb. Whenever the buffer is full, the data are archived to the disk of the PACS acquisition computer through the DR11-W host adapter. Other examples of this model are the SCSI data acquisition system manager (DASM) of the Fuji CR system (Fuji Medical Systems, USA Inc., Stamford, Connecticut) and the SCSI interface of the Lumiscan digitizer (Lumisys Inc., Sunnyvale, California).

The advantages of this model are: the interface units are commercial products, the connectivity is simple, the cost is affordable, and it has fast data throughput. We completed the integration of the AS film laser digitizer to our PACS in approximately two man-weeks. The price of this type of interface devices ranges from \$2,000 to \$8,000. The transfer rate between the AS film laser scanner and the SCSI disk of a Sun SPARC LX computer shown in Fig. 3 is greater than  $1.2 \text{ Mb s}^{-1}$ .

### 2.3. Memory access model

The memory access model is when a PACS acquisition computer connects to a medical imaging system through a dual-port RAM (random access memory) (4). The Imatron Cine CT scanner (Imatron Company, Oyster Point, California) utilizing an interface product called MegaLink to transmit images from the Cine CT scanner to the PACS acquisition computer belongs to this category. In Fig. 4, two MegaLink bus adaptor boards are linked by a

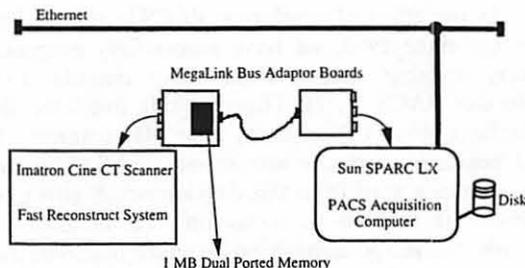


Fig. 4. The diagram demonstrates the MegaLink connecting the Imatron Cine CT scanner and the acquisition computer. Image data acquired by the acquisition computer are through the dual-ported memory.

pair of 25 ft ribbon cables (one for signal acknowledgement and one for data transfer). The adaptor board installed in the Fast Reconstruction System (FRS) of the Imatron CT scanner contains 1 Mb of dual ported RAM. The memory is accessible by both the FRS and the PACS acquisition computer. The FRS stores the header and reconstructed image data in this RAM which can be accessed by the PACS acquisition computer via the MegaLink bus adaptor board using the direct memory access mechanism. A simple semaphore-based interlock protocol software is used to synchronize the data transfer between the two computer systems (5, 6).

The direct memory access model provides fast data throughput because the data transfer is very similar to the scenario of writing data from the CPU memory of a computer to its own disk. Our measurement for this configuration (Fig. 4) is greater than  $1 \text{ Mb s}^{-1}$ . Since the interface devices consist of RAM memory, the cost is expensive (approximately \$15,000). We spent six man-weeks to complete the implementation of the automatic image acquisition from the Imatron CT scanner.

### 2.4. Shared disk model

We define the shared disk model as when a disk can be accessible by both the PACS acquisition computer and a medical imaging system. An example of this model is the Siemens (both Impact and Vision models) MR scanners (Siemens Medical Systems, Inc., Iselin, New Jersey). The Siemens MR scanner can be interfaced by using the network file system (NFS) protocol. The local disk of the imaging system can be remotely mounted by the acquisition computer through a network accessible by both computers. This scheme is shown in Fig. 5. Thus, whenever images are available in the local disk of the imaging

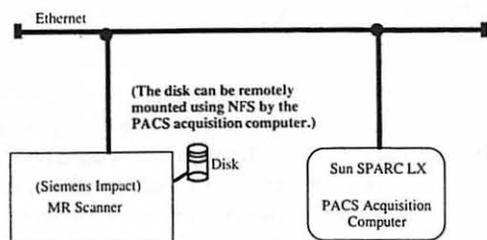


Fig. 5. An example of the shared disk model. The Sun computer using NFS remotely mounts the disk which is physically with the MR scanner. As a result, the acquisition computer can access the image data in the remote disk.

system, they are also available in the acquisition computer.

Of all the interface architectural models introduced in this paper, the shared disk model has the best image data availability. This is because no image data propagation is required in this model. In addition, the NFS-based shared disk model is a very low cost and easily implemented configuration because the NFS feature is commonly available in most of computer systems today. With one man-week, we were able to acquire images from the scanner to our acquisition computer. (In total, we spent 11 man-weeks to complete the software implementation to automatically acquire images from the Siemens MR scanner into our PACS.) However, there is a drawback to this NFS-based share disk model. Any interaction between the imaging system and the acquisition computer during data transfer requires the network and disk input/output (I/O) operations at the imaging computer system. Frequent I/O operations of peripheral devices may cause noticeable performance slowdown of the imaging system. This is an undesirable interruption to the operation of the imaging system; especially in the situation when the imaging system is in heavy use.

### 2.5. Interconnected network model

Interconnected network model is when a PACS acquisition computer and the medial imaging system host computer are connected in a network and communicate through standard communication protocols. The GEMS' newer CT and MR scanners, such as hi-speed spiral CT and Signa-5X MR scanners, are designed based on this model. Figure 6 shows the configuration of a GE imaging medical system's interconnected network currently installed

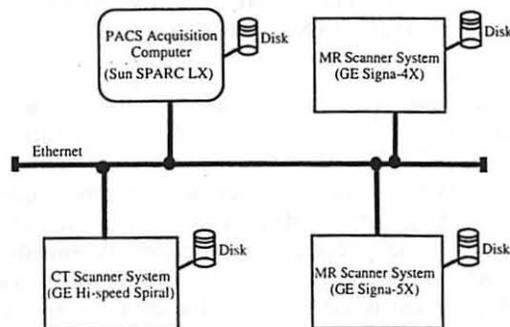


Fig. 6. The diagram shows the interconnected network model. The acquisition computer acquires images from the CT and MR scanners through standard network (Ethernet) and standard communication protocols (TCP/IP) and/or DICOM upper layer protocol for TCP/IP.

in our radiology department. The network follows the Open System Interconnection (OSI) standard layers specified by the International Standards Organization (ISO) (7). The layers of physical and data link are Ethernet, the network layer is Internet Protocol (IP), and the transport layer is Transmission Control Protocol (TCP). For the application layer, the GEMS' proprietary communication programs based on TCP/IP and a File Transfer Protocol (FTP) are used. We are in the process of replacing this GEMS program with the Central Test Node (CTN) software developed by Electronic Radiology Laboratory, Mallinckrodt Institute of Radiology (St Louis, Missouri) which is based on the DICOM V3.0 standard (8). In the meantime, we successfully acquired ultrasound images from Acuson Corporation's (Mountain View, California) Aegis-PACS using the CTN software (9).

Since the interconnected network model follows the industrial standards, its advantages include affordable cost, portable components, and easy implementation. In our experience, the major effort in configuring this model is laying down the network infrastructure. For an existing networking environment, there will be minimal effort (less than one man-week) required for configuration. However, in order to automatically acquire images, we spent approximately 12 man-weeks to complete the software implementation. The image data transmission performance (disk-to-disk) of the interface model ranges from 100 to 400 kb s<sup>-1</sup>. The result is so wide ranged because it depends on: (1) type of computer systems used (both the imaging system and acquisition computer); (2) type of system disks; (3) utilization of the network; and (4) workload of the computer systems.

### 3. AUTOMATED ACQUISITION FAULT TOLERANCE METHODS

A mechanism has to be designed in the implementation of image acquisition to detect and recover down-time occurrence of the acquisition process. If the period of the down-time is too long, the images in the imaging system may be purged, due to limited amount of disk space, before they are acquired by the acquisition computer. Potentially, this may cause the images to be lost forever if no backup system is used to store the data. To ensure the integrity and availability of patient images in a PACS system, automatic recovery from faults is a crucial implementation in the image acquisition design. In the following, we first briefly describe the major software involved in the acquisition computer. Then, we present the common scenarios causing the acquisition process to fail and the methods used to automatically recover the operation.

#### 3.1. Image acquisition software

In general, the image acquisition task consists of four programs, shown in Fig. 7. The acquiring program receives images from the imaging system to the acquisition computer. The formatting program organizes the acquired images and the associated text information based upon a standard format (e.g. DICOM). The sending program transfers the formatted image to one component of the PACS (see Fig. 1). This component is usually an image archiving system but depending upon the PACS architectural design, it can be an image display system as well. After the formatted image is properly stored in the PACS, the deleting program erases the acquired images and the formatted images from the acquisition computer to free-up the storage space. These four programs complete the chain of image acquisition data flow. If the chain is broken, the image acquisition task will fail to function. For simplicity, we will use the term "acquisition process" to

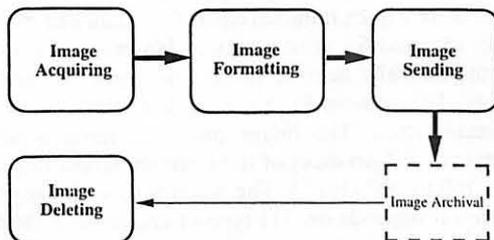


Fig. 7. Four major programs running in the PACS acquisition computer completes the chain of the image acquisition task.

represent these four programs in the following sections.

#### 3.2. Acquisition process recovery from errors

The acquisition process can be ended prematurely because of some fatal errors encountered. Examples include I/O errors of peripheral devices and not enough CPU memory. A monitoring process is required in order to recover the acquisition task from this type of fault. The monitoring process should possess three functions: (1) periodic alert times; (2) examination of the status of designated processes; which (3) restarts the designated processes if necessary. For example, the monitoring process wakes up at a given time and detects that the acquisition process has been terminated; it then restarts the acquisition process automatically. Usually, if the computer system has a clock daemon utility available, the periodical wake-up feature of the monitoring process can be replaced by the clock daemon. In this way, the monitoring program is ensured to be executed periodically as long as the acquisition computer is in operation.

#### 3.3. Acquisition process recovery from traps

Occasionally, after the acquisition program is launched by the monitoring process, the acquisition process is terminated right away. If this situation occurs repeatedly (i.e. the acquisition process is trapped), it will require human intervention to remove it. The most common example of this problem is that no space is left on the disk. This is a fatal error and requires immediate attention. To inform the service personnel automatically, a centralized dial-up and paging scheme is proposed. This dial-up scheme requires hardware configuration and software implementation. The hardware configuration of the central dial-up scheme is shown in Fig. 8.

The hardware involves a computer system to be configured with a modem device which is connected to a telephone line. This central dial-up computer can be any system in PACS, provided it is networked with the acquisition computers and others. From the task-oriented point of view, the central dial-up computer in this case is the server and the acquisition computers are the clients.

The software required in the server and clients mainly includes four modules: (1) sending of the service request; (2) receiving of the service request; (3) automatically dialing a person's page number; and (4) checking the availability of the computer systems networked. Consider the earlier example that no disk space is available, the acquisition process is trapped

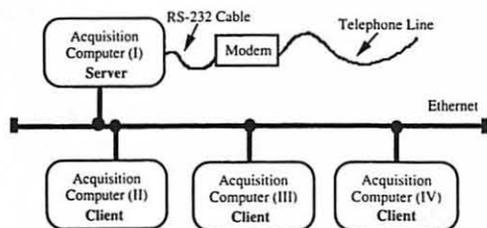


Fig. 8. The configuration of the central dial-up and paging scheme for recovery from traps. The modem is configured with the server and the clients. The clients send messages to the server through the Ethernet whenever the image acquisition is trapped. The server then sends the messages to a service personnel through the modem via a pager.

and it requires human intervention. In this situation, a message is sent from the client (which has no disk space) to the server through a network. The message contains the client's computer name, the process name, and the service engineer's page number. The computer name and the process name can be represented by numbers. Each pair of numbers is then coded to become a call-back number. In this way, the dial-up software module delivers the message to the engineer on call based on the page number and the coded call-back number. Besides the function of receiving request from the clients, the software also checks the availability of the clients via the network. This is based upon the fact that when a computer goes down, it is not available to other computers through the network.

#### 3.4. Acquisition computer recovery from down-time occurrence

The acquisition process cannot proceed if the acquisition computer goes down. A two-level scheme is presented here to handle this problem. The first level is to restart the computer system and the second level is to automatically restart the image acquisition programs.

On the first level, the acquisition computer can be programmed to reboot itself when the offending factors are removed. A common factor is a power outage. However, there are situations which cannot be recovered without human service such as the loss of the power supply. In this case, the dial-up and paging scheme described in Section 3.3 can be used to request the service.

In the second level of recovery scheme, two approaches can be used. The first approach is to take advantage from the computer process rebooting procedures. Basically, the computer starts up an

operating system kernel program, checks the configured peripheral devices, and executes necessary system processes during the computer rebooting. Since the system processes can be started up in the rebooting procedures, the image acquisition programs can be launched as well. The second approach is to combine the monitoring process and the computer clock daemon described in Section 3.2. Although either approach can be adopted, it is recommended that the first approach should be used as the primary and the second as a backup.

#### 3.5. Handling imaging system shutdown

It is possible that the imaging system is shut down while the image data is being transferred from the imaging system to the acquisition computer. This is because the technician who manages the imaging system has to observe the maintenance schedule irrespective of the acquisition process that may be in operation. It would be unacceptable if the acquisition process requires special care whenever the technician turns the imaging system on or off. Thus, the acquisition process must be transparent to the technician. In order to meet this requirement, the software in the image acquisition computer has to be designed to handle such an abrupt shutdown event. This can be accomplished by properly handling of the error events such as a broken network connection, a dropped network connection, or a connection time-out. If one of these events is detected, the acquisition process resets the uncompleted task back to initial status so that in the next execution the task can be processed from the beginning. This scheme may repeatedly handle the same uncompleted task if the imaging system remains in down status. However, this handling scheme ensures that the acquisition process will continue operation whenever the imaging system is turned back on.

## 4. DISCUSSION

In Table 1, we summarize the five interface models according to the three parameters, i.e. costs, rate of data transfer, and time required for implementation. In this table, an example imaging system and the associated interface mechanism for each model is given.

The PACS is an essential ingredient for the realization of digital radiology environment. One crucial task is to integrate various types of image scanners and modalities into the open system PACS network. Existing scanners, however, are developed based on "closed" architectural design and do not communicate with one another. In this paper, we

Table 1: The summary of the five interface models

		Models				
Parameters		Sequential chain	Direct interface	Direct memory access	Shared disk	Interconnected network
Cost		<\$40,000	<\$8,000	<\$15,000	\$0	\$0
Rate of data transfer		50 kb s <sup>-1</sup>	>1 mb s <sup>-1</sup>	>1 mb s <sup>-1</sup>	N/A	100-400 kb s <sup>-1</sup>
Time required for implementation (unit:man-week)		>24	2	6	11	12
Example	Imaging system Interface	GE CT-9800 IDNET-1	Abe-Sekkei film digitizer SCSI or DR11-W	Imatron CT MegaLink	Siemens MR impact NFS	GE MR Signa5 × Ethernet and CTN

have categorized the current interface architectures into five models. We believe that this is the first attempt to classify interface mechanisms of image acquisition systems. The examples described in each model provide both health care institutions and medical imaging manufacturers techniques for systematic modification or extension of the current scanners and to fuse them into one uniform PACS platform.

The open system's interoperability has been realized as an important feature in a digital medical system. To meet this requirement, following the standards is the sole approach. For this reason, the direct interface model and the interconnected network model are more favorable upon designing medical imaging systems.

The integrity and availability of patient images in a PACS system very much depend on the up-time of the image acquisition process. To optimize the up-time probability of the image acquisition process, automatic recovery schemes from faults have to be implemented in the image acquisition software design. This paper illustrates some common scenarios causing the acquisition task to fail and the methods used to automatically recover the operation.

In the near future, PACS will become a medical imaging input system to other PACS. To ensure PACS possesses the open system's interoperability, we suggest the following considerations in designing and implementing a PACS: (1) for acquired image data and associated text data already conforming to the DICOM standard, the acquired data format should remain unchanged; (2) if the acquired data does not adhere to the DICOM standard, the PACS image acquisition systems need to convert the data to the DICOM standard; and (3) to transmit image data out to other PACS, PACS image data-sending

programs (such as the CTN software) should follow the definitions in the DICOM upper layer protocol for TCP/IP.

## 5. SUMMARY

Automatic image acquisition is a major component in PACS. This component is tedious and laborious to implement because it involves interfacing to different manufacturers' imaging systems which are mostly designed with a closed architecture. Based on our 10 years of experience in integrating image acquisition systems to the PACS, we have categorized these interfaces into five models: sequential chain, direct interface, memory access, shared disk, and interconnected network model. We discussed the advantages and disadvantages of each model based on its cost, data transfer rate, and ease of implementation. Finally, we presented four fault tolerance methods for error recovery should the automatic image acquisition process malfunction.

*Acknowledgement*—This work was supported in part by the National Library of Medicine HPCC program N01-LM-4-3508 and the California Research and Education Network (CalREN) grant number ATMN-007.

## REFERENCES

- Huang, H.K.; Taira, K.R.; Lou, S.L. Implementation of a large-scale picture archiving and communication system. *Comput. Med. Imaging Graph.* 17(1):1-11; 1993.
- Huang, H.K.; Arenson, R.L.; Lou, S.L. Multimedia in the radiology environment: current concept. *Comput. Med. Imaging Graph.* 18(1):1-10; 1994.
- Lou, S.L.; Huang, H.K.; Mankovich, N.J. A CT/MR/US picture archiving and communication system. *Proc. SPIE, Med. Imaging III.* 1093:31-36; 1989.
- Tanenbaum, A.S. *Structured computer organization*, 2nd edn. NY: Prentice-Hall; 1984.

5. Maekawa, M.; Oldehoeft, A.E.; Oldehoeft, R.R. Operating systems—advanced concepts. MA: Benjamin/Cummings; 1987.
6. Dijkstra, E.W. Cooperation sequential processes. In: Geunys, F., ed., Programming languages. NY: Academic Press; 1968; 43-112.
7. Stallings, W. Handbook of computer-communications standards: the open system (OSI) model and OSI-related standards. In: Vol. 1. Sams; 1990.
8. RSNA and Washington University. 1993 RSNA DICOM standard connectivity demonstration—CTN imaging software users' guide. Documentation for the Radiological Society of North America (RSNA) 1993 digital imaging and communications in medicine (DICOM) demonstration; 1993.
9. Gould, R.G.; Huang, H.K.; Filly, R.A. Interface of a US PACS module with a departmental PACS infrastructure. Radiology. 141; 1994.

**About the Author**—S. L. LOU received a B.Sc. in Biomedical Engineering from Chung-Yuan Christian University at Taiwan, ROC in 1982 and a Ph.D. in Biomedical Physics from the University of California at Los Angeles (UCLA) in 1991. He was a visiting assistant professor in the Division of Medical Imaging, Department of Radiological Sciences at UCLA from 1991 to 1992. He is currently Assistant Professor in the Radiological Informatics Laboratory, Department of Radiology, University of California at San Francisco (UCSF). He also holds a joint faculty appointment in the UCSF/UC Berkeley Bio-Engineering Graduate Program. His major research activities in past seven years have been in design and implementation of Picture Archiving and Communications Systems (PACS), digital communication methodology applied in PACS, image data compression, and image mapping techniques.

**About the Author**—JUN WANG received a B.Sc. in Physics from Peking University in 1984, China, and an M.Sc. in Medical Physics from Chinese Academy of Sciences in 1987. She was a Lee Hysan Fellow at the University of Oxford, UK for two years. She is currently a Ph.D. student in the Biomedical Physics

Graduate Program at the University of California, Los Angeles. Her research interests include the film digitization artifacts and medical imaging processing.

**About the Author**—MICHAEL MOSKOWITZ received a B.Sc. in Physics from State University of New York, Binghamton in 1983, and a M.S.E.C.E. from University of Massachusetts, Amherst, in 1985. He was employed as a Microwave engineer at Raytheon Electromagnetic Systems Division in Santa Barbara, CA, from 1985-1987, and was a Member of the Technical Staff of Superconductor Technologies in S.B., CA from 1988-1990. Dr Moskowitz received his Ph.D. in Biomedical Engineering from Thayer School of Engineering, Dartmouth College, Hanover, NH, in 1994, and is currently a Postdoctoral Fellow at the Laboratory for Radiological Informatics, Department of Radiology, UCSF. His research interests include PACS networks, mammography image processing, and electrical impedance tomography.

**About the Author**—TODD BAZZILL received his B.Sc. in Electronic Engineering at DeVry Institute of Technology in 1989. He worked for General Electric Medical Systems from 1989 to 1990 and UCLA Radiology Department from 1990 to 1992. Mr Bazzill is currently the Principle Development Engineer at UCFS Radiology Department and is responsible for all aspects of networking and hardware for the second-generation PACS development.

**About the Author**—H. K. HUANG is Professor and Vice-Chairman, Director of the Laboratory of Radiological Informatics, UCSF. He also holds a joint faculty appointment in the UCSF/UC Berkeley Bio-Engineering Graduate Program. Dr Huang designed and implemented the UCLA PACS from 1982 to 1992, which is now in daily clinical operation. At UCSF he is establishing the Radiological Informatics Laboratory (LRI). Two primary functions in the LRI are hospital-integrated PACS and multimedia radiological database. His other research interests are in ATM-wide and local area networks, and image processing.

**Feature Selection in the Pattern Classification  
Problem of Digital Chest Radiograph Segmentation**

**Michael F. McNitt-Gray, *Member, IEEE*, H. K. Huang, *Senior Member, IEEE*, and James W. Sayre**

Reprinted from  
**IEEE TRANSACTIONS ON MEDICAL IMAGING**  
Vol. 14, No. 3, September 1995

# Feature Selection in the Pattern Classification Problem of Digital Chest Radiograph Segmentation

Michael F. McNitt-Gray, *Member, IEEE*, H. K. Huang, *Senior Member, IEEE*, and James W. Sayre

**Abstract**— In pattern classification problems, the choice of variables to include in the feature vector is a difficult one. We have investigated the use of stepwise discriminant analysis as a feature selection step in the problem of segmenting digital chest radiographs. In this problem, locally calculated features are used to classify pixels into one of several anatomic classes. The feature selection step was used to choose a subset of features which gave performance equivalent to the entire set of candidate features, while utilizing less computational resources. The impact of using the reduced/selected feature set on classifier performance is evaluated for two classifiers: a linear discriminator and a neural network. The results from the reduced/selected feature set were compared to that of the full feature set as well as a randomly selected reduced feature set. The results of the different feature sets were also compared after applying an additional postprocessing step which used a rule-based spatial information heuristic to improve the classification results. This work shows that, in our pattern classification problem, using a feature selection step reduced the number of features used, reduced the processing time requirements, and gave results comparable to the full set of features.

## I. INTRODUCTION

THE PATTERN classification problem consists of classifying an observation of unknown origin into one of several classes on the basis of one or more feature values [1]–[4]. In our segmentation research, the observation of unknown origin comes from a digital radiographic image. From a pixel and its neighborhood, multiple features are extracted to form a feature vector, which is input to a pattern classifier. The classifier generally consists of a set of discriminant functions  $g_i(\underline{x})$ ,  $i = 1, \dots, c$  where  $\underline{x}$  is the input feature vector, and  $c$  is the number of classes [2]. The classifier assigns an observation with input feature vector  $\underline{x}$  to class  $i$  if  $g_i(\underline{x}) > g_j(\underline{x})$  for all  $j \neq i$ . That is, the classifier computes  $c$  discriminant functions and selects the class corresponding to the largest discriminant. The discriminant functions are formed using either linear or nonlinear combinations of the input feature values. For example, the linear discriminant function uses a linear combination of the input feature values to form its  $c$  discriminant functions; the coefficients of each linear function

Manuscript received March 8, 1994; revised April 20, 1995. This work was supported in part by Public Health Service Grants T32GM08375 and P01CA51198 awarded by the National Institute of General Medical Sciences and National Cancer Institute, respectively. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was J. S. Duncan.

M. F. McNitt-Gray and J. W. Sayre are with the Medical Imaging Division, Department of Radiological Sciences, University of California, Los Angeles, CA 90024-1721 USA; e-mail: mmcnett-gray@mail.rad.ucla.edu.

H. K. Huang is with the Laboratory for Radiological Informatics, Department of Radiology, University of California, San Francisco, CA 94143 USA. IEEE Log Number 9413564.

are calculated to minimize the probability of misclassification [1], [2].

In many pattern classification problems, there is no clear method to choose which features should be included in the feature vector. Obviously, features which distinguish among the classes are desired. If there is a list of candidate features, we want to know which features best discriminate among classes: some features may provide good separation among the classes, while others may provide only noise or redundant information (i.e., by being highly correlated to other features). One method of measuring the ability of a feature to discriminate among classes is to use the one way analysis of variance  $F$ -statistic which measures the between class variance divided by the within class variance (see Table I). However, this statistic only calculates the ability of individual features to discriminate among the classes; there is no information about the interaction between features: which features are correlated, which features provide complementary information, etc.

Stepwise discriminant analysis [1], [5]–[7] overcomes these limitations by analyzing at each step the ability of each feature to discriminate among the classes, *given* the features that have been selected in previous steps (see the Appendix). In the stepwise selection procedure, variables are chosen to enter or leave the selected feature list according to the significance level of an  $F$  test from an analysis of covariance, where the features *already selected* act as covariates and the feature under consideration is the dependent variable [6]. At each step, the features in the selected list are examined. If the feature in the list that has the least discriminatory power fails to meet the criterion to stay, then that feature is removed from the selected list and placed back on the candidate list. Otherwise, the feature not already selected that has the greatest discriminatory power (and exceeds the criterion to enter) is added to the selected list. This continues until all features on the selected list meet the criterion to stay and none of the other features meet the criterion to enter the selected list. In the commercial algorithm that we have used (procedure STEPDISC in the SAS software package, SAS Institute, Cary, N.C., [6]), the stepwise selection has been employed. In this paper, we have investigated the use of stepwise discriminant analysis as a feature selection technique and the impact of using the reduced, selected feature set on classifier performance and computer resources (i.e., CPU and disk I/O time).

### A. Segmentation and Related Pattern Classification Problems

Segmentation in radiographic imaging is the problem of dividing an image into its component parts [8], [9]. This can

TABLE I  
ONE-WAY ANALYSIS OF VARIANCE (ANOVA) TABLE

Source of Variation	Sums of Squares	Degrees of Freedom	Mean Square	F ratio
Between Classes	$SS_B = \sum_{i=1}^I J_i (\bar{y}_i - \bar{y})^2$	$\gamma_B = I - 1$	$MS_B = \frac{SS_B}{\gamma_B}$	$F = \frac{MS_B}{MS_R}$
Within Classes	$SS_R = \sum_{i=1}^I \sum_{j=1}^{J_i} (y_{ij} - \bar{y}_i)^2$	$\gamma_R = n - I$	$MS_R = \frac{SS_R}{\gamma_R}$	
Total	$SS_T = \sum_{i=1}^I \sum_{j=1}^{J_i} (y_{ij} - \bar{y})^2$	$\gamma_T = n - 1$		

Where:

- i = index representing class.
- I = total number of classes.
- j = index representing observations within a class.
- J<sub>i</sub> = total number of observations within class i.
- n = total number of observations
- y<sub>ij</sub> = the jth observation of feature y within class i.
- $\bar{y}_i$  = mean of feature y within class i.
- $\bar{y}$  = mean of feature y over all classes.
- SS = the sums of squares.
- $\gamma$  = the degrees of freedom
- MS = the mean squares for the source of variation.
- subscript B = between classes
- subscript R = residual (within) classes
- subscript T = total

be accomplished by a process of pixel classification, where the image is segmented into regions by assigning individual pixels of the image to specific region classes (such as the heart region or lung region in a digital chest radiograph).

The pattern classification approach has been used to segment many different types of medical images. Hasegawa used a convolution neural network architecture for the segmentation of the lungs in digital chest radiographs using only pixel gray levels as inputs [10]. Vannier *et al.* [11]–[15] used pattern classification techniques to classify each pixel of an MRI brain image into one of several tissue categories—gray matter, white matter, cerebrospinal fluid (CSF), and other tissue types—on the basis of each pixel's T<sub>1</sub>, T<sub>2</sub>, and proton density value. Following the work of Vannier, several groups used neural networks to classify the pixels of MR images [16]–[20] instead of or in addition to the minimum distance-to-means and maximum likelihood classifiers.

Neural networks as pattern classifiers have also been suggested to perform other radiological tasks such as: differential diagnosis of interstitial lung diseases using the radiographic findings as inputs [21], signal detection [22], diagnosis and interpretation of neonatal chest radiographs—again using radiographic findings as inputs [23], location and classification of tumors within ultrasound images of the eye [24], detection of cardiac diseases from echocardiographic images [25], detection of liver disease in ultrasound images [26], [27], detection of microcalcifications in mammograms [28], [29],

and detection and localization of lesions in single photon emission computed tomography (SPECT) [30].

In each of these research efforts, features that were believed to be important in the classification problem were used to train the classifiers. There has been little discussion concerning which of the features are actually important to the classification problem or how to go about deciding which features are more important than others. If some features provide redundant information or are not useful to the classification, then there may be some subset of features which can yield equivalent or better results than those given by the entire feature set. If there is an adequate subset, then the number of features in the feature vector may be reduced, reducing the training and testing time of the classifier; this can be very significant for a neural network (which trains by repeated presentation of the entire training set).

### B. Overview of Our Approach

In this study, we used stepwise discriminant analysis as a method of feature selection in a pattern classification problem; with an application in the segmentation of digital chest radiographs. First, anatomical classes are defined. Next, a list of candidate features—consisting of gray level based measures, measures of local differences, and measures of local texture—is constructed. The list is then put through the feature selection process based on stepwise discriminant analysis. The resulting list of selected features form the feature vectors which are used to train and test the classifiers—a linear discriminator (which uses a linear combination of the feature variables) and a feedforward, backpropagation neural network (which uses a nonlinear combination of the feature variables). Both classifiers use supervised training techniques in which the class of each training observation is known. For the linear discriminator, the training data are used to calculate the coefficients of the discriminant functions in a deterministic manner [1], [3], [4]; for the neural network, the training data are repeatedly presented to the network and the weights are determined by propagating errors back through the network in an iterative manner [31]–[35]. The classifiers are evaluated by comparing their results to the classification made by a radiologist, on a pixel by pixel basis. This comparison results in a classification matrix, from which we calculate percent correct for each image and, by combining results across images, for the entire test set.

In addition, a postprocessing step is applied in which results from the pattern classifiers are reclassified according to heuristic rules based on spatial information. Because the classifiers use only locally calculated features (i.e., gray level, local differences, and local textures), some pixels are misclassified and their misclassifications are obvious because of the spatial location—for example, pixels classified as background that are located in the middle of the lung. Rules were formulated to first identify the obvious misclassifications and then to reclassify the results.

To determine the effects of feature selection on both classifier performance and computer resources, two other data sets are constructed by: 1) using the entire candidate set of

59 features, and 2) choosing eight features—which are not among the selected features—at random. These data sets are evaluated in the same manner as above by training and testing classifiers. The results of the classifications with and without feature selection are reported. Differences in performance in terms of both the classification results and the processing time are investigated. The postprocessing step in which heuristic rules based on spatial information are applied is also carried out and the same performance comparisons are made.

The pattern classification approach to segmentation is different from previous methods in that it is not as dependent upon assumptions or *a priori* information (such as location of structures, etc.). Though the heuristic methods currently provide better performance for segmentation than do the pattern classification methods, the latter provide a strong starting point to apply heuristic methods. This will be illustrated by the postprocessing step in which heuristic rules based on spatial location are applied. These heuristic rules are shown to improve the performance of the classifiers.

## II. METHODS

Several experiments were carried out to determine the effect of feature selection upon the performance of the classifiers used in segmenting chest radiographs. Images were collected, the classes were defined, and a method for extracting all of the candidate features to form observations were developed. From the full set of features, three data sets were formed: 1) the full set of candidate features with no feature selection performed; 2) the best eight features from feature selection; and 3) a random set of eight features not in the top eight of selected features. This last data set is not expected to give good performance and is used to show the performance of selecting features at random—without regard to the feature selection process.

Using these three data sets, the classifiers were trained and tested on observations from 33 digital chest radiographs. The results of the trained classifiers were compared to those of a radiologist for each classifier and each data set. The results of these classifications were then used as inputs to a postprocessing step which added spatial information to improve the classification. The individual steps are described below.

### A. Collecting Digital Chest Radiographs

Thirty-three adult chest radiographs were digitized by a laser scanner (Lumisys model DIS-1000, Sunnyvale, CA) at a resolution of 4096 pixels per 35 cm and at 12-b depth per pixel (4096 gray levels). The digitized images were stored on a Sun Sparcstation2 computer's disk (Sun Microsystems, Sunnyvale, CA) as 2048 by 2048 pixel images with 12-b depth. The 33 images were all images that were diagnosed as normal by thoracic radiologists for another study [36]. The 2048 × 2048 images were smoothed and subsampled for analysis to a 128 × 128 image matrix size by a simple averaging of pixels; the gray level range of the lower spatial resolution images was maintained at 12 b. The smoothing and subsampling were done to reduce the amount of data obtained for each image

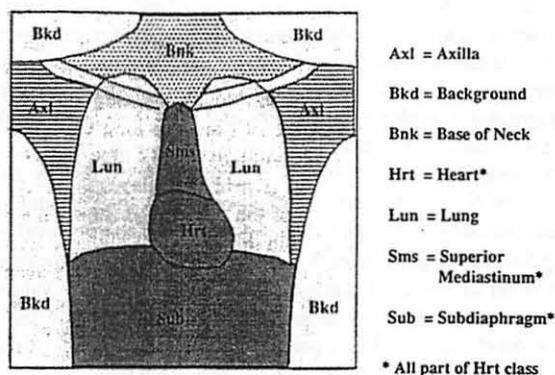


Fig. 1. Anatomic class definitions.

in order to speed up the computational process. An example image is shown in Fig. 5(a).

Images were randomly assigned to training and testing data sets, with an approximately equal number of images in each: 17 in the training set and 16 in the test set. Because these data sets were formed randomly, case selection is expected to have little effect on the results.

### B. Anatomic Class Definitions

In the chest radiograph, a projectional image is obtained that ranges superiorly from the bottom portion of the head, inferiorly to the subdiaphragm region of internal organs and laterally to the axilla (shoulder) and often includes parts of the arms. For this research, a classification is proposed that consists of several general regions described in Fig. 1: heart/subdiaphragm/superior mediastinum, lungs, axilla, base of the head/neck, and patient background. The heart, superior mediastinum, and subdiaphragm are combined into one anatomic class because of their similarity in gray level and local texture in the radiograph (preliminary tests show almost no separation among these three classes; a method other than pattern classification based on locally calculated features is necessary to discriminate among these regions). The structures of the axilla (arm and shoulder) are usually not clinically interesting to the thoracic radiologist, but these structures need to be accounted for in an automated segmentation of the image so that these pixels will not be confused with another anatomic class.

The five classes as outlined in Fig. 1 are:

- 1) heart/subdiaphragm/upper mediastinum;
- 2) lung;
- 3) axilla;
- 4) base of head/neck; and
- 5) background (area outside patient but within radiation field).

### C. Candidate Features

A list of candidate features was constructed. Three types of features were used in this approach: gray level based measures, local difference measures, and local texture measures. Features from these groups were chosen because of

the nature of the regions within a chest radiograph: the heart, superior mediastinum, subdiaphragm, and axilla are each very uniform areas with high gray levels; while the lung region has lower gray levels and more gray level variation (due to the radiographic complexity of the ribs and the lung vasculature). All 59 features are listed individually in [37, Appendix].

The gray level of a pixel is primarily determined by the amount of X-ray attenuation of the tissues and three candidate feature variables are included in this group: 1) the gray level of the pixel, 2) the gray level of the pixel divided by the average gray level of the entire image (to normalize the gray values to the overall brightness of an image), and 3) the gray level of the pixel expressed as a percentage of the histogram of the image (to account for the shape of the histogram).

Local difference measures are used to characterize the amount of homogeneity in a local neighborhood [8], [38]. The local difference measures used in the candidate list are:  $x$ -direction,  $y$ -direction, and total Sobel gradients (first-order difference) for both a  $3 \times 3$  and  $7 \times 7$  neighborhood; the standard deviation of pixels in a  $3 \times 3$  and  $7 \times 7$  neighborhood; and the Laplacian (second-order difference) for both a  $3 \times 3$  and  $7 \times 7$  neighborhood. The two neighborhood sizes were used for each measure to account for the scales at which important differences may occur.

Local texture measures are used to quantify the nature of local differences [39]–[43] and often take into account both distance and direction of the differences. This work uses four groups of local texture measures: microtexture masks, spatial gray level dependence matrices (SGLDM), oriented texture measures, and fractal dimension-based measures. The microtexture measures are four of those introduced by Laws [42] which are sensitive to: 1) horizontal edges, 2) high-frequency spots, 3) vertical edges, and 4) textures with low correlation (i.e., textures with more random arrangements). Four of the SGLDM measures are used [39], [41]: 1) energy and 2) entropy are measures of uniformity, 3) homogeneity measures the similarity between neighbors, and 4) inertia measures dissimilarity.

A third set of texture measures are the oriented texture measures, coherence and dispersion, described by Rao [43]. Coherence measures the consistency of directionality for local differences, using first-order differences, while dispersion measures the lack of this consistency. These features may be useful in distinguishing among regions where a strong directionality component exists, such as lung vasculature and airway branches, and regions where no directionality exist, such as the axilla.

The fourth texture measure is the fractal dimension of the image intensity surface [43]–[47]. The box-counting method for estimating fractal dimension was shown [43] to be faster than the intensity difference method suggested by Pentland [45] and is used in this research. The other texture measure related to fractal dimension is lacunarity [48], [49], which is used to measure second-order differences in fractal dimension. Keller *et al.* [49] proposed texture measures based on a description of lacunarity, which are derived from the box-counting method. These features may be useful in distinguishing among regions of complex texture, such as the lung

vasculature, and regions that are more homogeneous, such as the axilla and heart regions.

#### D. Image Display and Radiologist's Determination of Class Membership

A set of X-Windows based software tools were developed that allowed the display of radiographic images and the drawing (and saving) of region contours necessary to assign class membership for each pixel. First, the original  $2048 \times 2048 \times 12$ -b radiographic image is selected from a menu, read in, subsampled/smoothed to  $512 \times 512$ , and displayed at 8 b depth (256 gray levels). After the image is displayed, the radiologist used the mouse to draw the contours of the anatomical regions described above and saved that region's contour to a file. For each pixel within or on the contour, class membership was assigned to the anatomic class of that contour. In this way, class membership was assigned for each pixel used.

Only one radiologist was used to draw the anatomic regions as they are general and relatively obvious regions. Although we feel that interobserver variability is very small, studies are currently underway with multiple radiologists to measure possible variability among radiologists in this manual segmentation task.

#### E. Feature Extraction and Creation of Pixel Observations

It was necessary to create pixel observations with all of the associated information: location, class membership, and values of feature variables. The software tools described above also were used in the feature extraction step. In this step, the reduced size ( $128 \times 128 \times 12$  b) image was read in, then an anatomic region's contours were read in from the file saved above and scaled to the reduced format. For each pixel (or a subset of pixels—see below) within that contour, all of the candidate features described above were extracted for each sampled pixel. The extracted features were associated with the pixel location ( $x$  and  $y$  coordinates) and class membership to form the required pixel observation; the observations were all written to a data file.

In each anatomic region of each image, pixels at every fourth row and every fourth column were used to create the observations. The sampling was done to reduce the total amount of data to be processed. Some pixels were excluded from any of the anatomic regions; for example, areas that contained markers or patient identification labels were excluded. The result was that over 750 pixel observations were obtained from each image, with the observations being distributed among the anatomic classes.

#### F. Feature Selection and Creation of Data Sets

The stepwise discriminant analysis (see the Appendix) was carried out using observations as the inputs. Each observation consisted of the parameters described above: pixel class, pixel location, and feature values calculated for that pixel (in this case, all 59 were calculated). Observations from each of the 33 images were combined to form the full data set for feature selection. This consisted of 25 550 observations, distributed

among the classes with the following proportions: hrt—34.2%, lun—41.9%, axl—14.6%, bnk—5.2%, and bkd—4.1%. This data set was used as input to the stepwise discriminant analysis implemented in procedure STEPDISC [6] of the SAS statistical software package (SAS Institute, Inc., Cary, N.C.). This analysis created a list of selected features and showed them in the order in which they were selected. The first features selected possessed the most discriminatory power among the classes.

To evaluate the impact of feature selection on classifier performance and processing time requirements, three data sets were created and evaluated. The first data set consisted of all 59 candidate features—no feature selection was performed. For the second data set, the eight best features were selected. The third data set was a set of eight features selected at random from features that were not among the top eight features in the stepwise discriminate analysis. Thus, training and test data sets were created for each of these sets of features.

#### G. Training and Testing Data Sets

Because both classifiers require supervised training techniques, three training data sets were created corresponding to the three sets of features. Each data set consisted of pixel observations obtained from the 17 training images. For each anatomic region within each image, pixels at every fourth row and every fourth column were sampled and the appropriate features (full 59, best eight or lower eight) extracted. The resulting training data sets each consisted of 13 358 observations with the following class breakdown: hrt—35.2%, lun—39.9%, axl—15.6%, bnk—5.6%, and bkd—3.7%. The identical training data sets were provided to the linear discriminator and neural network for training.

Three test data sets were created that consisted of pixel observations from the 16 test images. Again, the appropriate features for each set were extracted from the sampled pixels and associated with pixel class and location to form the desired testing pixel observations. The resulting test data sets consisted of 12 192 observations with the following class breakdown: hrt—33.1%, lun—44.0%, axl—13.6%, bnk—4.8%, and bkd—4.5%.

#### H. Training and Testing the Classifiers

Two pattern classifiers were used: a linear discriminator and a feedforward neural network using the backpropagation learning rule. Both of these classifiers require supervised training techniques. Each classifier was trained with each of the three training data sets described above.

For all of the following neural network experiments, the number of input nodes was determined by the number of feature variables used—either eight (for the eight best or random eight data sets) or 59 (for the data set of all features)—and the number of nodes in the output layer was five (one for each class). Thus, the decision rule was to assign a pixel to the class corresponding to the output node with the highest value.

Other parameters used in the feedforward, backpropagation network were: 15 nodes in the hidden layer, a learning rate of 0.10, and a momentum term of 0.05. The number of nodes

in the hidden layer, the learning rate, and the momentum term were all determined empirically by a set of experiments in which each term was varied individually and its effect on performance was determined. While the network was relatively insensitive to small changes in these selected values, performance was slightly better at these values. The network was trained for 2000 iterations of the training data as it was empirically determined that performance (as measured by the mean squared error at the output nodes) was stable at this point. The order of training data was randomized for each iteration.

Following the training phase, each of the classifiers were put through a test phase. Test observations from the three test sets (eight selected features, the full 59 feature set, and the eight randomly chosen features) were used. The output decisions of the different classifiers are compared to the results obtained from the radiologist, who is considered to be the standard. The result of these comparisons is a series of classification matrices in which the classifiers output is compared to the radiologist's decisions. The overall percent correct is calculated by dividing the total number of correct pixel observations (pixels with class membership equal to that of the radiologist's classification) from each image by the total number of observations (pixels of all classes) in the test set. Comparisons based on the total percent correct as well as processing time used were carried out for each feature data set. To test whether differences in overall percent correct were statistically significant, we obtained an asymptotic estimate of the  $z$ -statistic and calculated the corresponding  $p$ -value. This was done for comparisons between each feature set and classifier combination.

#### I. Adding Spatial Information After the Classification

In this section, a postprocessing step is described in which results from the pattern classifiers are reclassified according to a set of heuristic rules based on spatial information. The methods to this point provide a segmentation based solely on locally calculated features and no spatial information has been utilized. Thus, there are some pixels that are misclassified and their misclassifications may be obvious only because of the spatial location—for example, pixels classified as background that are located in the middle of the lung. Rules were formulated to identify many of these obvious misclassifications and were applied to reclassify results from the pattern classifiers.

The application of these heuristic rules to add spatial information is carried out in two steps: 1) identification of several zones within the image where only one or two classes should exist, and 2) the results from the classifiers are examined by spatial location, compared to the classes that should exist at that location, and reclassified only if the classification is obviously incorrect. If the location of a pixel dictates that it can belong to only one of two classes and the classification result is not one of those two classes, then this pixel is considered to be obviously misclassified; rules are then applied to reclassify that pixel. If the pixel does belong to one of the two allowable classes for that location, then no change is made.

To identify the zones, the image is thresholded to approximately identify the lung and background regions. The

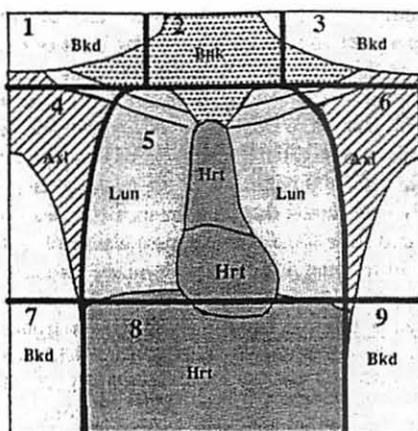


Fig. 2. Approximate definition of zones used to add spatial information.

TABLE II  
DEFINITION OF ANATOMIC CLASSES ALLOWED BY ZONE

Zone	Anatomical Classes Allowed
1	Axilla and Background (Axl, Bkd)
2	Base of Neck and Background (Bnk, Bkd)
3	Axilla and Background (Axl, Bkd)
4	Axilla and Background (Axl, Bkd)
5	Heart and Lung (Hrt, Lun)
6	Axilla and Background (Axl, Bkd)
7	Axilla and Background (Axl, Bkd)
8	Heart and Lung (Hrt, Lun)
9	Axilla and Background (Axl, Bkd)

approximate boundaries of the lungs are found by first locating the heart as the region above threshold in the central part of the image. For each row, the search proceeds laterally from the central region (each side is done independently) to locate the first pixel below threshold—this is the medial lung border. From the medial border, the search continues laterally to find the last pixel below threshold—this is the lateral lung border. This is repeated for each row until there is no area below threshold when searching from the central region—this criteria locates the top and bottom borders of the lung. Using the top and bottom of the lungs as well as the lateral borders of the lungs as boundaries, the image is divided into the zones shown in Fig. 2. In each zone, only two anatomic classes are allowed as listed in Table II. Because the zones are approximate and based on heuristic rules, pixels that are near the zone boundaries are excluded from the reclassification step.

In the reclassification step, results from the pattern classifiers are examined by pixel location; the class predicted for that pixel is compared to the allowed classes for that location (Table II). If the pixel class is one that is allowed for that location, then no change occurs. If the pixel class is not allowed for that location, then an obvious misclassification occurs and another set of heuristic rules are applied to reclassify that pixel. These rules reclassify the pixel to the allowed class with which the pixel's current class is most often confused. These rules were determined empirically and are different for each zone.

As an example, zone 5 is the central region which contains only lung and heart classes (see Fig. 2). If a pixel located

TABLE III  
EXAMPLE OF RECLASSIFICATION RULES FOR PIXEL IN ZONE 5 OF FIG. 2

Zone	Rules
5 (Hrt/Lun allowed)	if pixel class is Axl, then reclassify to Hrt if pixel class is Bnk, then reclassify to Hrt if pixel class is Bkd, then reclassify to Lun

TABLE IV  
STEPWISE DISCRIMINANT ANALYSIS RESULTS

Rank	Feature
1	Percent (Gray level expressed as % of histogram)
2	Ine2 (SGLDM measure Inertia at d=2, angle = 45)
3	Entv4 (SGLDM measure Entropy at d=4, angle = 90)
4	Glavg (Pixel gray level divided by image average)
5	Eh4 (SGLDM measure Energy at d=4, angle = 0)
6	Hv4 (SGLDM measure Homogeneity at d=4, angle = 90)
7	Inw2 (SGLDM measure Inertia at d=2, angle = 135)
8	Hne4 (SGLDM measure Homogeneity at d=4, angle = 45)
9	Lac4 (Fractal-based measure Lacunarity at scale of 4)
10	Lac2 (Fractal-based measure Lacunarity at scale of 2)
11	Eh2 (SGLDM measure Energy at d=2, angle = 0)
12	EntH4 (SGLDM measure Entropy at d=4, angle = 0)
13	Ev2 (SGLDM measure Energy at d=2, angle = 90)
14	G1 (Gray level of the pixel)
15	Entn2 (SGLDM measure Entropy at d=2, angle = 45)
16	Lac8 (Fractal-based measure Lacunarity at scale of 8)
17	Hh2 (SGLDM measure Homogeneity at d=2, angle = 0)
18	Hv2 (SGLDM measure Homogeneity at d=2, angle = 90)
19	Entw2 (SGLDM measure Entropy at d=2, angle = 135)
20	EntH2 (SGLDM measure Entropy at d=2, angle = 0)

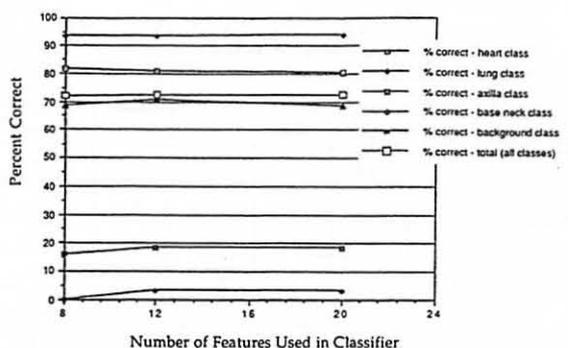
within zone 5 has already been classified as either lung or heart—the allowed classes for this region—then no change occurs. If a pixel located within zone 5 has been classified as background—a class not allowed for this region—then this pixel has obviously been misclassified. It is most likely being confused with lung and, therefore, will be reclassified to lung. Similarly, if a pixel in zone 5 has been classified as axilla or base of neck, it will also be reclassified (to heart). A summary of the rules for zone 5 is shown in Table III.

### III. RESULTS

#### A. Feature Selection Results

The results from the stepwise discriminant analysis are shown in Table IV. This table shows the first 20 variables selected. As one might expect, the first (Percent) and fourth (Glavg) features are related to image gray level, while the other six of the top eight are various spatial gray level dependence matrix (SGLDM) measures. Fractal-based measures show up as the ninth and tenth in the list and no other major feature groups enter into the top 20.

Because there are no clear rules about how many variables should appear in the selected feature list, a simple performance test was carried out. For this test, the linear discriminant was trained on observations from 10 images and tested on observations from a different set of five images. This train and test sequence was repeated using the top eight feature variables, the top 12 feature variables and, finally, the top 20 feature variables. In each situation, the percent correct for each class (number of observations correctly classified for that class/total number of observations for that class)



Percent correct for class =  $\frac{\text{number of observations correctly classified in this class}}{\text{number of observations in this class}}$

Percent correct total =  $\frac{\text{number of observations correctly classified in all classes}}{\text{number of observations in all classes}}$

Fig. 3. Percent correct for each class and total for the linear discriminator as a function of the number of features used.

TABLE V  
LIST OF LOWER EIGHT RANDOMLY PICKED FEATURES

Number	Feature
1	Gy3 (Sobel gradient of 3x3 in y-direction)
2	Disper (Dispersion measure in 5x5 neighborhood)
3	AvLaws3 (Laws' low correlation texture mask averaged over 3x3 neighborhood)
4	Entne2 (SGLDM measure Entropy at d=2 and angle = 135)
5	Ev2 (SGLDM measure Energy at d=2 and angle = 90)
6	Hh4 (SGLDM measure Homogeneity at d=4 and angle=0)
7	Iv4 (SGLDM measure Inertia at d=4 and angle = 90)
8	Frac (Fractal dimension of 8x8 neighborhood)

as well as the total percent correct (number of observations correctly classified for all classes/number of observations from all classes) are calculated. The results are reported in Fig. 3. This figure shows that for this training and test situation, there is almost no difference between using the top eight, 12, or 20 features.

We concluded that there was no advantage in using more than the top eight features and that significant computation time and data storage savings could be realized by using only these features. Therefore, the first eight features listed in Table IV represent the selected feature set.

The lower ranked data set was also created for comparison. This is a set of eight features selected at random from lower ranked features. These are listed in Table V.

#### B. Segmentation Results for Each Classifier/Each Feature Set

Each classifier makes a decision about the class membership of each input observation. This classification decision is compared to the known class in a classification matrix. These matrices are formed for each image. The results from all 12 162 of the test observations are summed to form one classification matrix for each classifier and for each feature set. Tables VI and VII show the classification matrix results from the neural network for both the full 59 feature data set and the eight best feature data set. From the matrices, the

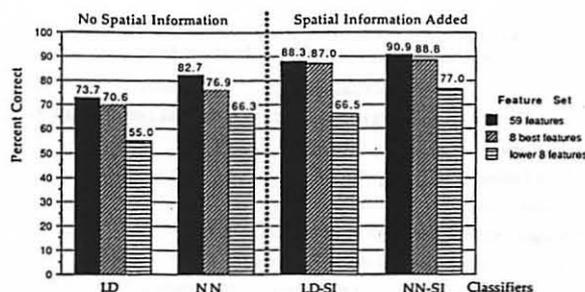


Fig. 4. Overall percent correct for each classifier and each feature data set before and after adding spatial information. (LD = linear discriminator; NN = neural network; LD-SI = linear discriminator after adding spatial information; NN-SI = neural network after adding spatial information.)

TABLE VI  
CLASSIFICATION MATRIX RESULTS FOR NEURAL NETWORK USING FULL 59 FEATURE DATA SET

True Class	Estimated Class					
	hrt	lun	axl	bnk	bkd	total
hrt	3440	231	293	65	2	4031
lun	304	4897	50	98	20	5369
axl	462	74	1055	39	27	1657
bnk	163	89	45	280	3	580
bkd	7	116	4	18	410	555
						12,192

Total Percent Correct =  $(3440+4897+1055+280+410)/12192 = 82.7\%$

TABLE VII  
CLASSIFICATION MATRIX RESULTS FOR NEURAL NETWORK USING BEST EIGHT FEATURE DATA SET

True Class	Estimated Class					
	hrt	lun	axl	bnk	bkd	total
hrt	3144	468	405	14	0	4031
lun	220	4999	111	27	12	5369
axl	664	157	804	32	0	1657
bnk	127	277	89	86	1	580
bkd	0	191	5	20	339	555
						12,192

Total Percent Correct =  $(3144+4999+804+86+339)/12192 = 76.9\%$

total percent correct is calculated to compare each data set and classifier. The comparisons are shown in graphical form on the left-hand side of Fig. 4.

This figure shows that while a feature set consisting of all candidate features provides slightly better performance, it is only 3.1% better for the linear discriminator (73.7% versus 70.6%) and 5.8% better for the neural network (82.7% versus 76.9%). The lower eight features yield much poorer performance for both classifiers: 18.7% worse for the linear discriminator and 16.4% worse for the neural network. Fig. 4 also shows that for this problem, the neural network performs

TABLE VIII  
RELATIVE PROCESSING TIME FOR TRAINING AND TESTING  
OF EACH CLASSIFIER FOR EACH FEATURE DATA SET

Data Set	Linear Discriminator*	Neural Network**
	Training Time/Testing Time	Training Time/Testing Time
All Features	7.6 / 2.3	15510 / 8.1
Best 8 Features	2.2 / 1.0	3044 / 1.0
Lower 8 Features	2.2 / 0.98	3044 / 1.0

\* Run using SAS on IBM ES/9000 mainframe.

\*\* Run using custom neural network on Sun Sparcstation2.

Note: Relative processing time is the processing time (including CPU and disk I/O time) for each task, normalized to the processing time required for the testing of the Best 8 Features. Each classifier is normalized separately as they are run on different machines.

Testing time is the time required to calculate results for all 16 test images. Training time is time required to perform all training iterations (1 iteration for linear discriminant, 2000 iterations for neural network).

slightly better than the linear discriminator (82.7% versus 73.7% for the 59 feature data set and 76.9% versus 70.6% for the best eight feature data set). Each of these differences is significant at the  $p < 0.001$  level.

The processing performance for each of the classifiers (as measured in relative processing time) is shown in Table VIII. For the linear discriminator, testing using all 59 features is more than twice as computationally expensive than using either the eight best or lower eight features, while training is more than three times as expensive. For the neural network, more features require not only more inputs, but more connections between the input layer and the hidden layer. This results in neural network training taking approximately five times as much processing time for the full data set as for the eight best features. Testing requires eight times as much processing for the full 59 feature data set as for the eight feature data set. While the linear discriminator and neural network classifications were performed on different machines, it is clear that training the neural network is very computationally expensive for a large network using a large training data set with a large number of iterations (2000, in our case). While the classification results are better for the neural network, the computational expense is much higher—especially in training.

### C. Applying the Spatial Information Reclassification Rules

The spatial information heuristic rules were applied to each of the test sets. The original classification results for each feature set and each classifier (from above) were used as inputs to the reclassifier rules and the resulting classification matrices are created. Tables IX and X show the classification matrix results from the neural network for both the full 59 feature data set and the eight best feature data set, after the rules adding spatial information were applied. From the matrices, the total percent correct is again calculated to compare each data set and classifier. These comparisons are shown in graphical form on the right-hand side of Fig. 4. In addition, Fig. 5(a) shows an example of an original digital chest radiograph and Fig. 5(b) shows the classification resulting from the eight best feature set applied to the neural network and after the application of the spatial information heuristic rules.

TABLE IX  
CLASSIFICATION MATRIX RESULTS FOR NEURAL NETWORK USING FULL  
59 FEATURE DATA SET AFTER ADDING SPATIAL INFORMATION

True Class	Estimated Class					total
	hrt	lun	axl	bnk	bkd	
hrt	3750	222	48	0	11	4031
lun	342	4810	63	63	91	5369
axl	48	14	1505	3	87	1657
bnk	19	22	38	498	3	580
bkd	0	0	24	14	517	555
						12,192

$$\text{Total Percent Correct} = (3750+4810+1505+498+517)/12192 = 90.9\%$$

TABLE X  
CLASSIFICATION MATRIX RESULTS FOR NEURAL NETWORK USING BEST  
EIGHT FEATURE DATA SET AFTER ADDING SPATIAL INFORMATION

True Class	Estimated Class					total
	hrt	lun	axl	bnk	bkd	
hrt	3505	471	43	3	9	4031
lun	219	4896	144	45	65	5369
axl	58	27	1410	1	161	1657
bnk	0	24	40	491	25	580
bkd	0	4	19	8	524	555
						12,192

$$\text{Total Percent Correct} = (3505+4896+1410+491+524)/12192 = 88.8\%$$

These results show how reclassification based on spatial information improves the performance of both classifiers. They also show that some differences due to feature selection are maintained, albeit reduced, even with a postprocessing step that improves the overall classification rate. For the linear discriminator, the full feature set is now only 1.3% better than the best eight feature set (88.3% versus 87.0%). For the neural network, the full feature set is only 2.1% better than the best eight features (90.9% versus 88.8%). The lower eight features still yield much poorer performance for both classifiers even after the reclassification: 17.8% worse for the linear discriminator and 13.9% worse for the neural network. Fig. 4 shows that, while the neural network still performs slightly better than the linear discriminator, the difference after reclassification is small (90.9% versus 88.3% for the 59 feature data set and 88.8% versus 87.0% for the best eight feature data set). Though the differences are reduced, they are still significant at the  $p < 0.001$  level.

## IV. DISCUSSION

The results for both classifiers, the linear discriminator and the neural network, show that while the full feature set may give slightly better results (82.7% correct for the neural network on the full feature data set), the results for using a selected subset of features are comparable (76.9% correct for the eight best features). Adding a postprocessing step,

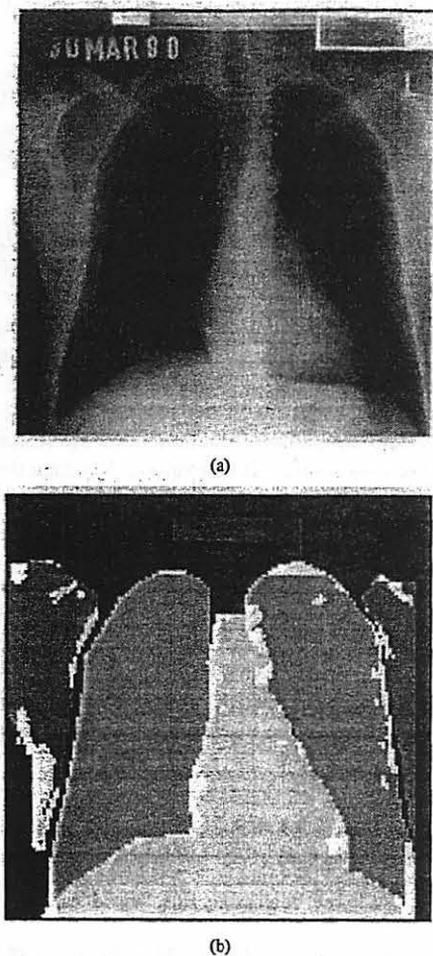


Fig. 5. (a) Original digital chest radiograph image and (b) resulting output image with pixels labeled after classification using neural network with eight best features and application of postprocessing reclassification rules. (Note: Each gray level represents a different class; white pixels represent misclassifications; most of the misclassifications are located near the border between regions.)

such as adding spatial information, improves the performance of all of the classifiers and reduces the differences between feature sets even further (90.9% for the full feature set and 88.8% for the best eight features using the neural network). These results for the selected feature set are obtained while saving computational processing expense in both training (1/3 as much time for linear discriminator, 1/5 as much time for the neural network) and testing (1/2 as much time for linear discriminator, 1/8 as much time for the neural network) the classifiers.

The differences between the eight best features and the lower eight features also show that the feature selection process itself—such as stepwise discriminant analysis—and not just the number of features is important in determining the composition of the subset.

For this particular problem, the neural network classifier gave better overall results than did the linear discriminator.

Before the spatial information was added, the difference was 9% for the full feature set (82.7% versus 73.7%) and 6.3% for the best eight features (76.9% versus 70.6%). After the spatial information was added through reclassification, the differences were reduced to 2.6% (90.9% versus 88.3%) for the full feature set and 1.8% (88.8% versus 87.0%) for the best eight features.

Using this pattern classification approach, the best results for the eight best features occur when the neural network classifier is used and the spatial information heuristic rules are applied. In this case, 88.8% of all pixels (from 16 test images) are correctly classified and 91.2% of all lung pixels are correctly classified. While it is difficult to compare these results directly with other segmentation approaches published in the literature, there has been some recent work along similar lines that sought to segment only the lungs. In the most directly comparable work, Duryea [50] used a rule-based heuristic to find the lung region on reduced-size ( $55 \times 64$ ) chest images and obtained an average percent correct of 95.7% for the left lung and 96.0% for the right lung. Pietka [51] used a rule-based heuristic to obtain 87.5% "correct" delineations of the lung field; where "correct" is defined as the case where the outline shows no penetration of the lung, and all other anatomic regions (head, arms, abdomen) are removed; no comparison to a radiologist was performed. Hasegawa [10] used a convolution neural network to identify lung structures, but reported only a subjective summary of these results. Others have used rule-based segmentation methods to identify specific regions and make anatomic measurements such as heart size and shape [52], delineation of the ribs [53], and the localization of the inter-rib spaces [54]. These efforts reported only the comparison of the anatomic measurements to manual measurements, and not the segmentation results themselves.

The question as to whether or not these results are accurate enough depends upon the demands of the application: if this segmentation is the basis for image enhancement or radiographic equalization, these results may be sufficient; if anatomic measurements are to be made, then the misclassification errors may prevent this from providing accurate enough results. The latter question will be addressed in future studies where multiple radiologists will be used to determine interobserver variability in both manually segmenting the image and in making anatomic measurements.

The difference in processing performance between the two classifiers was shown to be quite large. This was true even though we did not have directly comparable results: the linear discriminator (using SAS) was run on an IBM mainframe and the neural network was run on a Sun Sparcstation2. This difference arose largely because the neural network requires a large number of operations (multiplications, summations, and application of nonlinear functions at each node in each layer), compared to the linear discriminator. The differences were amplified in training the classifiers where the feedforward, backpropagation neural network requires an iterative approach compared to the deterministic approach of the linear discriminator.

The purpose of a feature selection step in the pattern classification process is to identify a subset of variables which

will give results comparable to the full feature data set. The results of this study show, for the problem of segmentation of chest radiographs, the ability of the feature selection step to identify the features which give results that are comparable to the full set while reducing the computational expense of training and testing classifiers. These computational savings can be quite significant for the training and testing of neural networks.

#### APPENDIX

##### BRIEF OVERVIEW OF THE STEPWISE DISCRIMINANT ANALYSIS

In this research, the stepwise discriminant approach [1], [5]–[7] was used to find the subset of candidate features which best discriminates among the classes. Using the forward selection method, the initial step of this approach is to calculate the one way analysis of variance  $F$  statistic (Table I) for each feature. The sources of variation are between classes ( $B$  subscript) and within classes or residual ( $R$  subscript). The  $F$  statistic is the ratio of the between class variation to the within class variation and is used as a measure of the ability of a feature to discriminate among all classes. If the variation in the feature value is large between classes and yet low within the classes, then this would be a good discriminator and would give a large  $F$  value. The feature with the largest  $F$  value is added to the selected feature list.

After the initial step, the ability of a feature to discriminate between classes, given that other features have already been selected, is measured. The partial  $\Lambda$  statistic [55], [5], [6] is calculated to measure the change in ability to discriminate between classes due to the addition of a feature variable to the selected variable list, given the feature variables that are already on that list. The corresponding  $F$  statistic [5] is used to test the significance of the change in  $\Lambda$  resulting from the addition of a feature variable to the selected variable list. This partial (conditional)  $F$  statistic is used as the basis for comparing candidate features, given the list of features already selected.

In the stepwise selection procedure, variables are entered or removed from the selected feature list according to the significance level of the partial  $F$  statistic. At a given step, if the feature on the selected list with the lowest partial  $F$  statistic has a value lower than a specified criterion (called the  $F$ -to-remove), then that feature is removed in that step. Otherwise, the feature on the candidate list with the highest partial  $F$  statistic which exceeds a criterion (called the  $F$ -to-enter) is entered onto the selected feature list. This continues until all of the features on the selected list exceed the  $F$ -to-remove criterion and none of the other features exceed the  $F$ -to-enter criterion.

Obviously, the choice of the  $F$ -to-enter and  $F$ -to-remove can significantly impact the feature selection. Using Monte Carlo simulations to determine stopping rules, Costanza and Afifi [56] recommend a value of approximately 4.0 for the  $F$ -to-enter, to provide a significance level which allows variables of moderate discriminatory power to enter. To keep features from entering and removing in endless loops, the  $F$ -to-remove is set at a value somewhat less than this (e.g., 3.96).

The forward stepwise procedure is as follows (and is described in more detail in [1], [6], and [7]):

- Step 1: For each of the candidate feature variables, calculate the one way analysis of variance  $F$  statistic, along with the degrees of freedom ( $\gamma_B, \gamma_R$ ). The variable with the highest value of  $F$  is selected and put onto a list of selected variables.
- Step 2: For the candidate variables not on the list of selected variables, calculate the partial (conditional)  $F$  statistic for each variable, given the variable(s) already on the selected list.
- Step 3: The variable with the highest partial  $F$  value is chosen and added to the selected variable list, if it exceeds the  $F$ -to-enter criterion.
- Step 4: For the variables on the selected feature list, calculate the partial  $F$  statistic. Find the feature with the lowest value. If that value is less than the  $F$ -to-remove, then remove that feature from the selected list and place back on the candidate list, and repeat Step 4. Otherwise, repeat Steps 2–4.
- Step N: Steps 2–4 are repeated until all features on the selected list exceed the  $F$ -to-remove and none of the other features exceed the  $F$ -to-enter. At this point, no more variables will be entered as there are no variables which provide significant additional information to discriminate among classes.

#### ACKNOWLEDGMENT

The authors wish to thank M. Razavi for his effort in manually segmenting all of the training and test images.

#### REFERENCES

- [1] A. A. Afifi and S. P. Azen, *Statistical Analysis: A Computer Oriented Approach*. New York: Academic, 1972.
- [2] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley-Interscience, 1973.
- [3] P. A. Lachenbruch, *Discriminant Analysis*. New York: Hafner, 1975.
- [4] A. A. Afifi and V. Clark, *Computer-Aided Multivariate Analysis*. Belmont, CA: Lifetime Learning, 1984.
- [5] R. I. Jennrich, "Stepwise discriminant analysis," in *Statistical Methods for Digital Computers*, K. Einslein, A. Ralston, and H. S. Wilf, Eds. New York: Wiley, 1977.
- [6] *SAS/STAT User's Guide, Version 6*, 4th ed., SAS Institute Inc., Cary, N.C., 1989.
- [7] W. J. Dixon, M. B. Brown, L. Engelman, and R. I. Jennrich, *BMDP Statistical Software Manual*. Berkeley, CA: Univ. of California, 1990.
- [8] R. C. Gonzalez and P. Wintz, *Digital Image Processing*. Reading, MA: Addison-Wesley, 1987.
- [9] A. Rosenfeld and A. C. Kak, *Digital Picture Processing*. San Diego, CA: Academic, 1982.
- [10] A. Hasegawa, S.-C. B. Lo, M. T. Freedman, and S. K. Mun, "Convolution neural network based detection of lung structures," in *Proc. SPIE Med. Imag.*, vol. 2167, pp. 654–662, 1994.
- [11] M. W. Vannier et al., "Multispectral analysis of magnetic resonance images," *Radiology*, vol. 154, pp. 221–224, 1985.
- [12] L. D. Schertz, M. W. Vannier, M. H. Gado, and R. L. Butterfield, "Statistical assessment of MR image classification techniques," in *Proc. IEEE Annu. Conf. EMBS*, 1985, pp. 585–592.
- [13] M. W. Vannier et al., "Multispectral magnetic resonance image analysis," *CRC Crit. Rev. Biomed. Eng.*, vol. 15, no. 2, pp. 117–144, 1987.
- [14] M. W. Vannier et al., "Validation of magnetic resonance imaging (MRI) multispectral tissue classification," in *Proc. Int. Conf. Pattern Recognition*, vol. 1, pp. 1182–1185, 1988.
- [15] M. W. Vannier et al., "Validation of magnetic resonance imaging (MRI) multispectral tissue classification," *Comput. Med. Imag. Graphics*, vol. 15, no. 4, pp. 217–223, 1991.

- [16] L. P. Clarke, W. C. Naylor, K. Brown, and M. Silbiger, "MRI tissue classification: Comparison of MLM and ANN systems," *Med. Physics*, vol. 17, no. 4, p. 742, 1990.
- [17] K. Oshio and M. Singh, "Automatic segmentation of magnetic resonance head images using neural nets," in *Proc. IEEE Nucl. Sci. Symp.*, vol. 2, pp. 1430-1434, 1990.
- [18] L. P. Clarke, "MR image segmentation using MLM and artificial neural nets," *Med. Physics*, vol. 18, no. 3, p. 673, 1991.
- [19] M. Ozkan, H. G. Sprengels, and B. M. Dawant, "Multi-spectral magnetic resonance image segmentation using neural networks," in *Int. Joint Conf. Neural Networks*, vol. 1, pp. 429-434, 1991.
- [20] L. P. Clarke, R. P. Velthuisen, L. O. Hall, J. C. Bezdek, A. M. Bensiad, and M. Silbiger, "Comparison of supervised pattern recognition techniques and unsupervised methods for MRI segmentation," in *Proc. SPIE Med. Imag. VI*, vol. 1652, pp. 668-677, 1992.
- [21] N. Asada *et al.*, "Potential usefulness of an artificial neural network for differential diagnosis of interstitial lung diseases: Pilot study," *Radiology*, vol. 177, pp. 857-860, 1990.
- [22] J. M. Boone, V. G. Sigillito, and G. S. Shaber, "Neural networks in radiology: An introduction and evaluation in a signal detection task," *Med. Physics*, vol. 17, no. 2, pp. 234-241, 1990.
- [23] G. W. Gross, J. M. Boone, V. Greco-Hunt, and B. Greenberg, "Neural networks in radiologic diagnosis II: Interpretation of neonatal chest radiographs," *Invest. Radiol.*, vol. 25, pp. 1017-1023, 1990.
- [24] R. H. Silverman and A. S. Noetzel, "Image processing and pattern recognition in ultrasonograms by backpropagation," *Neural Networks*, vol. 3, pp. 593-603, 1990.
- [25] K. J. Cios, K. Chen, and R. A. Langenderfer, "Use of neural networks in detecting cardiac diseases from echocardiographic images," *IEEE Eng. Med. Biol.*, vol. 9, no. 3, pp. 58-60, 1990.
- [26] J. S. DaPonte and P. Sherman, "Classification of ultrasonic image texture by statistical discriminant analysis and neural networks," *Comput. Med. Imag. Graphics*, vol. 15, no. 1, pp. 3-9, 1991.
- [27] J. S. DaPonte, J. A. Parikh, and D. A. Katz, "Detection of liver metastasis using the backpropagation algorithm and linear discriminant analysis," in *Proc. SPIE Med. Imag. V*, vol. 1469, pp. 441-450, 1991.
- [28] Y. Wu, K. Doi, M. L. Giger, and R. M. Nishikawa, "Computerized detection of clustered microcalcifications in digital mammograms: Applications of artificial neural networks," *Med. Physics*, vol. 19, no. 3, pp. 555-560, 1992.
- [29] Y. Wu, M. L. Giger, K. Doi, C. J. Vyborny, and R. A. Schmidt, "Artificial neural networks in mammography: Application to decision making in the diagnosis of breast cancer," *Radiology*, vol. 187, pp. 81-87, 1993.
- [30] G. D. Tourassi and C. E. Floyd, "Artificial neural networks for single photon emission computed tomography: A study of cold lesion detection and localization," *Invest. Radiol.*, vol. 28, no. 8, pp. 671-677, 1993.
- [31] D. E. Rumelhart and J. L. McClelland, *Parallel Distributed Processing, Explorations in the Microstructure of Cognition*, vol. 1. Cambridge, MA: MIT Press, 1986.
- [32] R. P. Lippmann, "An introduction to computing with neural nets," *IEEE ASSP Mag.*, pp. 4-22, Apr. 1987.
- [33] G. P. Gorman and T. J. Sejnowski, "Analysis of hidden units in a layered network trained to classify sonar targets," *Neural Networks*, vol. 1, pp. 75-89, 1988.
- [34] B. Widrow and R. Winter, "Neural nets for adaptive filtering and adaptive pattern recognition," *IEEE Comput. Mag.*, pp. 25-39, 1988.
- [35] R. C. Eberhart and R. W. Dobbins, *Neural Network PC Tools: A Practical Guide*. New York: Academic, 1990.
- [36] D. R. Aberle, F. Gleeson, J. W. Sayre, K. Brown, P. Batra, D. A. Young, B. K. Stewart, B. K. T. Ho, and H. K. Huang, "The effect of irreversible image compression on diagnostic accuracy in thoracic imaging," *Invest. Radiol.*, vol. 28, pp. 398-403, 1993.
- [37] M. F. McNitt-Gray, "Segmentation of chest radiographs using a pattern classification approach," Ph.D. dissertation, UCLA, 1993.
- [38] D. Marr and E. Hildreth, "Theory of edge detection," in *Proc. Royal Soc. Lond.*, vol. 207, pp. 187-217, 1980.
- [39] R. M. Haralick, K. Shanmugan, and I. Dinstein, "Textural features for image classification," *IEEE Trans. Syst. Man Cyber.*, vol. SMC-3, no. 6, pp. 610-621, 1973.
- [40] R. P. Kruger, W. B. Thompson, and A. F. Turner, "Computer diagnosis of pneumoconiosis," *IEEE Trans. Syst. Man Cyber.*, vol. SMC-4, no. 1, pp. 40-49, 1974.
- [41] R. M. Haralick, "Statistical and structural approaches to texture," *Proc. IEEE*, vol. 67, no. 5, pp. 786-804, 1979.
- [42] K. I. Laws, "Textured image segmentation," Univ. of Southern California, USCIP Tech. Rep. 940, 1980.
- [43] A. R. Rao, *A Taxonomy for Texture Description and Identification*. New York: Springer-Verlag, 1990.
- [44] S. Peleg, J. Naor, R. Nartley, and D. Avnir, "Multiple resolution texture analysis and classification," *IEEE Pattern Anal. Machine Intell.*, vol. PAMI-6, no. 4, pp. 518-523, 1984.
- [45] A. Pentland, "Fractal-based description of natural scenes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-6, no. 6, pp. 661-674, 1984.
- [46] S. Majumdar and R. R. Prasad, "The fractal dimension of cerebral surfaces using magnetic resonance images," *Comput. Physics*, vol. 2, no. 6, pp. 69-73, 1988.
- [47] C. C. Chen, J. S. DaPonte, and M. D. Fox, "Fractal feature analysis and classification in medical imaging," *IEEE Trans. Med. Imag.*, vol. 8, no. 2, pp. 133-142, 1989.
- [48] B. B. Mandelbrot, *The Fractal Geometry of Nature*. San Francisco, CA: Freeman, 1988.
- [49] J. M. Keller, S. Chen, and R. M. Crownover, "Texture description and segmentation through fractal geometry," *Comput. Graphics Vision Image Processing*, vol. 45, pp. 150-166, 1989.
- [50] J. Duryea and J. M. Boone, "A fully automated algorithm for the segmentation of lung fields on digital chest radiographic images," *Med. Physics*, vol. 21, no. 6, p. 942, 1994.
- [51] E. Pietka, "Lung segmentation in chest radiographs," *J. Digital Imag.*, vol. 7, no. 2, pp. 79-94, 1994.
- [52] N. Nakamori, K. Doi, V. Sabeti, and H. MacMahon, "Image feature analysis and computer-aided diagnosis in digital radiography: Automated analysis of sizes of heart and lung in chest images," *Med. Physics*, vol. 17, no. 3, pp. 342-350, 1990.
- [53] S. Sanada, K. Doi, and H. MacMahon, "Image feature analysis and computer-aided diagnosis in digital radiography: Automated delineation of posterior ribs in chest images," *Med. Physics*, vol. 18, no. 5, pp. 964-971, 1991.
- [54] G. F. Powell, K. Doi, and S. Katsuragawa, "Localization of inter-rib spaces for lung texture analysis and computer-aided diagnosis in digital chest images," *Med. Physics*, vol. 15, no. 4, pp. 581-587, 1988.
- [55] C. R. Rao, *Linear Statistical Inference and Its Applications*. New York: Wiley, 1965.
- [56] M. C. Costanza and A. A. Afifi, "Comparison of stopping rules in forward stepwise discriminant analysis," *J. Amer. Stat. Assn.*, vol. 74, no. 368, pp. 777-785, 1979.

## Temporal Image Database Design for Outcome Analysis of Lung Nodule

Xiaoming Zhu,\* Ki-Nam Lee\*†, D.L. Levin\*, STC Wong\*, HK Huang\*, and K. Soo Hoo‡

\*Department of Radiology, University of California at San Francisco, San Francisco, CA 94143

†Department of Diagnostic Radiology, Dong-A University, Pusan, 602-103 Korea

‡Bioengineering Graduate Program, UCSF/UC Berkeley, San Francisco, CA 94143

**Abstract:** This paper presents the design of a temporal image database system and its application in thoracic imaging. The design of this information system is based on the client/server architecture. The system consists of a chest imaging database server, a library of image processing modules, a link to the picture archiving and communication system (PACS) archive, and a low end client workstation with Motif-based graphic user interface (GUI). The database system can be used to aid the radiologists in quantitating solitary metastatic lung nodules and in assessing effectiveness of therapeutic treatment procedures for lung cancer or lesions. The GUI allows a user to retrieve any patient study from PACS. After a nodule is visually identified, it will be segmented automatically to obtain relevant features, such as the center of mass, volume, and surface area. Such 3D nodule information, together with the patient textual information, is subsequently organized in the chest imaging database to facilitate outcome analysis.

### Introduction

Current database management systems (DBMS) are designed to manage and query large amounts of textual information. With the rapid growth of digital imaging modalities, medical database management is not so straightforward as digital medical data are now increasingly multimedia. In addition to textual data such as diagnostic reports, physical exams, and laboratory tests, radiological and pathologic images are also important constituents of a patient's record. How to manage this broad spectrum of data types efficiently for clinical applications is a challenging research issue. The new requirements for efficient health care delivery, coupled with the diversity of data types, is demanding a new methodology for medical database management.

Textual data are represented in symbolic objects while image data in pixel values. This fundamental difference of text and image accounts for their variation in information contents and processing methods. The content of a text can be abstracted and indexed in a few keywords. Efficient tools for keyword search are widely available and allow a user to retrieve the right information from a large database efficiently. Images, on the other hand, contain information of higher abstraction. The content of an image must be extracted by image processing tools which are usually slow and inefficient. Furthermore, no corresponding generalized image processing tools exist that enable a user to extract key features reliably from any image.

The purpose of this study is to develop an application-specific temporal image DBMS that can provide clinicians with necessary image processing tools to collect data for outcome analysis. Without these tools, it is formidable for a radiologist to provide a quantitative description of a lesion, such as its volume. The conventional way for a radiologist to infer the size of a lesion is to measure the typical dimension of a lesion and to get a rough estimation of its volume, with certain assumption of its geometrical shape. The temporal image database management system, on the other hand, enables an accurate lesion quantitation by using a computer algorithm to check for every pixel in estimating the volume. Furthermore, it stores the extracted features and other relevant information for future query.

With few parameter adjustments this DBMS can be made applicable for other lesion volume assessment. In this initial phase, however, we develop the system specially for lung nodule volume estimation using spiral CT images. Metastatic lung nodules are nearly spherical lesions having sharp contrast with its surroundings. They can be easily identified visually either on a monitor screen or film. They can be distinguished from blood vessels which are in cylindrical shape.

The result of lesion quantitation is important in several aspects. It can be used for clinicians to assess the effectiveness of ongoing therapeutic treatment, or to determine the characteristics of the lesion. For example, for a solitary nodule, the time interval of doubling its volume is an important indicator to decide whether the nodule is benign or malignant (1). When a nodule is small and is located deep inside the lungs, the usual way to determine a benign nodule is to see if its volume is doubled within a short (<30 days) or very long (>490 days) period of time. For metastasis of lung cancer, the effect of drug treatment is reflected in the variation of the total nodule volume. If a drug is effective, it should be able to reduce or to halt the division of the cancer cells and thus the total volume of the malignant nodules should not increase during the treatment period.

## System Architecture

The design of the temporal medical image database is similar to the architecture proposed in [Wong and Huang](2). As shown in Fig. 1, the design is based on a three-tiered client/server architecture. The image database server (IDBS) is a centrally located core of application programs for accessing, processing and managing CT chest images and associated textual reports from radiological and hospital information systems stored in the central data archive of the UCSF PACS. This middleware database server, resided in between PACS and the client medical workstation, consists of an image processing engine, a relational database engine (in our case, Sybase, Emeryville, CA), and a database containing processed multimedia chest data for temporal lung nodule analysis.

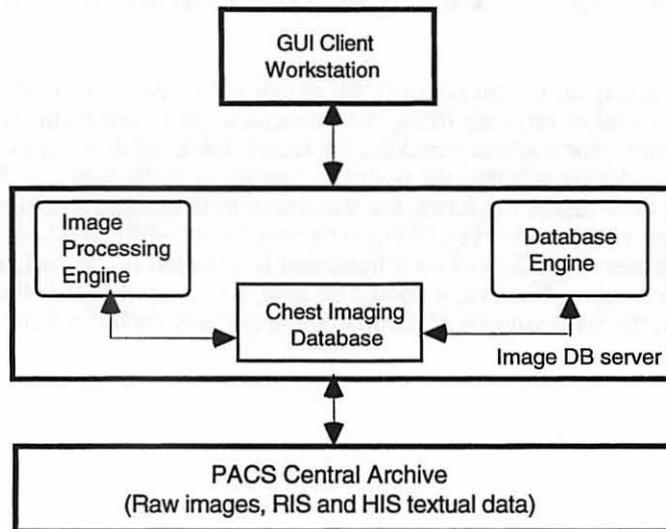


Fig. 1. The 3-tiered client server architecture of the temporal chest imaging database system.

Medical images and text are retrieved from the picture archiving and communication system (PACS) central archive in the Department of Radiology of UCSF. UCSF PACS is developed to manage data acquired from various digital imaging modalities, such as computed tomography (CT), magnetic resonance imaging (MRI), and computed radiography (CR), located in various locations of UCSF hospitals in the San Francisco Bay Area. PACS automates medical imaging handling and communication in a hospital and offers to maintain the quality of healthcare services on timely bases. With PACS, digital images can be quickly retrieved from central archive and be sent across the computer network to remote display workstations for diagnosis, bypassing a lot of middle steps that is necessary in a film-based environment. By providing a digital environment, PACS allows computer technology to become more accessible in healthcare. It not only provides efficiency and convenience, but also defines new functionality that is otherwise not possible in the film-based environment.

To retrieve any patient data from PACS, one needs to specify either a patient name or a patient identification number (ID). A summary of archived patient studies will be presented to a user after the patient name or ID are entered. Based on the summary, the user can select a particular study into the temporal image database after retrieving it from the PACS.

## Image Processing Tools

Different from traditional DBMS, an image database system integrates sophisticated image processing routines into its database server in order to manipulate pixel based image data. The set of image processing tools developed for the lung nodule application is discussed as follows:

**Identify 3-D nodules:** Regions of interests (ROI) can be quantitated and extracted from 3-D CT image data and stored in the database for later retrieval and analysis. In this work, we developed the tools to accomplish such tasks and emphasize on user friendliness and easy to use.

There are several ways to do image segmentation. For lung metastasis, the nodule has a fairly large CT number (in the range of Hounsfield unit 0 or above) compared with its air-like surroundings (-900). Therefore, the pixel values provide adequate information to distinguish the lesion from its surroundings. Complication occurs, however, when a nodule attaches with other tissue such as blood vessels, thoracic wall, or another nodule (they all have similar CT numbers). In such circumstances, a mechanism to edit out the connection manually is implemented in the software.

Image segmentation is carried out after a user places the mouse pointer within a nodule and clicks the left mouse key. Suppose a user identifies a nodule at slice  $n$  and the mouse position  $(x_m^n, y_m^n)$ , the software first extracts a subwindow (51 by 51 pixels, which is usually 3 to 5 times the dimension of a nodule) of the image centered at the mouse position  $(x_m^n, y_m^n)$ . It then attempts to find all closed contours  $C_m^n$  in the 51 x 51 subwindow at a predetermined pixel value of  $L=-500$  in slice  $n$ . After all contours are identified, the software attempts to find which contour encloses the mouse pointer position  $(x_m^n, y_m^n)$ . If there is none, then the mouse pointer is not considered to be within a nodule and the calculation stops.

If contour  $C_m^n$  was identified to enclose the mouse pointer, then segmentation proceeds to the next slice. We first work in minus direction, in slice  $n - 1$ . The center of the subwindow in slice  $n - 1$  was determined by the center of mass of all points enclosed by contour  $C_m^n$ . After all contours in the subwindow of slice  $n - 1$  were identified, the software inspects each contour to find out if it is in contact with contour  $C_m^n$  in slice  $n$ . We define contact such that if a point  $(x, y)$  is both enclosed by contour  $C_i^{n-1}$  and contour  $C_m^n$ , then contours  $C_i^{n-1}$  and  $C_m^n$  are in contact. All contours that are in contact with  $C_m^n$  are considered as part of nodule extending to slice  $n - 1$ .

After they were all identified, we proceed to slice  $n - 2$ . The center of the 51 x 51 subwindow in slice  $n - 2$  is located at the center of mass of all points enclosed by all contact contours in slice  $n - 1$ . If any contour in slice  $n - 2$  is in contact with any contact contour in slice  $n - 1$ , we defined it as contact contour and consider it as part of the nodule in slice  $n - 2$ . Similar calculation proceeds until we reach the slice that contains no contour at the specified pixel level or all its contours are not in contact with any contour in previous slice. After all slices in the minus direction have been examined, the software proceeds to work in the other direction, first in slice  $n + 1$ , then  $n + 2$ , etc.

**Accuracy Assessment:** To assess the accuracy of the image processing tools, we applied our image processing tools to a phantom study. Inside the phantom we placed 12 spherical objects which have similar CT number as the lung nodule. The size of the objects varies from 3 mm to 20 mm in diameter which is also roughly the expected size of the nodule. The phantom was scanned with 5 mm collimation using a GE 9800 Hi-Speed spiral CT scanner (GE Medical

Systems, Milwaukee, Wis) at 250 mAs and 120 KV. The images were reconstructed in 1 mm and 5 mm slice spacing, respectively. Table 1 shows the result of the phantom study. In the first column, we listed object number. The capital letter, L or R, in front of a number indicates whether the object is placed in the left or right lung of the phantom. The second and the third columns list the diameter and the corresponding volume calculated using the diameter of the object. The fourth column shows the volume calculated by the image processing tools. The last column gives the relative error of the calculation.

**TABLE 1A Phantom Study Result  
5 mm collimation with 1 mm reconstruction**

number	D (mm)	Real Vol (mm <sup>3</sup> )	Cal. Vol (mm <sup>3</sup> )	err %
R-1	19.0	3590.9	3584.5	0.2
R-2	12.7	1066.1	1049.9	1.5
R-3	9.5	445.3	434.3	2.5
R-4	6.3	130.9	137.9	5.3
R-5	6.3	130.9	123.8	5.4
L-1	19.0	3590.0	3519.7	2.0
L-2	12.7	130.9	1032.0	3.2
L-3	9.5	445.3	424.6	4.6
L-4	6.3	1066.1	120.1	8.3
L-5	6.3	130.9	126.2	3.6

**TABLE 1B Phantom Study Result  
5 mm collimation with 5 mm reconstruction**

number	D (mm)	Real Vol (mm <sup>3</sup> )	Cal Vol (mm <sup>3</sup> )	err %
R-1	19.0	3590.9	3692.5	2.8
R-2	12.7	1066.1	1102.7	3.4
R-3	9.5	445.3	437.4	1.8
R-4	6.3	130.9	171.0	30.6
R-5	6.3	130.9	120.1	8.3
L-1	19.0	3590.0	3417.8	4.8
L-2	12.7	1066.1	1035.6	2.9
L-3	9.5	445.3	431.0	3.2
L-4	6.3	130.9	122.1	6.7
L-5	6.3	130.9	126.2	3.7

Table 1 only lists the results from 10 nodule studies. The missing two nodules all have diameter 3 mm and the image processing tools are not designed to measure the volume of such small nodule. As will be discussed, small nodules do not contribute significantly to the total volume estimate. For nodules of 6 mm or large, the accuracy of the software for nodule quantitation is quite acceptable, with error margin well within 10 percent except for one case (R-4, Table 1B) which has the relative error as large as 30 %. The diameter of the sphere for R-4 is 6.3 mm, comparable to slice spacing of 5 mm. Since small nodules do not contribute significantly to the total volume estimate, such kind of error may not cause any significant problem in total volume estimate.

### Database Design

Images in the temporal chest imaging database are stored according to image slices. The first two columns of each image table contain patient ID and study number. Each of the rest column in the image table contains an image slice appended with a header file of 1536 bytes. The header file contains detailed description of the image slice. Since the image processing tools process each slice at a time, this structure helps to optimize the image retrieval process.

Additional tables store the processed results, such as nodule position, volume, and its link to nodules in previous study and the later study. Table 2 lists an instance of such the DB table.

**Table 2 An Instance of DB table containing processed nodule information**

Column name	Type(Length)	Comments
pat_id	char(12)	patient ID
study_no	char(8)	Study Number
l_name	char(25)	Nodule_Name (xc, yc, zc)
xc	int	Nodule Center of Mass in X
yc	int	Nodule Center of Mass in Y
zc	int	Nodule Center of Mass in Z
xd	int	Nodule width in X
yd	int	Nodule width in Y
zd	int	Nodule width in Z
vol	float	Volume of the nodule
sur	float	Surface Area
mask	image	Nodule Mask
p_name	char(25)	Nodule name it associates with the previous study
n_name	char(25)	Nodule name it associated with the next study

Suppose a table named Nodule1 contains the nodule information for a patient with patient ID = 1234567 and study number = 1994\_2, one can execute the following SQL query to the database for the volume information:

```
select sum(vol) from Nodule1 where pat_id='1234567' and
study_no='1994_2'
```

After the SQL is executed, the chest image database will return the following result, giving the total volume of the patient study:

```
-----
68278.240540
```

A more complicated query can be executed, for example, to find the total volume of the nodules that are present in the previous study, thus allowing us to assess the information how the nodules grow between two studies.

```
select sum(vol) from Nodule1 where pat_id='1234567' and study_no
= '1994_2' and p_name != NULL
```

This time, the database returns the following result:

```
-----
60961.405400
```

If we make another query in the previous study (1994\_1) for the total volume of nodules that are also present in the next study,

```
select sum(vol) from Nodule1 where pat_id='1234567' and study_no
= '1994_1' and n_name != NULL
```

We get the total volume:

```
-----
20742.445060
```

which enables us to find out that for those nodules that are present in both studies, their volume increases by a factor of 3 in the second study. The results from two queries indicate that the current patient treatment plan is not effective to control the disease. We should point out there that above results are derived from a real patient study that will be discussed in the later section.

## Graphic User Interface

Fig. 2 shows the layout of the GUI software. There are two menu bars at the top. Under menu bar **File**, one is given the options to add a new image file into the database; to delete an existing image file from the database; and to list and to select a particular image file for display. Under menu bar **Nodule**, one has the option to delete a nodule record that contains 3D information of a nodule; the option to display an image slice that contains a particular nodule of interest; and the option to print the total volume of the identified nodules in a text window below the displaying area. Below the menu bar is the main display area. When a particular patient is selected, the image of a predetermined spiral CT slice is displayed there. If one clicks the right mouse button inside the display area, the mouse position coordinates and the pixel value at that position are displayed in the text window right below the display area. Since the spiral CT images usually have gray scale of 12 bits, a user is given the choice to select window and level using the scales below the text window. The third scale in the right allows the user to select a particular image slice for display.

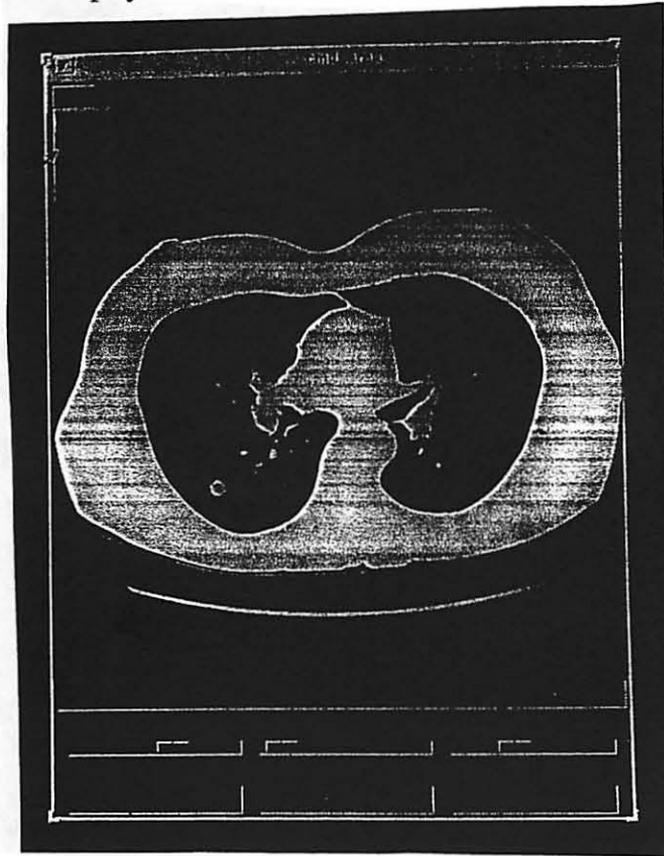


Fig. 2. The layout of the GUI of the software

When a nodule is identified visually, the radiologist initiates the image processing tools by moving the mouse pointer anywhere within the nodule and then clicking the left mouse button. This activates a series of image processing steps to segment out the nodule and to calculate 3D information of the nodule automatically. After the calculation is completed, a popup window will appear to allow the user to inspect the result of automatic segmentation, as shown in Fig. 3.

## Graphic User Interface

Fig. 2 shows the layout of the GUI software. There are two menu bars at the top. Under menu bar **File**, one is given the options to add a new image file into the database; to delete an existing image file from the database; and to list and to select a particular image file for display. Under menu bar **Nodule**, one has the option to delete a nodule record that contains 3D information of a nodule; the option to display an image slice that contains a particular nodule of interest; and the option to print the total volume of the identified nodules in a text window below the displaying area. Below the menu bar is the main display area. When a particular patient is selected, the image of a predetermined spiral CT slice is displayed there. If one clicks the right mouse button inside the display area, the mouse position coordinates and the pixel value at that position are displayed in the text window right below the display area. Since the spiral CT images usually have gray scale of 12 bits, a user is given the choice to select window and level using the scales below the text window. The third scale in the right allows the user to select a particular image slice for display.

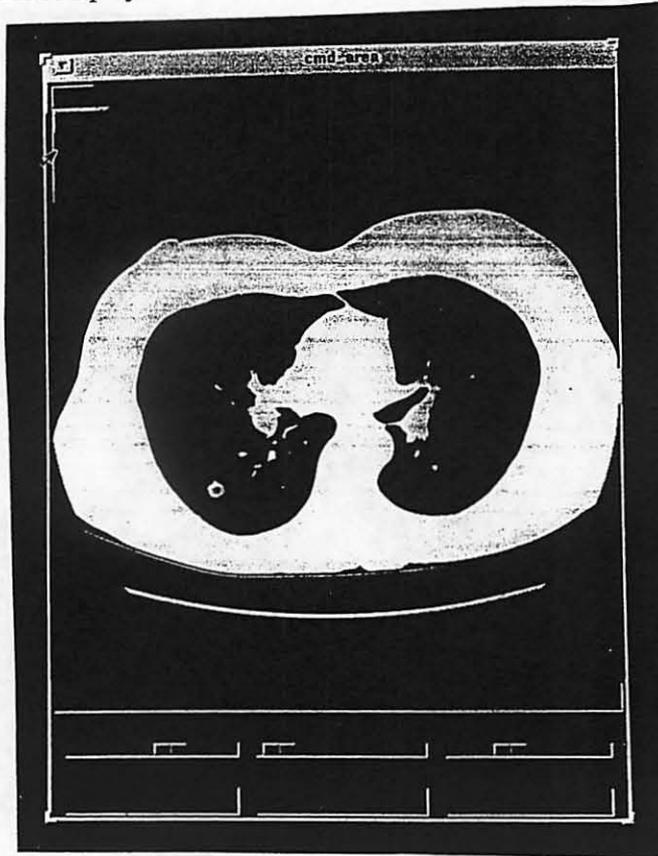


Fig. 2. The layout of the GUI of the software

When a nodule is identified visually, the radiologist initiates the image processing tools by moving the mouse pointer anywhere within the nodule and then clicking the left mouse button. This activates a series of image processing steps to segment out the nodule and to calculate 3D information of the nodule automatically. After the calculation is completed, a popup window will appear to allow the user to inspect the result of automatic segmentation, as shown in Fig. 3.

The top row in the popup window shows the images in 51 by 51 subwindow for different slices to allow a radiologist to examine the original image data. The bottom row shows the masks of segmented nodule for the radiologist to find out if the result is satisfactory or not. If the result is satisfactory, the 3D nodule information will be stored into the chest imaging database. Otherwise, the user can manually correct any segmentation errors and then stores the result into the database. If the result in any slice contains the segmentation errors, the radiologist can click the mouse pointer anywhere inside the slice. A second popup window appears after the mouse click, as shown in Fig. 4. The popup window amplifies the slice the radiologist clicked in Fig. 2. In the window. All pixels belonging to the nodule are marked with a plus (+) sign. A radiologist can easily include or exclude one pixel by clicking the mouse key at the pixel position.

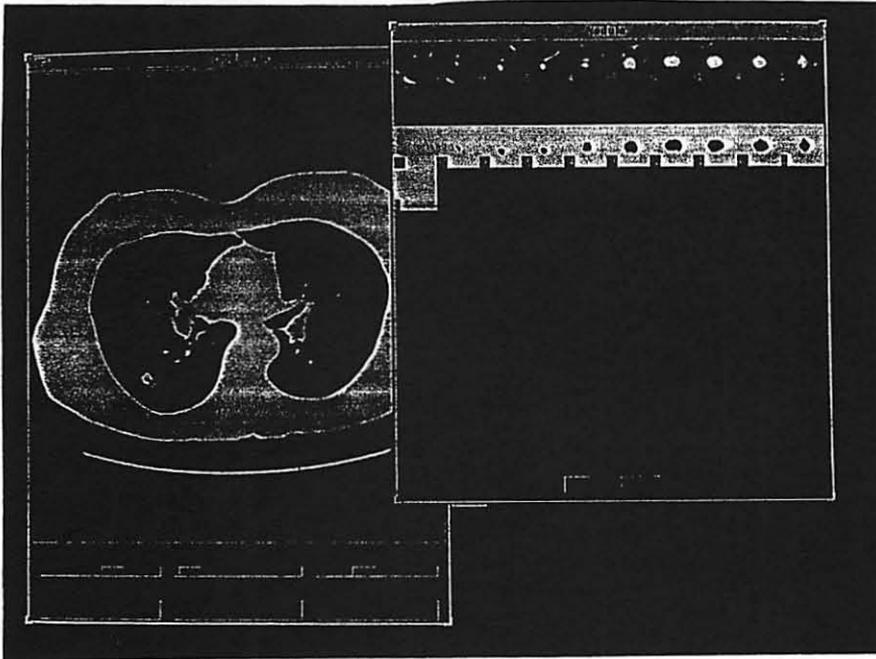


Fig. 3. The first popup window that displays the segmentation results. Top row: different slices of the selected nodule. Bottom row: mask of the segmented nodule for quantity inspection.

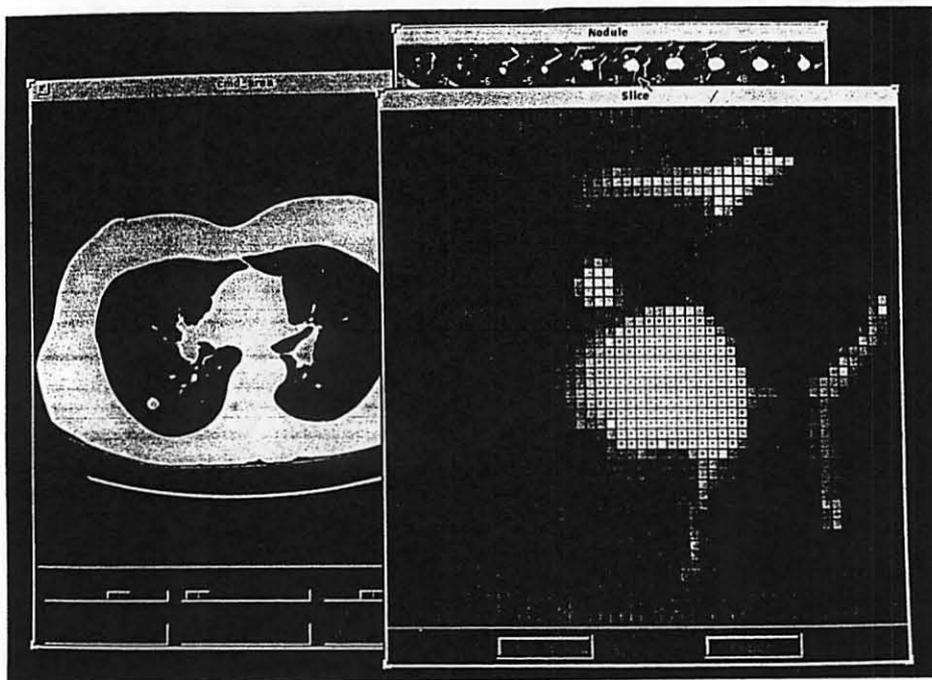


Fig. 4. The second popup window that allows a user to correct the segmentation result on pixel bases. A pixel can be included or excluded by a simple mouse key click anywhere inside the pixel.

To ensure a nodule not to be calculated twice, the processed nodules will turn into red color on display. This feature is found to be helpful for metastasis application since a patient with lung metastasis usually has many nodules. It is difficult for a radiologist to keep track which ones have been processed.

### An Example of Application

**Case History:** We illustrate a patient case of lung metastasis to demonstrate the usefulness of the temporal database system when it was applied to it. The patient is a female with a family history of breast cancer. In 1989, she complained a small lump in her breast which was found benign after aspiration. Two years later (1991), she complained two small masses one in each breast. Excisional biopsy and pathology showed that the right one was breast carcinoma. She was treated by radiation therapy and chemotherapy. In 1993, multiple small lung nodules were delineated in following chest radiographs. Biopsy of the nodules showed they are adenocarcinoma. The patient was under various chemotherapy treatments and was under stable condition. Between September 1994 and March 1995, the patient tried some experimental medication for her cancer therapy. In her March 1995 check up, marked progression of lung metastatic nodules on chest CT and diffuse bone metastasis was noted. After March 1995, the patient received Taxol treatment (chemotherapy) and in June radiation therapy (4500-4000 cGy) with nodule dose on C7-T8 (vertebral body) and C1-C6 (cervical vertebral body). After radiation therapy, adriamycin (chemotherapy) was applied.

**Temporal Assessment:** Since the UCSF PACS is in operation in February, 1994, a total of 5 different spiral CT studies of the patient has been archived. The first archive was dated in August, 1994 while the most recent one was in November, 1995. The first two studies were

scanned at 5 mm collimation, and the images were reconstructed with 1 mm slice thickness. The last three studies were scanned at 7 mm collimation and were reconstructed with 7 mm slice spacing. The change of protocol between the first two and the last three studies is due to the change of the radiologists involved in the case. These studies were retrieved from PACS into the temporal image database management system. Nodule segmentation and volume estimate were performed by a radiologist aided with the software tools. Total volume and the center of mass of each nodule were saved into the database. After all nodules in 5 studies have been segmented, the progression of the lung disease of the patient can be visually assessed in Fig. 5. The horizontal axis plots time. The vertical axis plots both number of nodules (scaled in left) and total volume of the nodules in each study (scaled in right). The volume is measured in unit of cubic millimeter. Below the horizontal axis, we draw two lines to denote the intervals during which the patient was under different treatment plans. In the first interval, the patient was taking an experimental medication. Unfortunately both nodule volume and nodule number grew with time. Total nodule volume is tripled during this experimental period. After March 1995, the patient is under a combination of chemotherapy and radiation therapy which were effective to control the development of the lung cancer. From the figure we note that the total volume of the cancer reduces by a factor of 5 from its peak value near March 1995.

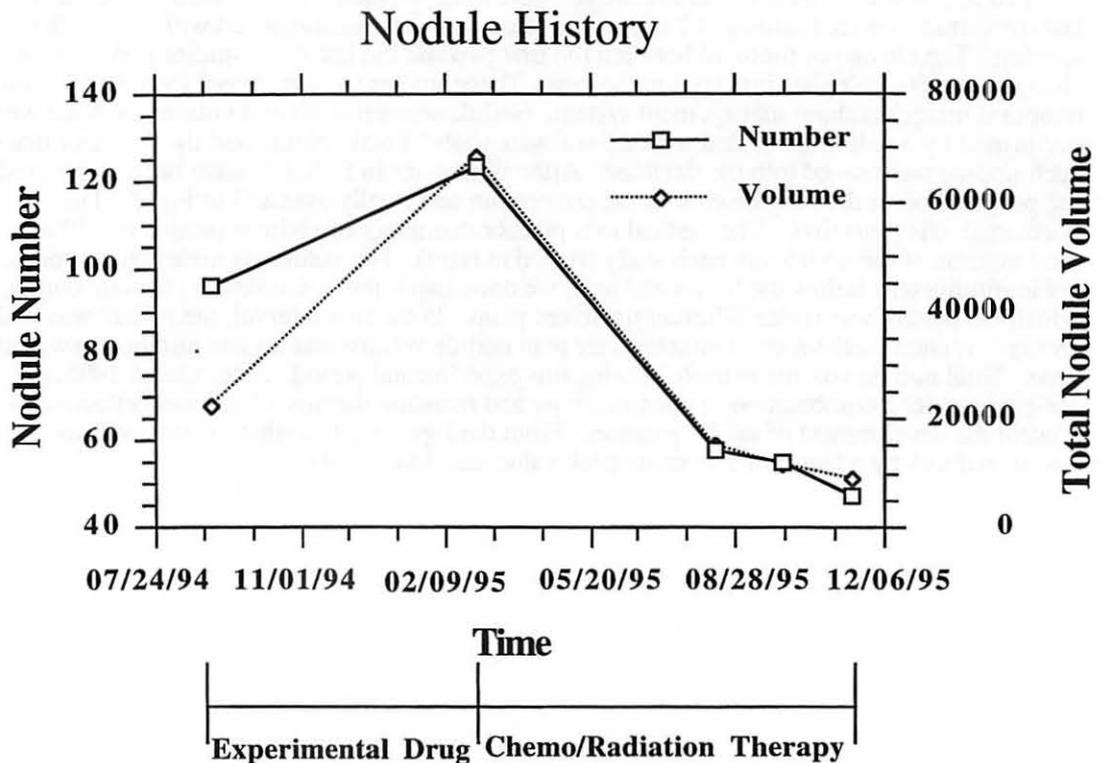


Fig. 5. Temporal history of the patient nodule number and the total nodule volume.

Fig. 6 plots the number of nodules as function of volume and time. The z axis marks the number of nodules. The x axis denotes the volume of the nodules. All nodules with volume between 0 and 100 cubic mm are classified in bin one, all nodules with volume between 100 and 200 cubic mm are classified in bin two, and so on. Bin ten contains those nodules whose volume exceeds 900. The y axis plots the time. From Fig. 6 we note that in the second patient examination, the number of the large volume nodules is significantly increased. This attributes to the enhancement of the total nodule volume found in Fig. 5 around the March 1995 examination.

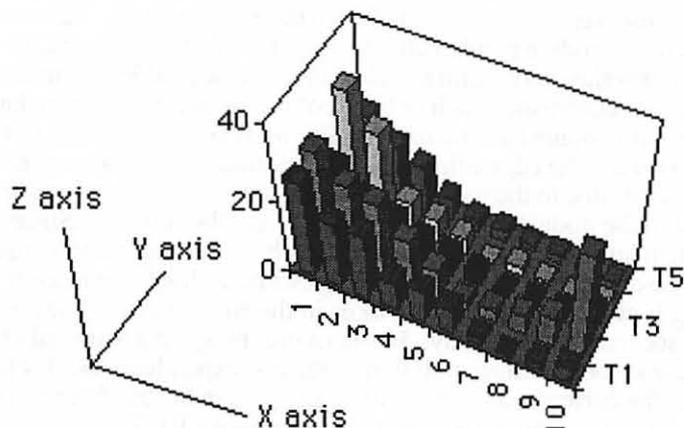


Fig. 6. Number of nodules in 10 different volume ranges in different examination dates. The x axis denotes the volume bins where one denotes volume between 0 and 100 cubic mm and so on. The y axis denotes the patient examination time. The z axis plots the number of nodules in each volume bin and each examination date.

### Intra- and Inter- Observer Variations

In this preliminary study, we asked two radiologists to use the image processing tools to segment images from the patient case previously described. About a month after the first radiologist finished all 5 studies using the software, he repeated the image segmentation with output stored in a separate table in the database. This allows us to study the intra-observer variation. After the first radiologist finished the case twice, the second radiologist joined in to segment the image using the same image processing software, with his output stored in the third table in the database. By comparing the results in three tables in the database, we also obtained preliminary result for inter-observer variations.

**Intra-observer variations:** Table 3 shows the result of the intra-observer variations. The first column lists the patient study number. The second (third) column lists the corresponding results for the first (second) estimate the radiologist made. The last column lists the relative error of the total volume between the two estimates.

**Table 3 Intra-observer variations**

	The First Estimate (number) volume	The Second Estimate (number) volume	Volume Error %
Study 1	(88) 24054	(96) 22520	6.5
Study 2	(118) 68278	(123) 68448	0.2
Study 3	(52) 17517	(58) 15125	14.7
Study 4	(59) 11384	(55) 11627	2.1
Study 5	(40) 8681	(47) 8780	1.1

The number in parenthesis in column 2 and 3 is the number of nodules found in each study, and the number immediately following it is the total volume in the unit of cubic millimeter. If we sum the number in the last column and divide the result by 5 we obtain the average percentage volume error for the intra-observer variation, which is 6%. From columns 2 and 3 we also note that between the two estimates, there is no or even anti-phase relation between the number of nodules and the total

volume. In the first study, for example, the first estimates found 88 nodules and the second 96. The total volume in the second estimate, however, is smaller than that in the first estimate. This is because the causes of the variations in the nodule number and the total nodule volume are different. For nodule numbers, they depend on whether very small nodules, barely discernible on the computer screen, are counted or not (which do not contribute much to total volume as we will discuss later), and on whether two attached nodules are counted as one or two. For total nodule volume, it depends on how to define the edge of each nodule. The edge often changes gradually across several pixels from high CT number to low CT number due to the partial volume effects (3, 4), and sometimes it is hard to define due to the attachment of the nodule with other vessels or another nodule. Since the users can manually correct the segmentation results, the volume of each nodule can vary according to where one defines its edge. In any case, the result from the second estimate should be more reliable because the radiologist was learning how to use the software tools in the first estimate study and he was much more experienced in the second estimate study. Furthermore, he spent additional effort to examine each processed nodule in the second estimate to link it to the corresponding nodules in the studies immediately before and after the current one. Our result, therefore, indicates that the total volume estimate error between a novice and a more experienced user is about 10 %.

**Inter-observer variations:** After the first radiologist finished the two estimates, the second radiologist joined us to use the software to process 5 studies. Based on the results from the two radiologists, the inter-observer variations can be calculated, shown in Table 4. The first column lists the study number. The second column lists the results obtained by the second radiologist. The third and the fourth columns list the percentage volume estimate errors in the inter-observer variations when the results from the second radiologist are compared with the ones from the first radiologist in his first and second estimates, respectively.

**Table 4 Inter-observer variations**

	(number) volume	Volume Error % With First Estimate	Volume Error % With Second Estimate
Study 1	(88) 23994	0.2	6.3
Study 2	(109) 67433	1.2	1.5
Study 3	(49) 15307	13.5	1.2
Study 4	(47) 12255	7.4	5.3
Study 5	(37) 7435.5	15.5	16.6

From the table we note that except for the last study (study 5), the inter-observer variation is comparable to the intra-observer variation, with total volume error less than 10 %. The large variations in study 5 are due to the fibrotic changes after therapeutic treatment. Although in this study, it was agreed by both radiologists that fibroc changes will be excluded from the nodule calculation, it is not easy to determine whether a nodule with fibroc change will be included or excluded. This attributes to the large inter-observer variation in study 5.

## Discussion and Conclusion

**Error Analysis:** As discussed before, partial volume effects tend to blur the edge of a nodule which affects volume estimate. It should be emphasized these effects are more prominent for smaller nodules than large ones. On the other hand, the major contribution to total volume comes from large nodule. Therefore, even though partial volume effects pose some problem in volume determination for individual nodules, its overall effect should not be a major concern for total volume estimate; unless all nodules for a patient are very small.

For simplicity, let us assume a nodule with a spherical shape so that its volume can be expressed as

$$V = \frac{4}{3} \pi R^3 \quad (1)$$

where V is the volume and the R is the radius of the sphere. If there is an error  $\delta R$  in estimating the radius, the volume estimate also has an error  $\delta V$  which relates to  $\delta R$  by taking the derivatives of the above equation:

$$\delta V = 4\pi R^2 \delta R \quad (2)$$

The relative error, by dividing the two equations, is

$$\frac{\delta V}{V} = 3 \frac{\delta R}{R} \quad (3)$$

No matter whether the nodule is large or small, the partial volume effects always cause the edge of a nodule to be blurred in a few pixels. Therefore,  $\delta R$  is fixed and is of the order 1. Once  $\delta R$  is fixed, we learn from above equation that large nodules should have smaller relative error since large nodules have large R. This rule should be intuitive. If the uncertainty in radius determination is a significant part of the whole radius of the sphere, one may expect this can cause large uncertainty in the volume determination. Close examination of the Table 1 allows us to see the trend: the relative error in volume estimates in the phantom study does increase as the nodule size gets smaller. While for large nodules, the volume estimate can often have less than 3 percent relative error.

Because the volume is proportional to  $R^3$ , large nodules contribute dominantly to the total volume estimate. Therefore as long as we can make good volume estimate on large nodules, total volume can be made fairly accurately. In previous paragraph, we have already discussed that the relative errors for large nodule is small. The overall total volume estimate, therefore, are not seriously affected by uncertainties of small nodule estimates.

**Nodule Segmentation Time:** Nodule segmentation time varies by the user and also depends on the complexity of the nodule, slice spacing, and the available computer CPU time. For each individual nodule, it varies between 30 seconds to 5 minutes. One has to spend more time on those complex nodules which are attached to a blood vessel or to costal and mediastinal pleural surfaces, etc. since at this preliminary stage, the software is not sophisticated enough to handle these circumstances fully automatically. For the patient example which has nodules of order 100, it takes a radiologist 0.5 - 1 hour to segment all nodules in a study with 7 mm slice spacing (the last 3 studies). The time a radiologist has to spend on studies with 1 mm reconstruction is much longer, typically in several hours since more slices need to be examined for a nodule of a given size.

Shortening of image processing time is possible by improving image processing algorithm (5), and this is underway. By designing more sophisticated image processing tools, we expect that the segmentation process can be speeded up and the amount of human intervention can be reduced, thus increase the productivity. Further, as database expands, it allows large scale analysis of interpatient data under the similar treatment plans. These topics will be discussed in future.

Since large nodule contributes more to the total nodule volume, another way to speed up the nodule quantitation is to select a few large nodules in the initial patient study and trace them in subsequent ones to find the time evolution of these nodules. By working on only a few (say 10) selected nodules, a radiologist can spend much less time in segmentation. Fig. 7 plots 8 selected nodules that are present in all five studies. From Fig 7. we note that all nodules follow similar temporal evolution as that in Fig. 5

## Selected Nodule Volume

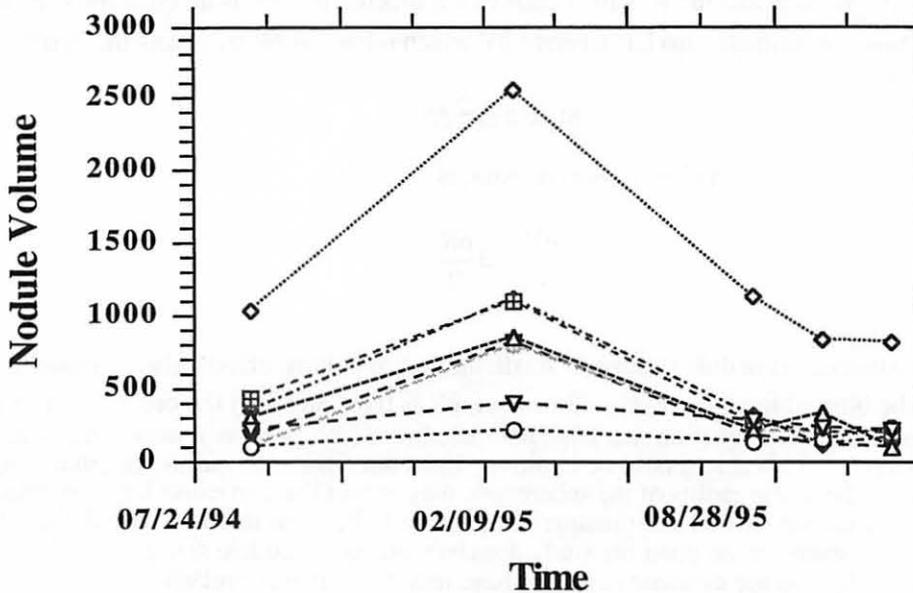


Fig. 7 Temporal volume evolution of selected nodules present in all 5 studies

**Other Issues:** Accurate volume estimate is possible only with digital images. In film-based environment, volume can be estimated by measuring the dimension of a nodule after assuming it is a sphere or ovoid, which in most circumstances is prone to errors since the shape of a nodule is irregular. The estimate is also affected by the window and level setting for the image which determines the edge of the nodule. Our image processing tools, on the other hand, always use the full 12 bits greyscale information to segment a nodule. Therefore, the volume estimate is less affected by the window and level settings for display.

In conclusion, the temporal image database management system we developed helps the radiologists to collect data for outcome analysis. This kind of study is possible only when digital medical images are available which is conveniently provided on-line in the PACS environment.

**Acknowledgments** -- We would like to thank Dr. G. Gamsu of Department of Radiology, Cornell University Medical College and Dr. W. R. Webb of the Department of Radiology, UCSF for their suggestions and clinical help for the project.

## References

1. Fraser, R.G.; Pare, J.A.P.; Pare, P.D; Fraser, R. S.; Genereux, G.P., *Diagnosis of Diseases of the Chest*, Vol 2, 3rd Edition, W.B. Saunders Company, Philadelphia, 1989.
2. Wong, STC; Huang, HK. Design Methods and Architectural Issues of Integrated Medical Image Database Systems. *Computerized Medical Imaging & Graphics*, this issue 1996.
3. Baxter, B.S; Sorenson, J.A, Factors affecting the measurement of size and CT number in computed tomography, *Investigative radiology*, 16(4) 337-341, 1981.
4. Kalender, W. A.; Polacin, A; Suss, C., A comparison of conventional and spiral CT: an experimental study on the detection of spherical lesions, *J. Computer Assisted Tomography*, 18(2) 167-176, 1994.
5. McInerney, T.; Terzopoulos, T. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis, *Computerized Medical Imaging & Graphics*, 19(1): 69-83, 1995.

## Figure Captions

- Fig. 1. The client server architecture of the temporal chest imaging database system.
- Fig. 2. The layout of the GUI of the software
- Fig. 3. The first popup window that displays the segmentation results. Top row: different slices of the selected nodule. Bottom row: mask of the segmented nodule for quantity inspection.
- Fig. 4. The second popup window that allows a user to correct the segmentation result on pixel bases. A pixel can be included or excluded by a simple mouse key click anywhere inside the pixel.
- Fig. 5. Temporal history of the patient nodule number and the total nodule volume.
- Fig. 6. Number of nodules in 10 different volume ranges in different examination dates. The x axis denotes the volume bins where one denotes volume between 0 and 100 cubic mm and so on. The y axis denotes the patient examination time. The z axis plots the number of nodules in each volume bin and each examination date.
- Fig. 7 Temporal volume evolution of selected nodules present in all 5 studies.

# A Hospital Integrated Framework for Multimodality Image Base Management

Stephen T. C. Wong, *Member, IEEE*  
H. K. Huang, *Senior Member, IEEE*  
The University of California, San Francisco (UCSF)

## Abstract

*The trend in healthcare information technology is increasingly digital and multimedia oriented. The next generation of health care information systems will consist of a vast network of heterogeneous, autonomous, and distributed imaging scanners, databases, information systems, knowledge intensive applications, and large quantities of multimedia medical data. A key challenge facing system researchers and builders is to provide a new organizational framework that can integrate this varied collection of resources into what appears to be a uniform and logical conglomeration of data and knowledge store in order to increase the availability of global or previously non-accessible information and to address demanding new information processing requirements for diverse image-assisted medical applications.*

*The purpose of this paper is to present our research towards the development of a hospital integrated framework of multimodality image base management (MIBM) for digital radiology of the future. This evolutionary framework consists of three hierarchical components: a hospital-integrated picture archiving and communication system (HI-PACS), a medical image database system (MIDS), and a set of image-based medical applications that relies on the support of MIDS and PACS. In this paper, we describe the system architecture, guiding principles, and design specifications of HI-PACS and MIDS and illustrate their functions and capabilities with three implemented applications, namely, patient folder workflow, distributed object management, and multimodality imaging studies. In addition, we conclude our findings with a summary of challenges and research directions.*

## 1. INTRODUCTION

Medical images form the cornerstone of patient records and are at the heart of the patient's diagnosis, determination of therapy, and follow-up. They are used not only by the diagnostic and interventional radiologists, but also by medical oncologists, radiotherapists, surgeons, dermatologists, pathologists, neurosurgeons, family physicians, and other medical professionals. The trend in medical imaging is increasingly digital and multimedia oriented [31,46]. The basic motivation is to represent medical images in digital form supporting image transfer and archiving and to manipulate visual diagnostic information in useful and novel ways, such as image enhancement and volume rendering.

In the past two decades, we have witnessed an explosion of primary digital modalities: film scanners, ultrasound, x-ray computed radiography (CR), magnetic resonance imaging (MRI), computed tomography (CT), positron emission tomography (PET), and single positron emission computed tomography (SPECT), to name just a few [48,37]. These modalities, currently constitute 30% of the imaging examinations of the nation, have revolutionized the means to acquire patient images, provide flexible means to view anatomical cross sections and physiological states, and reduce patient radiation dose and examination trauma. The other 70% of examinations on skull, chest, breast, abdomen, and bone are done in conventional x-rays. Different kinds of film digitizers, such as laser scanner, solid-state camera, drum scanner, and video camera, can be used to convert X-ray films into digital format for processing [29]. The new digital modalities, at the same time, have generated a large volume of multimodality image data such that traditional methods employed by the clinicians and health care providers to manage their data with papers and films are no longer adequate. Therefore, the classical way to deal with these data has appeared often sub-optimal with regard to their quality and to

their medical objective which often tend to exclude complementation information from other clinical information systems.

Modality	Image Dimension	Gray Level (bits)	Avg. Size/Exam
Nuclear Medicine (PET & SPECT)	128 × 128	8 or 16	1 ~ 2 MB
Magnetic Resonance Imaging (MRI)	256 × 256	12	8 MB
Ultrasound (US)	512 × 512	8	5 ~ 8 MB
Digital Subtraction Angiography (DSA)	512 × 512	8	4 ~ 10 MB
Computed Tomography (CT)	512 × 512	12	20 MB
Digitized Electronic Microscopy (DEM)	512 × 512	8	Varies
Digitized Color Microscopy (DCM)	512 × 512	24	Varies
Digitized X-rays	2048 × 2048	12	8 MB
Computed Radiography (CR)	2048 × 2048	12	8 ~ 32 MB
Digitized Mammography	4096 × 4096	12	32 MB (one image)

Table 1. Medical images sizes (unit: MB = megabytes).

In addition to the multimodality variation, another key characteristic of medical images is their large data sizes. A two-dimensional (2-D) radiologic image has a size of  $M \times N \times k$  bits, where  $2^k$  equals the gray level range. Table 1 lists the average number of megabytes (MB) per examination generated by medical imaging technologies, where any 12-bit image is represented by 2 bytes in memory. The plain X-ray films of images of the higher resolution requirement can be digitized by  $4K \times 4K \times 12$ -bit digitizers [29]. The size of an image and the number of images taken in one patient examination vary with modalities. Except for DEM and DCM that are pathologic and histologic images of microscopic tissues, all the other modalities are generally classified as radiologic images and are used for diagnosis and treatment planning purposes. In contrast with most other types of biomedical images, such as DCM, a radiologic image is monochrome, and each radiologic examination involves a well defined procedure. For example, one examination (about 40 image slices) of x-ray computed tomography (CT) with each image slice size of  $512 \times 512 \times 12$  bits is around 20 MB while one image of digital mammography usually generates 32 MB of data.

Despite the broad range of digital imaging applications in diagnosis, therapy, and prognosis, tools for archiving, communicating, classifying, organizing, and retrieving medical images to take full advantage of their rich information content are virtually nonexistent [78]. The lack of mechanisms to manage and distribute complex, interrelated, diagnostic images and text data hinders the effective use of primary imaging modalities in routine medical practice and in providing cost-effective patient care. During the last decades, individual investigators of medical imaging have been working on their own modalities in isolation, and, on the other hand, the medical informatics community focused primarily on textual information processing [63,64,62]. Recent advances in computers and networks, demands for cost-effective delivery of health care, and the presence of recognized and implemented standards, bring together the imaging and informatics researchers in the effort of integrating primary digital modalities and clinical databases into digital radiology. This effort also aims to facilitate the "mining" or extraction of the rich information embedded in medical images in a way that exceeds the individual sum of each of the imaging technologies.

The purpose of this paper is to describe our research towards the development of a framework for totally digital radiology departments, managing medical images, associated clinical text and voice data electronically. This framework of *multimodality image base management (MIBM)* is a logical progression from the primary digital modalities and from the stand-alone medical information systems towards the networked multimedia environment for digital radiology. Building such a comprehensive framework that encompasses generation, processing, management, and communication of multimodality, multidimensional images for diverse medical applications is an evolutionary process. System design must provide the capability to phase out old components and incorporate new devices and technology as the system evolves. Therefore, our development principle is first to implement an open systems infrastructure for managing and archiving multimedia files in the heterogeneous

distributed systems. Then, we retrieve these files from the consolidated data store and process them into more sophisticated data structures at a middleware layer for distributed data management. This middleware enables the user to access the multimedia data of the underlying heterogeneous distributed systems as if they are stored in a single database for diverse medical applications and services.

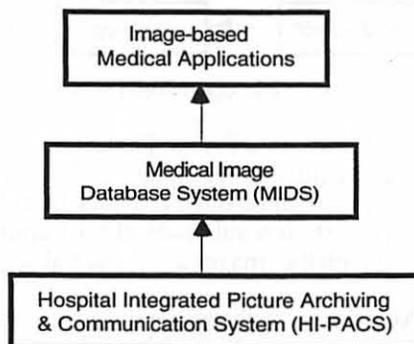


Figure 1. Three key components in the hierarchical framework of MIBM.

As shown in Figure 1, the MIBM framework consists of three hierarchical components: a *hospital-integrated picture archiving and communication system (HI-PACS)*, a *medical image database system (MIDS)*, and a set of *image-based medical applications* that relies on the information content and computational support of MIDS and PACS. HI-PACS provides an open systems infrastructure for acquiring, storing, and communicating raw image and text data. MIDS is the middleware that adds value to such consolidated data sets by processing and organizing them into more sophisticated data structures for targeted medical applications. The set of image-based medical applications can be localized in the medical center where MIDS resides or applied to remote sites using the wide area networking (WAN) technology; the latter enables better delivery of health care and sharing of expertise between major medical centers and smaller community hospitals. Further, the design and development of the MIBM framework is a bottom up endeavor, progressing from building the hospital integrated PACS infrastructure to providing multimodality image database management and services.

This paper is organized around the following major headings. Section 2 introduces the notion of PACS and discusses the implementation methodology at the University of California, San Francisco (UCSF) to extend the stand alone PACS into a hospital integrated environment. Section 3 describes the architecture of medical image database system (MIDS). Sections 4-6 illustrate three implemented examples of image-based medical applications at UCSF, namely, patient folder workflow, distributed object management, and multimodality neuroimaging analysis. Each of these applications assesses different capabilities of the MIDS server and serves specific medical functions that were previously not possible. Section 7 concludes this presentation and discusses the research direction.

## 2. PICTURE ARCHIVING AND COMMUNICATION SYSTEMS

Picture archiving and communication systems have been the prevalent means for acquisition, storage, communication, and display of medical images related to radiologic practice during the last decade [30]. A picture archiving and communication system (PACS) is a system integration of many components, including image acquisition devices, computers, communication networks, image display workstations, and database management systems. They attempt to utilize computing and communication power to increase the efficiency of the radiology department in a cost-effective fashion.

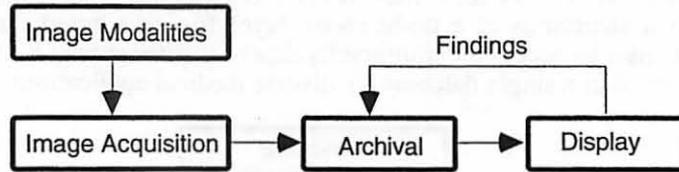


Figure 2. The operational flow of PACS.

Figure 2 illustrates the operations of a PACS. Digital images of an imaging scanner generated in an examination are captured by an acquisition computer and sent immediately to a central database for permanent archive. These images then will be retrieved later from the central archive and transmitted to display at a remote medical display station automatically or upon the user's request. The findings of viewing these images is appended with the images at the central archive.

The complexity of a PACS varies with applications. For example, a PACS for an intensive care unit simply comprises a video camera for digitization of radiographs, a baseband video system to transmit the images, and a video monitor to receive and display images. On the other hand, a hospital-integrated PACS is comprehensive and requires careful planning and large capital investment. During the past decade, several large-scale picture archiving and communication systems have been developed and are in clinical trial and use [28,30,45,24,34]. But fundamental system integration and operational issues have not yet completely been resolved, and the concept of digital (filmless) radiology yet to be materialized. These first generation system installations have been difficult to use in clinical practice. They are often applied to one or two modalities, closed to other medical information systems, and do not implement recognized standards of data exchange formats and communication protocols.

The research and development of the hospital integrated PACS at UCSF aims to remedy the shortcomings of first generation systems [53,2,31]. The fiber-optic backbone of HI-PACS has been installed at UCSF interconnecting its four major facilities, i.e., Moffitt and Long hospitals, ambulatory care clinic, and campus library, to provide a fully connected data bank that links clinical databases of different departments and contains patient data in various forms and states. The formation of this rich data bank thus enables the research into multimodality image database management. Figure 3 provides a schematic diagram of the installed HI-PACS, which emphasizes the following five new features: standardization, open systems connectivity, hierarchy memory management, database integration, and security. We discuss these features in Sections 2.1-2.5.

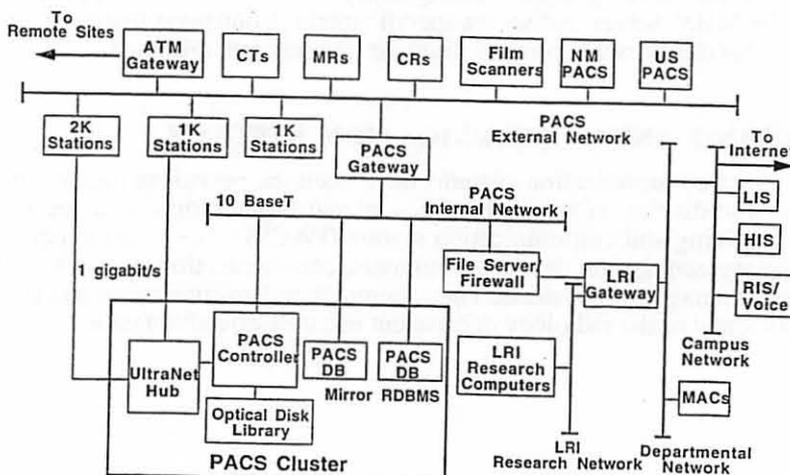


Figure 3. UCSF Hospital-Integrated PACS Architecture.

## 2.1 STANDARDIZATION

**Imagery Data Exchange:** Images generated from various radiological modalities give rise to different types of image file formats because of the manufacturer's own standards. Images archived in the MIBM framework must have a standard image file format so that images can communicate between system components freely. To this end, we have adopted the established ACR/NEMA 2.0 [1] and its updated version DICOM 3.0 (Digital Imaging and Communications in Medicine) [19] image file format as the standard for diagnostic images. The standard is to be applied right after the image file is acquired from an acquisition device, and the image can then be encoded in the standard file format. The standard file, however, can be decoded to a non-standard internal format of certain system component at a later stage. For example, different display stations often have different designs, and the standard file format might not be efficient for the workstation to display images rapidly. Thus, the workstation first receives the standard formatted images and converts them to the appropriate internal form for display. Its built-in interpreter will convert these images into the standard form again before transmitting to other external components.

**Textual Data Exchange.** Standard interface to other clinical information systems consists of two components: data format and communication protocol. For information other than images we will adhere to a standard format for file transfer. The Health Level 7 (HL7) [27] is an accepted application protocol for electronic data exchange in health care environments. In addition to the images, the clinician would require clinical text information from the hospital information system (HIS) [5] and radiology information system (RIS) [40], such as patient history, physical examination records, and radiology consultation reports, during the delivery of care. The information transferred from the HIS and RIS to PACS database should also follow the HL7 standard. The communication protocol between these databases will be standardized using the TCP/IP (transmission control protocol/Internet protocol) protocols [7]. The interfaces between RIS, HIS, and PACS controller in the HI-PACS implement the HL7 and TCP/IP [69,9] standards.

**Medical Vocabulary.** A standardized, uniform data set defines the central core of clinical records needed on a routine basis by the physicians and care providers. It establishes standard elements, definitions, and classifications of this core to minimize ambiguity and miscommunication. Current work attempts to use the vocabulary provided by the Systematized Nomenclature of Medicine (SNOMED), a subset of the Unified Medical Language System (UMLS) metathesaurus [32,33,21] from National Library of Medicine for describing anatomical and other annotations of medical images and embedded them into the fields defined in the DICOM 3.0 image header. If this application of standardized vocabulary is proven successful, we will extend the metathesaurus data set for radiologic report generation.

## 2.2. OPEN SYSTEMS CONNECTIVITY

**Local Area Networking (LAN).** The UCSF HI-PACS interconnects various acquisition subnetworks scattered at various hospital sites and acquires their images into the PACS optical jukebox [69], using the standardized data formats and protocols (see Section 2.1). These acquisition networks cover a variety of imaging scanners including CT, MRI, CR, film digitizers, and ultrasound scanners from the ultrasound PACS. The nuclear medicine PACS subnet consists of positron emission tomography (PET) and single positron emission computed tomography (SPECT) scanners. Nuclear medicine PACS is scheduled to be integrated into HI-PACS in the near future. Currently, the PET and SPECT image files are downloaded into the PACS database using the file transfer protocol (ftp) [7].

HI-PACS also supports distributed display station networks of UltraNet and Ethernet. The former network contains the 1 Gb/s UltraNet hubs for transmitting images to Unix-based image display stations each with two to four 1000 (1K) or 2000 (2K) lines of resolution monitors for concurrent image display in different modes and forms. The 2K display stations at the reading rooms of neuroradiology, chest, and pediatric, and informatics laboratory, as well as 1K stations at the pediatric intensive care unit and coronary care unit are already in clinics to support primary diagnosis. Besides

the central archive in the PACS cluster, each of these display workstations also maintains a 5 GB of RAID (reliable arrays of inexpensive disks) level 5 disk array [55] that provides the short term storage of most recent medical images of a given body region (5 - 10 days) relevant to its subspecialty location to reduce the workload of the networks and to provide fast on line transaction processing. For example, neuroradiology display station keeps only MRI and CT head images in its local RAID disks, and these images can be retrieved and displayed within 1 to 2 seconds. Moreover, over fifty Macintosh image review stations on a departmental Ethernet are also connected through a multimedia file server (see Section 2.5) that allows the clinicians to review low resolution images from their offices [60].

In addition, HI-PACS links the internal radiology PACS network with other external clinical and academic networks (see the lower right portion of Figure 3). These connections allow the PACS network to access the hospital information system (HIS), radiologic information system (RIS), and library information system (LIS), as well as to communicate with the outside world, including the library network and the Internet. Section 2.5 discusses the security measures of this open network.

**Wide Area Networking (WAN).** UCSF Medical Center is affiliated with four other hospitals and clinics: Mt. Zion Hospital, San Francisco VA Medical Center, San Francisco General Hospital, and San Francisco Magnetic Resonance Center, in the San Francisco Bay Area. For wide area communication, HI-PACS has installed a 155 Mb/s per sec (Mb/s) high speed ATM (Asynchronous Transfer Mode) wide area broadband network and a T1 frame relay carrier line (1.54 Mb/s) between Mt. Zion and UCSF Medical Center. The ATM connections of the other three hospitals with UCSF are scheduled in 1995.

In the last few years, T1 has appeared as the main carrier to provide high bandwidth between remote hospital locations [31,46,53]. For real time diagnostic imaging applications, however, ATM promises to be the ultimate on-premise internetworking technology [18,3]. ATM, proposed by international standards organizations [3,4], uses 53 byte cells to transmit data in multiples of OC-1 rates (51.84 Mb/s). It differs from the traditional local networking technology, such as Ethernet, where all computers share one communication medium. A switch-based ATM network is capable of supporting multiple connections simultaneously. The aggregate bandwidth of an ATM switch may be several Gb/s or more [18]. The ATM connection between UCSF and Mt. Zion Hospital is based on the OC-3c (155.52 Mb/s) ATM switches. We are also implementing the ATM technology to UCSF local PACS networks [41]. In this configuration, PACS computers and devices are connected to one or more ATM LAN switches, and the communication between each pair of computers is established through a shared LAN switch.

### 2.3 HIERARCHICAL MEMORY MANAGEMENT

A traditional PACS adopts a centralized approach in image data storage such that the system performance degrades rapidly during busy operations and any software or hardware problem occurs in the only one database may halt all of the radiological operations [15,47]. Figure 4 illustrates the new approach of hierarchical multimedia data archive in the HI-PACS, which implements a composite image staging mechanism of multiple storage media: magnetic disks, erasable magneto-optical disks, and write-once-read-many (WORM) disks. It is composed of the PACS controller (a multiprocessor archive server), a mirrored relational database management system (DBMS), a optical disk library (ODL) with a 1.3-terabyte (TB) optical jukebox which will be double its capacity to 2.6 TB in mid-1995, and a dual speed communication network [69].

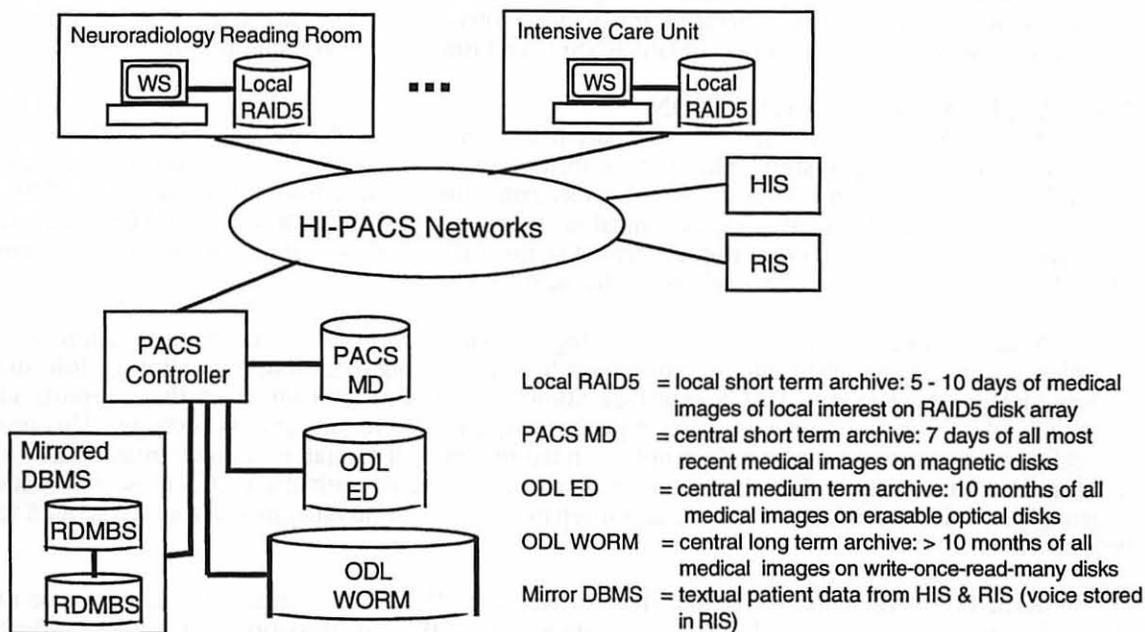


Figure 4. Implementation of hierarchical memory management in HI-PACS.

The PACS controller is configured with four central processing units (CPUs) to support multiprocessing and 14 gigabytes (GB) of magnetic disk memory to provide instant access to all short term images (about 7 days). The mirrored DBMS is configured with two SPARC computers with 64 MB RAM and four 2.1 GB high-speed differential SCSI (Small Computer System Interface) disks, running structured query language (SQL) utilities. All database transactions on managing the images stored in the optical jukebox and related textual information are mirrored on these two database servers. The mirroring feature provides non-interrupt operations that ensures that the minimal system vulnerability to data loss in case of disk failures. Moreover, the two database servers allow the users on the hospital-integrated PACS to access the patient images and records from heterogeneous client machines according to different levels of access privilege. The optical disk library (ODL) features four multi-function optical drives which support erasable magneto-optical disks for longer term storage and WORM disks for permanent storage. This optical jukebox is connected to the archive host via two small SCSI buses which support concurrent archive and retrieve operations. All patient data and images recorded by the PACS controller are tagged with the time and date of transaction, thus making the PACS data a continuous chronological medical history of patients.

The local display network consists of a standard Ethernet network and a high-speed fiber-optic UltraNet network. The Ethernet network is designated for receiving images from acquisition nodes and for transmitting images to low resolution or less time critical display stations. The one gigabit/sec (Gb/s) UltraNet network is designated for distributing images and relevant information to high-resolution 2K (2,048 × 2,048 × 2 bytes) display stations, as well as certain time critical 1K stations, such as intensive care units [28]. Each display station contains a RAID level 5 disk array capable to store up to 5 - 10 days of images of interest to the local operation. For example, the neuroradiology display station will archive only head CT and MRI images locally to provide 1 ~ 2 seconds response time of recent images. For another modality or an older image, the clinician has to query the PACS cluster, through the display station's graphical interface. The retrieval would take from a few seconds to a minute, depending the storage location and size of images requested. Further, all display stations have the same software functionalities such that the local clinician can also query the PACS cluster through the display station installed in a section of different subspecialty. This design provides the added flexibility to support mobile patient care in the hospital environment and avoids a single point of failure stopping the local

radiological section's operations. Thus, the implementation of data placement in HI-PACS maintains a balance between response times, network traffic, fault tolerance, and storage costs. The PACS memory management software is written in standard ANSI C and run under UNIX environment.

## 2.4 DATABASE INTEGRATION

The UCSF PACS interfaces with many heterogeneous databases, clinical or administrative, such as voice reporting system, radiology information system (RIS), hospital information system (HIS), library information systems (LIS), and electronic mail. Our integration strategy is to filter and consolidate the medical data of relevant clinical information systems into the central PACS cluster while providing on line access to other non-clinical systems, including such library information systems as Melvyl and KnowledgeFinder, as well as electronic mail.

**Voice Reporting.** Traditionally, radiological reports are first recorded on an audio cassette recorder and then transcribed into a textual form before inserting them into the radiology information system (RIS). The RIS and PACS interface allows for sending and inserting these reports into a centralized database such that a user can request to display these reports at a workstation. This process is inefficient, as the textual report does not reach the referring physician in a timely manner due to the delay in the transcription. The digital voice database associated with the PACS cluster enables the referring physician to look at the images and listen to the report simultaneously before the written report becomes available [8].

**RIS, HIS, and PACS.** The interface of RIS and HIS can be handled through the use of the HL7 (Health Care 7 [27]) standard for text data and TCP/IP (transmission control protocol/Internet protocol) protocol suite [7]). The interface between RIS and PACS controller assumes the following [9]: (1) the characteristics of both RIS and PACS, remain unchanged in its configuration, data, and functions performed, (2) each system has hardware and software extended for allowing communication with other systems, and (3) only data is shared, functions remain local. Based on these assumptions, successfully interfacing the PACS and the RIS includes the following tasks: (1) identify and uniformly represent for each system the subset of data that it agrees to share with the other systems; (2) make the shared data meaningful; and (3) define the communication protocol of data transfer and data format standard, i.e., TCP/IP and HL7. As the medical image base management framework evolves, HI-PACS will integrate more clinical databases into its networks.

## 2.5 FILE SERVER AND FIREWALL

UCSF HI-PACS is a dual-use healthcare information infrastructure that not only provides patient data and images to aid clinical diagnosis but also serves as a large data bank to facilitate biomedical research. The department of radiology at UCSF operates an Ethernet network that supports hundreds of Unix and Macintosh computers for electronic mail, software and file transfer, and library literature search, and is connected to other academic networks and Internet.

To maintain data integrity while allowing the clinicians and the researchers to access the PACS data archive, we developed a multimedia file server linking the radiology departmental network and the Internet to the UCSF PACS networks [59,60]. This system is based on a Unix SPARC 20 workstation that is physically located in the locked PACS cluster room at the Laboratory of Radiological Informatics of UCSF. This file server functions also as a firewall machine to isolate internal PACS networks from the outside world [14,50]. The firewall program of the file server allows only limited kinds of messages from the computers registered in the file server's database to pass through and rejects all unwarranted or unauthorized messages from the outside (whether they're mail, ftp, or attempts to break in) that are directed to internal PACS machines. The kinds of messages that are supported for registered computers include the SQL query to PACS textual databases, the retrieval of medical images, and the composition of a multimedia patient record of predefined formats. To further ensure patient confidentiality, the multimedia file server allows access only to those with a need to know and then certifies positively their identity before granting access.

### 3. MEDICAL IMAGE DATABASE SYSTEM

The intent of picture archiving and communication systems is to provide image file management and distribution in a digital radiology department. However, not only is the analysis and archiving of images for any single patient not well addressed, but there is little effort to gather data from images of different patients and coupled them with the information carried in the clinical text reports for the purpose of obtaining relevant disease knowledge. What is missing from PACS is the means to organize, synthesize, and present medical images and associated text to the users in a structured way for diverse image-based medical services and research. Furthermore, as PACS networks grow and user populations diversify, performance, transaction management, organization workflow, and imaging processing requirements surface and demand resolution. System developers must respond with an appropriate solution for distributed-application development, or the management, maintenance and evolution of the multimedia healthcare information infrastructure will become increasingly difficult.

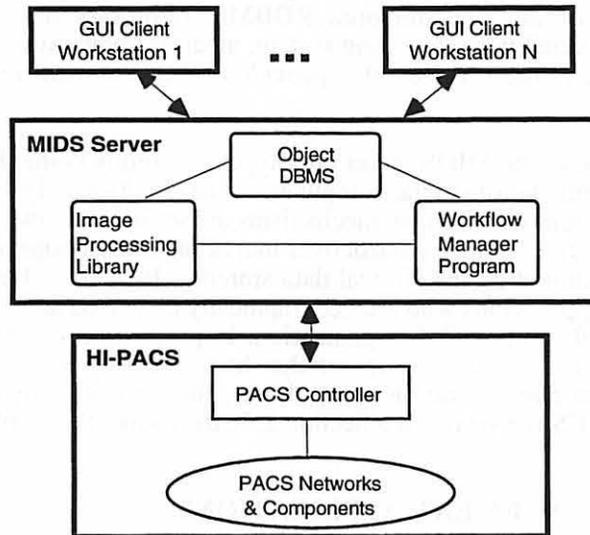


Figure 5. A system architecture of medical image base management with MIDS server over HI-PACS.

To address these issues, we are developing a middleware on top of HI-PACS to support distributed-application developments in health care organizations (see Figure 5). This middleware, or the MIDS server, is a centrally located core of resources and application routines. It provides much needed database management capabilities for PACS and serves to filter all client requests through various, centrally controlled levels of security, such as IP (Internet Protocol) firewalls, access control, and user authentication. The key components in Figure 5 are described as follows.

- (1) **Object DBMS** that contains a collection of object modules and interfaces based on an object oriented data model [12]. This DBMS treats other system components and the underlying multimedia data as a logical conglomeration of distributed, interacting objects of various levels of granularity. This model of distributed objects adheres to the concept of object request brokers (ORBs) proposed by the object management group (OMG) forum [51,52] for open systems computing.
- (2) **Image Processing Library (IPL)** that provides various image processing routines for different medical applications. This set of routines includes computational intensive algorithms, such as automatic image registration and feature extraction. The decision to locate heavy duty imaging application routines in the more powerful MIDS server is to reduce the hardware requirements and software complexity on the desktop client workstations.

- (3) **Workflow manager program** that schedules and controls the flow of related global events or activities reliably based on high-level workflow specifications and procedures. All three components in the MIDS server interact with each other during system operations.
- (4) **GUI Client workstations** each provides a graphical user interface (GUI) for the remote user to request services of the central MIDS server to perform operations, provide the arguments to the operations, retrieve results, and present them in a windowing environment. Each client workstation contains simple image processing capabilities, such as thresholding and window and leveling, and accesses the MIDS server for more complex and computational intensive imaging tasks. Some of these client workstations can be modified from the existing PACS display stations containing multiple high resolution display monitors. The GUI design of each workstation is tailored to the local medical application and requires the GUI builder tools such as GainMomentum (Gain Technology, Sybase, Palo Alto, CA), OpenLook, and Motif windowing toolkits.
- (5) **HI-PACS** (described in Section 2) consists of the PACS controller and a collection of distributed heterogeneous medical databases, imaging acquisition devices, and information systems, such as PACS optical disk library and mirrored RDBMS, radiology information system, hospital information system, digital voice reporting system, medical library system, image display stations, and electronic mail and file services, that provide the environment in which the above system components run.

For research purpose, the MIDS server prototype is currently being developed on a 4-processor SUN 690 MP SPARC computer of similar computational capabilities and environment as the computer that the PACS controller runs on. Security mechanisms are set up to allow the MIDS server accessing the PACS controller with fine-grained control over individual transactions and data items, yet without running the risk of contaminating the clinical data stored in PACS databases. Security is especially important to serve user applications who are geographically dispersed and highly diversified, possibly including clients external to the UCSF organization. In production environment, the MIDS server software will be ported to the same computer that hosts the PACS controller software for better performance and system administration. In addition, the security function of the MIDS server supersedes that of the PACS file server (see Section 2.5), the former thus replaces the latter.

#### 4. WORKFLOW MANAGEMENT APPLICATIONS

The system integration environment of the MIBM framework enables automated workflow management to streamline and coordinate information-based radiologic activities at levels of quality, throughput, and reliability which were never before possible. That is especially significant in this era of reducing health care costs and greater demands for shortened examination time and faster patient responsiveness. As an example, sometimes, the request of a physician may involve the tasks concerning not just the past medical history of patients, but also future activities. Consider the following case where John Smith's doctor invokes the following request:

*(4.1) Alert me once all the pathology reports of John Smith are ready, and show me these reports along with the radiology reports and the images. Then, save them for two weeks. Do not alert me again until the authorized final reports are ready.*

Any long data transaction of this kind that involves multiple operations spread over a prolonged period is outside the functional scope of DBMS. Instead, MIDS will have the workflow manager program to handle such a request. The workflow manager will spawn active rule objects, each contains one or more specific *Event-Condition-Action* (ECA) rules based on the method developed in [74]. Such a rule takes the form:

ON *event* IF *condition* THEN *action*

The *event* part of an ECA rule specifies events related to general operations, temporal events, or signals from other component systems. The *condition* part specifies a set of queries to the object DBMS

of the MIDS server. The action part specifies a function call or set of calls by the workflow manager program. When the event occurs in the distributed system, the rule is fired, and the condition is evaluated. If the condition is satisfied, the action is executed. Rules, events, conditions, and actions will be included as object types. For Query (4.1), the triggering event will be that the pathology reports are ready and the checking condition is whether the reports are authorized. The actions of this query call a set of functions to display images, pathology and radiology reports and to save them in the local MIDS database (not PACS database) for two weeks.

ECA rule objects differ from active triggers in databases [16], the latter do not address the full cooperation in the large distributed system. The workflow manager program of the MIDS server uses the declarative rule objects to support the design, analysis, and maintenance of global flow of information and to manage a variety of resources, such as related database transactions, display programs, and image processing and transfer, rather than memory and CPU cycles of conventional multidatabase transactions [25]. It also uses them to oversee the flows of global tasks while allowing individual component systems to focus on specific, modular tasks.

An example of enterprise-level operations in a digital radiology organization is the need to pre-fetch the images before the arrival of the patient, append to current examinations, and deliver to the display stations as soon as the current examination is completed. This requires the coordination of nine interrelated events of information processing and transfer among heterogeneous distributed systems (see Table 2) [31]. In the MIBM framework, the workflow manager program would monitor the occurrence of these events and coordinate the desired sequence of tasks through a group of ECA rule objects. Figure 6 shows such an intricate flow of events and actions between the workflow manager and related information systems in HI-PACS, i.e., databases, information systems, image scanners, voice dictaphone, and display workstations.

ECA RULES		
EVENT	CONDITION	ACTION
(1) Patient Admission	Previous images& reports of that patient exist	Select appropriate images and reports from PACS optical jukebox to assigned workstations at hospital wards and reading rooms.
(2) Order Entry	Event (1) occurred.	Make sure action (1) is completed.
(3) Patient Arrival	Events (1) and (2) occur	Make sure action (2) is completed.
(4) Examination completion	New images scanned	Store these images in temporary optical archive in PACS. Send them to the same assigned workstations in wards and reading rooms.
(5) Voice Dictation	Dictation of the physician recorded	Add digital voice report to temporal archive. Send them to the same workstations.
(6) Written Transcript	Preliminary written report done	Add preliminary report to temporary archive. Send it to RIS, and workstations. Erase voice report from temporal archive.
(7) Signature	Final report done	Add final report to temporary archive. Send it to RIS and assigned workstations.
(8) Transfer	Action (7) completed, patient transfer	Transfer all relevant image files to workstation at the specified new location
(9) Patient Discharge	Images and reports are ready	Copy patient information from temporary archive to permanent archive in PACS. Erase patient information from workstations.

Table 2. A table of major event-condition-action rules from patient admission to patient discharge.

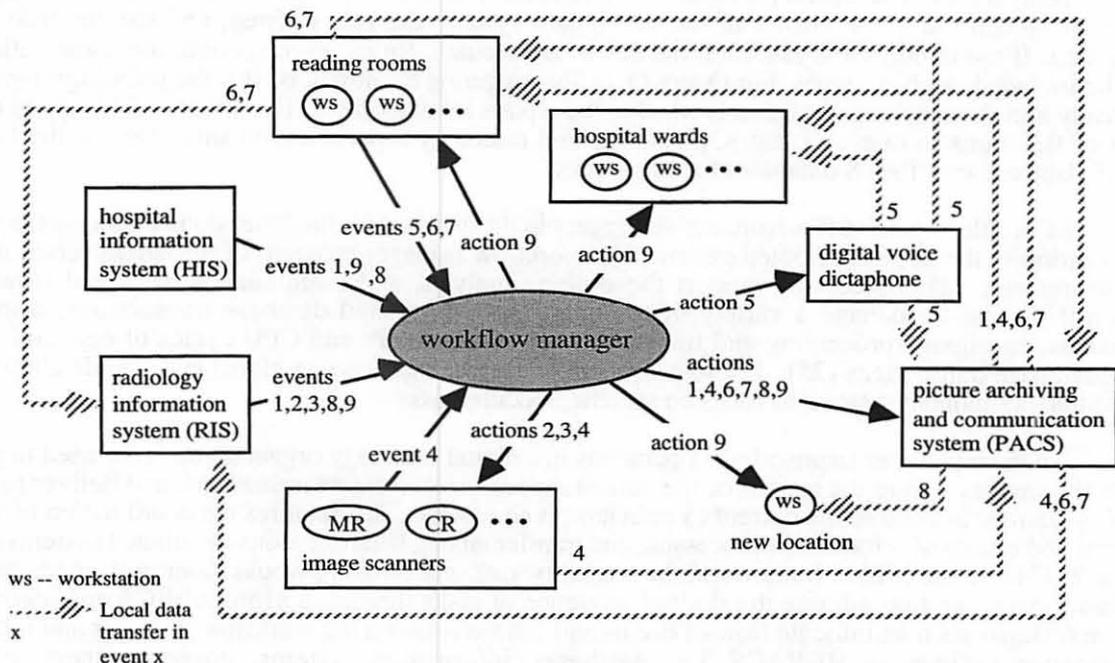


Figure 6 The corresponding flow of events and actions between the workflow manager of the MIDS server and the underlying system components of HI-PACS specified by the event-condition-action rule table in Table 2.

Also shown in Figure 6, heterogeneous component systems of HI-PACS remain autonomous in that they would manage localized transactions and data transfer during these events. In Event (1), for instance, the PACS controller will also send the patient image files acquired from previous visits to the assigned workstations in the hospital wards and reading rooms for the study by the physicians and the radiologists, respectively. The software prototype for patient folder management has been completed and tested. Currently, we are coordinating with the clinical sections of the UCSF hospitals to implement the production system.

## 5. DISTRIBUTED OBJECT MANAGEMENT

The Object DBMS at the MIDS server contains a commercial DBMS and a library of software programming objects that extends its capability for biomedical object modeling and interfacing external databases of HI-PACS. These objects are coded in an object data model as in object oriented systems [12,71]. We adopt the object-oriented paradigm due to its suitability for multimedia data representation [10,35,49] and open systems computing [51,73,43,54,42]. The object model is simply summarized: (1) basic modeling primitive of medical images and data is the object; (2) objects can be categorized into types, and all objects of a given type exhibit common behavior and a common range of states; (3) the behavior of objects is defined by a set of operations that can be executed on an object of the type; and (4) the state of objects is defined by the values they carry for a set of properties, which may be either attributes of the object itself or relationships between the object and one or more other objects [11].

In this way, the MIDS server can treat the underlying multimedia medical records and system resources as a logical conglomeration of distributed, interacting objects of various levels of granularity. These include images and free text, production rules, system services such as image processing, data archive, and workflow management. As an example, the workflow management discussed in Section 4 relies heavily on ODBMS for storing routing instructions, procedures, and rules. The ODBMS in turn can be used to track the status of workflow processes and maintain a historical audit trail of each transaction. Further, the ability of ODBMS to define independent, generalized operations allows

procedures of arbitrary complexity to be included. This is important for open multimedia computing environments, such as that of medical image base management, where some of the "objects" in the heterogeneous distributed systems will be software programs or other non-database components. The MIDS server treats both database and non-database components uniformly as objects and invokes them through defined object implementations.

Distributed object management can be used to facilitate intelligent generation of medical records, especially those involving a combination of patient images and data retrieved from the underlying distributed databases. Figure 7 shows a simple example of type-composition structure of multimedia medical record. The type of generic radiologic document containing subtypes such as medical images, diagnostic reports, and study details and history. These object subtypes are further decomposed into simpler object data types, such as thumbnail image slide and radiologist name, as well as more complex document components, such as image, report and findings.

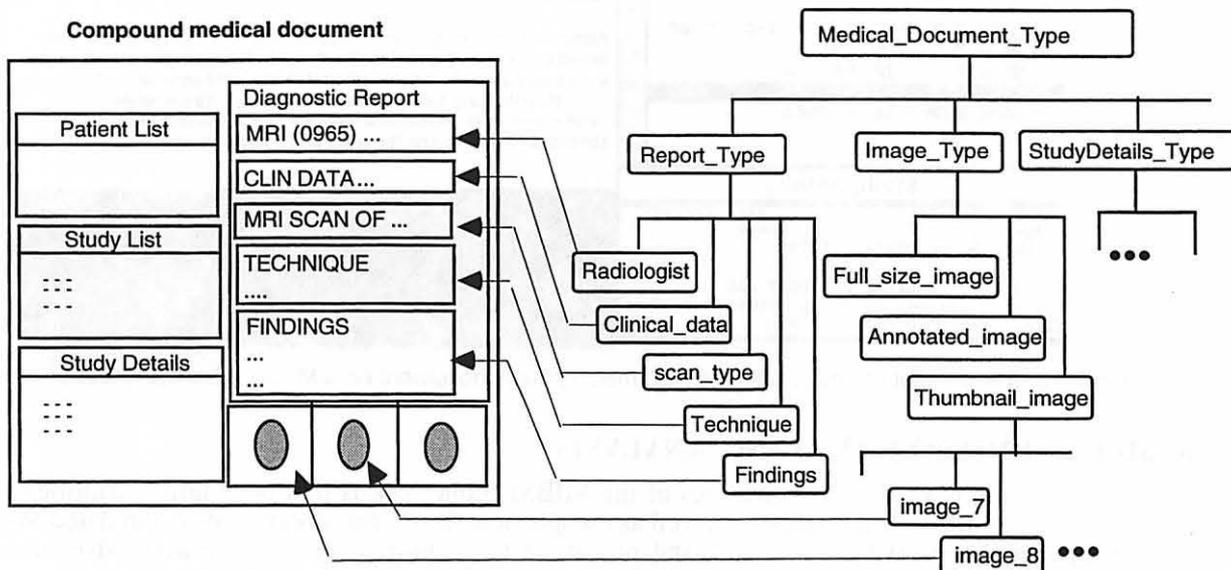


Figure 7. A multimedia medical record and its associated object type structure.

In Figure 8 of the following page, we show an object instance of this multimedia document type containing thumbnail images and with menu options available to select other types of images, i.e., full size and annotated images, for more detailed display. The patient name can be changed in this example to maintain the confidentiality. This mechanism of intelligent documents currently supports the Macintosh [59] and Unix computer platforms [72]. Such an object-oriented multimedia record structure [70] enables flexible retrieval, display, view, and presentation of clinical data to meet varied reporting needs of practitioners. A type-subtype hierarchy of document classes allows the dynamic composition of patient records to satisfy individual preferences and work habits, and increases user acceptance of electronic patient record [7]. In addition, the distributed object approach can be used to explore different formats of information display on screen, including the chronological presentation of patient information [23] or problem-oriented composition of medical records [44].

Utilities Information Tools Reference Tools Educational Tools

Enter Pt's. Name or ID#:

64x64 Retrieve Requested Full-Size Images R

**Patient List**

DOE, JANE	9999999	14 y/o F	↑
SEARCH 1			

**Diagnostic Report**

MRI BRAIN W/WO (0965) -Dictating Radiologist  
CLINDATA: ?rabies

MRI SCAN OF BRAIN: 3/26/94. No comparison films

TECHNIQUE:  
Sagittal T1, axial first and second echo T2, coronal fast spin echo T2, axial T1 pre- and post-gadolinium and coronal T1 post-gadolinium.

FINDINGS:  
Abnormal T2 prolongation with subtle mass effect is seen within the medial aspect of the right temporal lobe involving the hippocampus and the amygdala. More subtle areas of T2 prolongation are possibly seen within the left temporal lobe. No enhancing abnormalities are identified. The ventricular size is within normal limits. Normal flow voids are seen at the skull base. No abnormality is noted within the orbits or the paranasal sinuses.

**Study List**

03/26/94	05:56:08	MR	SAG-T1	↑
03/26/94	05:56:08	MR	AXIAL T2	
03/26/94	05:56:08	MR	AXIAL T2	
03/26/94	05:56:08	MR	T1 AXIAL PRE/POST GAD	
03/26/94	05:56:08	MR	FSE2 COR	
03/26/94	05:56:08	MR	FSE2 COR	
03/26/94	05:56:08	MR	CORONALS	↓

**Study Details**

Name: DOE, JANE ID#: 9999999  
Age: 14 y/o F BDate:  
Hx: CHANGE MENTAL STATUS

Study: MR T1 AXIAL PRE/POST GAD  
Date: 03/26/94 05:56:08 Std#: 999-99  
Slice: 24 ( 5.0 mm) Dr. Req. Physician

7/24 8/24 9/24

Figure 8. A snap shot of object-oriented multimedia patient document on a Macintosh display station.

## 6. MULTIMODALITY IMAGING ANALYSIS

One of the most attractive features of the MIBM framework is that once large quantities of patient images of different modalities, as well as complemented clinical reports, are accumulated over time, they can be retrieved, combined, and processed to support diverse image-assisted medical applications. The image database approach has been explored among the PACS researchers. A representative work of this approach retrieved the computed radiography (CR) hand images from a stand alone PACS and performed the phalangeal and carpal bone analysis to assess skeletal age and maturity. The analysis involved a set of low-level image processing tasks, including thresholding, spatial filtering, and morphological filtering, to extract bone objects and their attributes based on a predefined object data model of human hand bone structure [56,57,10]. Our experience gained from that research, however, is that without the integration of clinical patient data and the interplay between a variety of digital modalities, the applicability of image database systems for routine clinical service is rather restricted. This motivates our implementation of a hospital-integration PACS before embarking the design of medical image database systems.

Figure 9 illustrates a general flow of GUI client workstation in the MIBM framework for neuroimaging analysis. The function of this workstation is to provide an automated tool for the physicians to access, organize, and process the multifaceted, multidimensional neuroimaging data and present critical diagnostic information to them effectively. In this setting, the radiologist or the imaging specialist queries the MIDS server for the patient images and records according to certain attributes, such as patient name, medical history, disease types, image modality, and size of certain brain regions or objects. He or she can create and visualize electronic multimedia reports consisting of single or multiple images (static, cine, video, or 3D), and accompanied by voice and/or text using the graphical user interface provided by the 3D display workstation. Often, the tomographic images from a single patient subject taken at different times or of different modalities are needed to be reoriented into a

common coordinate space or registered for fair comparison. The registration algorithm also applies to the scenario of studying inter-subject images.

Another important feature of the 3D neuroimaging workstation is the ability to measure, quantitatively, regions or volumes of brain that are of important clinical relevance. There is no convenience means to do so in the current film-based radiologic practice. Finally, when the analysis of a patient case is completed, the medical workstation would then prompt the user whether to archive the extracted data, e.g., volumes of certain brain segments as well as metabolic counts of those volumes in PET scans, into MIDS object database automatically. In this way, the content of the image database will be enriched with the accumulation of a large number of diagnosed cases. These cases in turn can be used for supporting other data-intensive clinical studies, such as measuring the effectiveness of neuroimaging procedures or analyzing the patient outcomes.

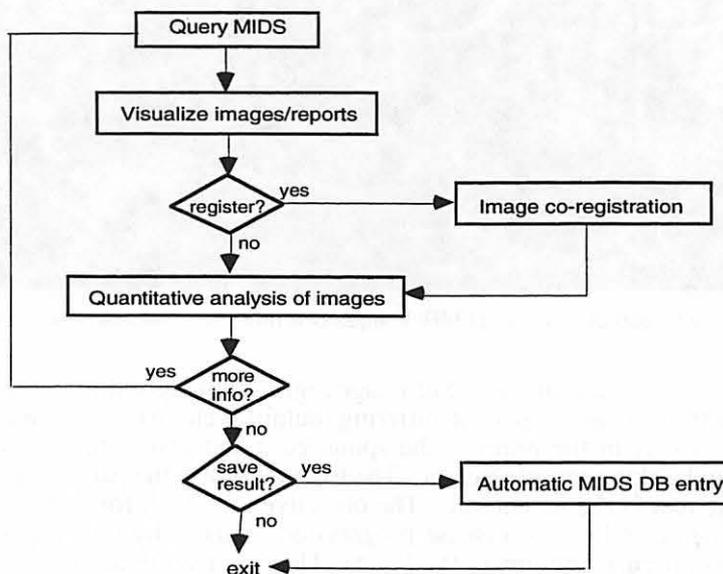


Figure 9. Operation flow of a 3D neuroimaging workstation in the MIBM framework.

To support the image retrieval by content, the MIDS server requires to pre-process patient images and reports and extract relevant information before sending them to the PACS cluster for archive. The MIDS server keeps only the indexing information of extracted image objects and their attributes in a brain data model of its ODBMS. These indices enable the MIDS server to retrieve any specific or any combination of image segments from the large store of image data files archived in the PACS ODL. Since most of the relevant clinical information have already been contained in the header files of medical images (following the DICOM or the ACR/NEMA standard, see Section 2.1), as well as the text of the patient reports stored in RIS and HIS, simple software routines can be written to convert or extract these text data into a pre-defined object-oriented model in ODBMS. For more complex brain images, however, an interactive method is used to identify regions or volumes of interest. Automated segmentation programs are being developed to improve the extraction process, these programs are similar to the one we developed for x-ray bone analysis [56,57] but are more attune to the neuroimaging modalities, such as CT, MRI, and PET.

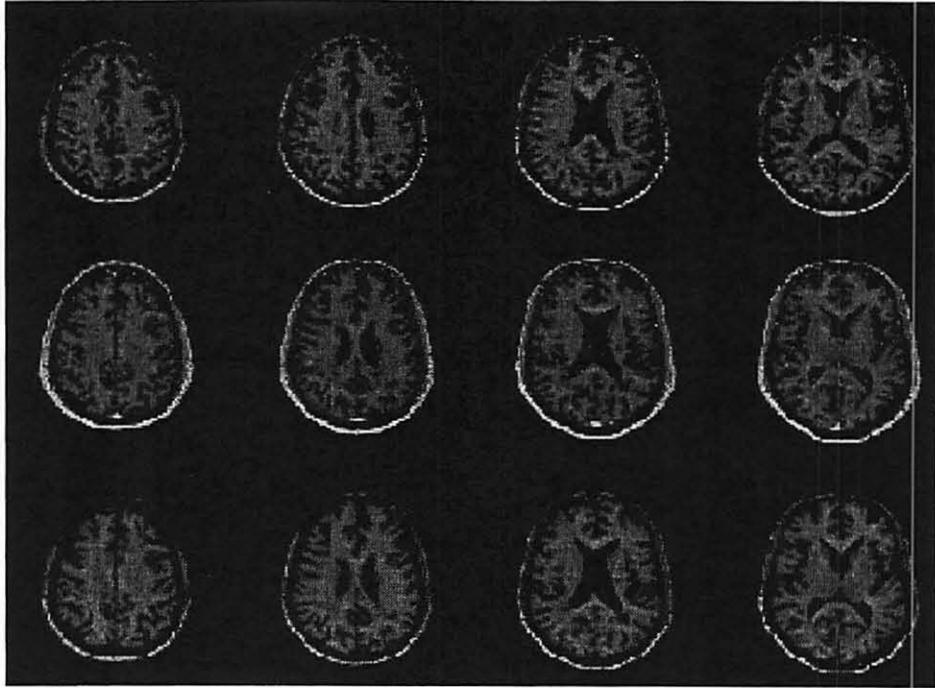


Figure 10. Image coregistration of two sets of MRI images of a multiple sclerosis patient taken a year apart [72].

Figure 10 provides a case of the use of image registration algorithm to register two sets of MRI images taken from a thirty-year old patient suffering multiple sclerosis, a disease condition marked by patches of hardened tissue in the brain or the spinal cord and associated especially with partial or complete paralysis and jerking muscle tremor. The top row shows the patient's brain scan taken a year ago while the middle row is the recent one. The objective is to look for differences between the two scans to evaluate changes related to disease progression (marked by the sizes of low intensity dark regions around the ventricular regions of the brain). Using conventional films, the radiologist has to rely on imagination to correct for discrepancies in head position that result from different positioning in two scans; even then, he or she cannot compare changes of lesions in top scan to middle scan.

Within the MIBM framework, the client workstation can request the MIDS server to align the first scan based on its 3D dataset with the recent one, and the result is displayed in the bottom row. The result indicates that the disease condition of the patient is not worsen, as the size of low intensity regions, i.e., lesions, around the ventricular areas remain the same. The set of image registration algorithms are modified from [75] and can be extended to intrasubject and intermodality registration [72,76]. The MIDS server makes possible the on-line access of a large image store for various registration purposes, including looking for subtle changes in signal across time, combining different types of functional (such as PET and SPECT images) and structural information (such as CT and MRI images), and comparing different subjects or patients to one another for better clinical decision making.

Currently, this 3D neuroimaging workstation is being used in the clinical research of presurgical epilepsy evaluation [72]. We are also implementing additional image processing software to enable large scale experimentation on the neurological disorders of preterm infants caused by abnormal brain myelination (white matters) growth.

## 7. CONCLUSIONS

The next generation of health care information systems will consist of a vast network of heterogeneous, autonomous, and distributed imaging scanners, image acquisition systems, databases,

information systems, knowledge intensive applications, and large quantities of multimedia medical data. The diffusion of multimedia information management will enable the development of tools available for effective decision making and open up many opportunities in advancing healthcare services. To transmit and link patient information, however, presumes the existence of an infrastructure. Currently, most medical information systems, such as hospital information systems and radiology information systems, handle only textual patient information and lack the computing mechanisms and networking resources to handle different kinds of digital modalities. A key challenge facing system researchers and builders is to provide a new organizational framework to integrate the current distributed information resources into what appears to the users as a uniform and logical conglomeration of databases supporting diverse image-assisted medical applications.

In this paper, we have presented our research towards the development of a framework for multimodality image base management (MIBM). This framework can serve as a blueprint for developing digital radiology departments for the hospitals of the future. The implementation of the MIBM framework is a bottom up endeavor, with the foundation of a comprehensive infrastructure for multimedia file management, archival, and communication before developing sophisticated image database management and applications. Applications development in MIBM is evolutionary. Its open systems connectivity and object interfaces support the incremental integration with standardized information systems as well as legacy systems of other hospital departments.

This hierarchical information framework consists of three components: a hospital-integrated picture archiving and communication system (HI-PACS), a medical image database system (MIDS), and a set of image-assisted medical applications. HI-PACS circumvents the shortcomings of the first generation PACS by implementing recognized standards, open systems connectivity, hierarchical memory management, database integration, and security. The value-added middleware, MIDS, enables the user to access the multimedia images and text of the underlying heterogeneous distributed systems as if they are stored in a single database. In addition, the MIDS server also provides image processing and workflow management for various medical applications. We have also presented three implemented examples in electronic patient folder, intelligent medical documents, and multimodality neuroimaging analysis. Our future work will enrich the object interfaces and object schema of the medical image database server, as well as develop object data models for different medical applications.

MIBM is the solution to the progressing need to computerize the radiology department and hospital and to support the increasing demands for image-based medical services in diagnosis, therapy, and surgery. MIBM research is just in the beginning. It will continue to grow with each new technology and device. The quest will continue to improve system performance including communication networks, storage I/O, volumetric image processing, and distributed data management. Challenging research areas are summarized in the following.

First, the advance of communication technology must be made to transmit large volumes of 3D medical image data in real time and to provide efficient interfaces between the proliferated local area and wide area networks. Our experience in building a large scale PACS in another healthcare campus of the University of California system (the use of Ethernet at 10 Mb/s, FDDI (fiber distributed data interface) at 100 Mb/s, and Ultraset at 1 Gb/s) concluded that the efficient utilization of bandwidth, upward scalability, reliability, and coexistence of various LAN and WAN technologies are crucial requirements for multimedia medical communications [31,28]. Ubiquitous deployment of high performance networks will allow the PACS developers to better resolve network infrastructure issues and focus more on the overlying, value-added medical services and applications. The emerging broadband integrated services digital network (B-ISDN) standard, Asynchronous transfer mode (ATM), promises to be the ultimate internetworking backbone. Its scalable, high bandwidth and uniform switching can transfer images, text, voice, and video from application to application at much higher speeds than now available [18]. However, much work still is needed to optimize the communication protocols and hardware interfaces to take advantage of this technology [41,3,68,77].

Second, the I/O processing speed of storage systems (either magnetic or optical medium) must be enhanced to remove this performance bottleneck. Current emphasis in disk storage system design is on capacity rather than I/O performance. While disk storage densities are improving impressively (60 - 80 percent compounded annually), I/O performance improvements have been occurring at only about 7 to 10 percent compounded annually in the last decade [61,65]. As a result, disk system performance is fast becoming a dominant factor in networked multimedia healthcare systems, which requires us to focus on designing I/O systems that can handle real-time applications of different I/O demands. Retrieval and display of images in PACS require guaranteed real-time I/O throughput, but telemedicine conferencing demands very small latencies of delivery. The research into high performance, large capacity, and highly reliable RAID systems may result in a storage server architecture for the heterogeneous computing environment of the future digital hospitals [13,38,55].

Third, image processing techniques must continue to reduce the computational overheads of volumetric image visualization [39,58] and multimodality 3D image registration [22,75,76] as well as to improve the accuracy of identifying the contents of images of various body parts [67,66]. The image understanding mechanisms remain very complex and rely not only upon the matching of image data but also upon the medical knowledge about the organs and the relationships between anatomy and function. Extensive work, especially 3D graphical visualization, has been carried out in the past to prove access and usage of medical images. A better use of medical images, however, requires more research into image data fusion [6,22]. This concerns the fusion of multimodality information, the fusion of information coming from different patients or from a priori knowledge (e.g., inter-subject image registration and diagnostic expert systems) and the identification and extraction of anatomical structures.

Fourth, database management systems must be designed and developed to manipulate and represent distributed, multimodal images effectively [78,36,2], for both local clinical use or collaborative medical work. Current standardization efforts in distributed object management [51,52] and object oriented DBMS [12], complemented by the recognized and implemented standards in healthcare industry [19,1,27], may lead to a practical solution to support demanding applications of mass image archiving, modeling, manipulation, and retrieval. Furthermore, as electronic workflow continues to infiltrate health care provider organizations, we will expect the image database system technology to undergo an even greater metamorphosis, by fusing task-specific enterprise workflow with data-centric transaction processing [17].

#### **ACKNOWLEDGEMENTS:**

This work is supported by the following funding sources: National Cancer Institute NCI PO1 CA51198 and NCI RO1 CA 40456, National Library of Medicine NO1-LM-4-3508, CalREN, Pacific Bell Corp. ATMN-007, and IBM Almaden Research Center Grant Title: Image Folder Manager Concept. The authors wish to thank the technical members of the Laboratory of Radiological Informatics and the clinical faculty of the Departments of Radiology and Neurology, UCSF involved in this work.

#### **REFERENCES**

- [1] ACR/NEMA Digital Imaging and Communication Standard Committee. Digital Imaging and Communications ACR-NEMA 300-1988. Washington, DC: National Electrical Manufacturers Association, 1989
- [2] L. Allen and O. Frieder, Exploiting Database Technology in Medical Arenas: A critical assessment of PACS, *IEEE Engineering in Medicine and Biology*, March 1992;11(1):42-49.
- [3] *ATM User-Network Interface Specification, Ver 3.0*, September, ATM Forum, 1993.
- [4] *The ATM Forum BISDN Inter Carrier Interface (B-ICI) Specification, Ver 1.0*, ATM Forum, August, 1993
- [5] AR Bakker, et al. Towards New Hospital Information Systems, Amsterdam, 1988.

- [6] C Barillot, D Lemoine, L Le Briquer et al. Data fusion in medical imaging: merging multimodal and multipatient images, identification of structures and 3D display aspects. *European J Radiology*. 1993;17:22-27.
- [7] U Black. *Computer networks*. 2nd ed., Prentice Hall, 1993.
- [8] CM Breant, RK Taira, HK Huang. Integration of a Voice Processor Machine in a PACS. *Comp. Med. Image Graphics* 1993;17:13-19
- [9] CM Breant, RK Taira, HK Huang. Interfacing Aspects Between PACS, RIS and HIS. *J. Digital Imaging*. 1993; 6:88-94.
- [10] AF Cardenas, IT Jeong, et al. The knowledge-based object-oriented PICQUERY+ language, *IEEE Transactions in Knowledge and Data Engineering*, Aug. 1993.
- [11] RGG Cattell. *Object data management: object-oriented and extended relational database systems*. Addison-Wesley, 1994.
- [12] RGG Cattell (ed.) *The Object Database Standard: ODMG-93*, Morgan Kaufmann, 1993.
- [13] PM Chen, EK Lee, AL Drapeau et al. Performance and design evaluation of the RAID-II storage server. *Distributed and Parallel Databases*. July 1994;2(3):243-260.
- [14] B Cheswick. The design of a secure Internet gateway. *USENIX Con. Proceedings*, Anaheim, Summer, 1992:233-37.
- [15] PS Cho, et. al. Centralized vs. distributed PACS for intensive care units. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
- [16] U Dayal, M Hsu and R Ladin, Organizing long-running activities with triggers and transactions, *Proc. ACM SIGMOD Int. Conf. on Management of Data*, 1990.
- [17] U Dayal, H Garcia-Molina, M Hsu, B Kao, and MC Shan, Third generation TP monitors: A database challenge, *Proc. 1993 ACM SIGMOD Int. Conf. on Management of Data*, 1993:393-397.
- [18] M De Prycker. *Asynchronous transfer mode: Solution for broadband ISDN*. 2nd Edition, Ellis Horwood, 1993.
- [19] *Digital Imaging and Communications in Medicine (DICOM) Version 3.0*, American College of Radiologists/National Electrical Manufacturers Association, 1993.
- [20] J Engel (ed.) *Surgical Treatment of the Epilepsies*, Raven Press, 1993.
- [21] DA Evans, WR Hersh, IA Monarch, RG Lefferts, SK Handerson. Automatic indexing of abstracts via natural-language processing using a simple thesaurus. *Medical Decision Making*. 1991;11(suppl):S108-115.
- [22] A Evans. 3D multimodality human brain mapping: past, present and future. In *Quantification of Brain Function*. K Uemura et al. (eds.), Elsevier Science, 1993.
- [23] AR Feinstein. Quality of data in the medical record. *Computer and Biomedical Research*. 1970;3:426-435.
- [24] F Georinger. Medical diagnostic imaging support system for military medicine. In *Picture Archiving and Communication Systems (PACS)*. Huang HK, Ratib O, Bakker AR, Witte G, eds. NATO ASI Series F, PACS in Medicine, Vol. 74. New York: Springer-Verlag, 1991:213-230.
- [25] J Gray and A Reuter. *Transaction processing: Concepts and techniques*, Morgan Kaufmann, 1993.
- [26] JO Greenberg. *Neuroimaging: A companion to Adams and Victor's Principles of Neurology*. McGraw-Hill, 1995.
- [27] Health Level Seven (HL7). *An Application Protocol For Electronic Data Exchange in Health Care Environments Version 2.1*. Health Level Seven, Inc. Ann Arbor, Michigan, 1991.
- [28] HK Huang, WK Wong, SL Lou, BK Stewart. Architecture of a comprehensive radiologic imaging network. *IEEE J. S.A. Communication*. 1992;10(7):1118-1196.
- [29] HK Huang. *Elements of digital radiology: A Professional Handbook and Guide*. Prentice-Hall, Inc., New Jersey, April, 1987.
- [30] HK Huang, O Ratib, et al. (eds.) Picture archiving and communication system (PACS). *J. Digital Imaging*. 1992;5(1):22-25.
- [31] HK Huang, RL Arenson, SL Lou, AWK Wong, KP Andriole, TM Bazzill, and D Avrin, Multimedia in the radiology environment: current concept, *J. Computerized Medical Imaging and Graphics*, 1994;18(1):1-10.

- [32] BL Humphreys and DAB Lindberg, Building the unified medical language system, in L Kingsland ed., *Proceedings of 13th Symposium on Computer Applications in Medical Care (SCAMC)*, Washington DC, IEEE Press, 1989:475-80.
- [33] BL Humphreys, DAB Lindberg, WT Hole. Assessing and enhancing the value of UMLS knowledge sources, *Proc. of 15th Symp. Computer Applications in Medical Care (SCAMC)*, 1991:78-82.
- [34] G. Irie. Clinical experience - 16 months of HU-PACS. In *Picture Archiving and Communication Systems (PACS)*. Huang HK, Ratib O, Bakker AR, Witte G, eds. NATO ASI Series F, PACS in Medicine, New York: Springer-Verlag, 1991;74:213-230.
- [35] H. Ishikawa, F. Suzuki, F. Kozakura, et al. The model, language, and implementation of an object-oriented multimedia knowledge base management system, *ACM Trans. on Database Systems*, March 1993;18(1):1-50.
- [36] R. Jain, (ed.) NSF Workshop Report on visual information management systems, Feb., 1992.
- [37] AE James, JH Anderson, and CB Higgins, eds., *Digital image processing in radiology*, Williams & Wilkins, Baltimore, MA, 1985.
- [38] RH Katz, PM Chen, AL Drapeau, et al. RAID-II: Design and implementation of a large scale disk array controller. *1993 Symp. Integrated Systems*, University of California at Berkeley UCB/CSD 92/705, 1993.
- [39] A Kaufman, D Cohen, and R Yagel. Volume graphics. *IEEE Computer*. July 1993; 26(7):51-64.
- [40] GM Kolodny. A Cost-Effective Radiology Reporting System. *Applied Radiology* 61-64, 1981.
- [41] HU Lemke, STC Wong, and HK Huang. Systems design criteria for ATM oriented broadband medical services. *Proc. European Symp. Advanced Networks and Services-95*, March, Amsterdam, 1995 .
- [42] W. Litwin, L. Mark, and N. Roussopolous, Interoperability of multiple autonomous databases, *ACM Computing Surveys*, 1990;22(3): 267-293.
- [43] F. Manola, S. Heiler, et al., Distributed object management, *Int. Journal of Intelligent and Cooperative Information Systems*, 1992;1(1): 5-42.
- [44] CZ Margolis. Problem-oriented record: A critical review. *Paediatrician* 1979;8:152-162.
- [45] H. Masser, A. Mandl, M. Urban, H. Hradil, W. Hruby. The Vienna Project SMZO in Picture Archiving and Communication systems (PACS). In Huang HK, Ratib O, Bakker AR, Witte G, eds. *NATO ASI Series F, PACS in Medicine*, New York: Springer-Verlag, 1991;74:247-250.
- [46] T.P. McGarty. Multimedia communications technology in diagnostic imaging. *Investigative Radiology* 1991;26:377-381.
- [47] C Morioka, et al. Interfacing a computed radiography system in a centralized PACS system through a microcomputer. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
- [48] JD Newell, Jr. and CA Kelsey, eds., *Digital imaging in diagnostic radiology*, Churchill Livingstone, New York, 1990
- [49] E. Nishihara, An overview of medical image data base, *Journal of Information Processing Society of Japan*, May, 1992, pp. 478-485.
- [50] NRC (National Research Council). *Computer at risk: Safe computing in the information age*. Washington, D.C., National Academy Press, 1991.
- [51] Object Management Group, *Object management architectural guide*, OMG Document 92.11.1, 1992.
- [52] Object Management Group, *The common object request broker: Architecture and specification*, OMG Document 91.9.1, Sept. 3, 1991.
- [53] M Osteaux. *A second generation PACS concept* (Ed.). Springer-Verlag, 1992.
- [54] MP Papazoglou, SC Laufmann, and TK Sellis, An organizational framework for cooperative Intelligent information systems, *Int. Journal of Intelligent and Cooperative Information Systems*, 1(1):169-202, 1991.
- [55] D Patterson, G Gilson, and RH Katz. Reliable arrays of inexpensive disks (RAID), *ACM SIGMOD Conference 1988*, Chicago, June 1988:1089-116.
- [56] E Pietka, F McNitt-Gray, ML Kuo, HK Huang, Computer-assisted phalangeal analysis in skeletal age assessment, *IEEE Trans. on Medical Imaging*, Vol. 10, No. 4, 1991:616-620.

- [57] E Pietka, L Kaabi, ML Kuo, HK Huang, Feature extraction in carpal-bone analysis, *IEEE Trans. on Medical Imaging*, Vol. 12, No.1, 1993:44-49.
- [58] SM Pizer. *Tutorial: Multiscale geometric image analysis, 3D biomedical imaging*. First IEEE-EMBS International Summer School, July 2-10, France. 1994.
- [59] MR Ramaswamy, DE Avrin, KP Andriole, AWK Wong, TM Bazzill, RL Arenson, and HK Huang. Use of personal computer technology in supporting a radiological review workstation, *SPIE Proceedings on medical imaging*, vol. 2165, 1994.
- [60] MR Ramaswamy, WK Wong, JK Lee, HK Huang. Accessing a PACS's text and image information through personal computers. *Am J Roentgenology*, 1995 (to appear).
- [61] C Ruemmler and J Wilkes. An introduction to disk drive modeling. *IEEE Computer*, March, 1994:17-28.
- [62] *Symposium on Computer Applications in Medical Care (SCAMC)*, 18th Annual Conf Proceedings. Washington DC, Nov 5-9, 1994.
- [63] E Shortliffe, L Perrault, G Wiederhold, and L Fagan. *Medical informatics*, Wokingham, Addison Wesley, 1990.
- [64] EH Shortliffe, LE Perreault, LM Fagan, and G Wiederhold (eds.) *Medical informatics computer applications in health care*. Reading, Mass., Addison-Wesley, 1990.
- [65] PA Steenkiste. A systematic approach to host interface design for high-speed networks, *IEEE Computer*, March, 1994, 47-57.
- [66] MR Stytz, G Frieder, and O. Frieder, Three-dimensional medical imaging: Algorithms and computer systems, *ACM Computing Surveys*, 1991;23:421-499.
- [67] P Suetens, P Fua, and AJ Hanson. Computational strategies of object recognition. *ACM Computing Surveys* 1992;24:5-61.
- [68] SM Walters, DS Burpee, and GH Dobrowski. Evolution of fiber access systems to ATM broadband networking, *Proc. IEEE*, Nov., 1993;81(11):1588-1593.:
- [69] WK Wong, RK Taira, HK Huang, Digital archival center: Implementation for a radiology department, *American Journal. Roentgenology*, 1992;12(159):1101-1105.
- [70] STC Wong, JL Wilson, Extending knowledge-based systems through closely-coupled graphics and windows, *Proc. 2nd Int. Conf. on Industrial and Engineering Applications of AI and Expert Systems*, Tullahoma, TN, 1989.
- [71] STC Wong and JL Wilson. A set of design guidelines for object-oriented deductive systems, *IEEE Trans. Knowledge and Data Engineering*, October, 1993:895-900.
- [72] STC Wong, R Knowlton, HK Huang. The Integration of Automated Image Registration in a PACS environment. *Radiological Society of Northern American (RSNA) - InfoRAD Exhibit*. Chicago, Nov. 1994.
- [73] STC Wong. An open system knowledge framework and its bridge evaluation application. *IEEE Trans. Systems, Man, and Cybernetics*, 1994;24(60):901-917.
- [74] STC Wong. COSMO: A communication scheme for cooperative knowledge-based systems, *IEEE Trans. on Systems, Man, and Cybernetics*, 1993;23(4):809-824.
- [75] RP Woods, JC Mazziotta, and SR Cherry. Automated image registration. In *Quantification of Brain Function*. K Uemura et al. (eds.), Elsevier Science, 1993:391-400.
- [76] RP Woods, JC Mazziotta, and SR Cheery. MRI-PET registration with automated algorithm. *J. Computed Assisted Tomography*, July/Aug., 1993;17(4):536-546.
- [77] R Yavatkar, K Lakshman. Communication support for distributed collaborative applications. *Multimedia Systems*. 1994;2:74-88.
- [78] S. Zink and CC Jaffe. Medical image databases, *Investigative Radiology*, 28(4):366-372, 1993.

## **Design Methods and Architectural Issues of Integrated Medical Image Database Systems**

Stephen T.C. Wong, H.K. Huang

Laboratory for Radiological Informatics, the University of California, San Francisco  
Box 0840, 530 Parnassus Avenue, Rm 158-CL, San Francisco, CA 94143-0628.

### ***Abstract***

The past twenty years have seen tremendous changes in medical imaging techniques. New modalities and protocols are expanding the available digital image data at a rapid rate. Yet a framework for gathering, managing, and using multimodal image information in an integrated database environment is missing. The purpose of this paper is to present the experience of implementing integrated medical image database system at UCSF. We discuss the general system architecture, software design methods, and specific database tools and illustrate them with application examples.

Two immediate issues confounding the building of medical image database systems are: lack of supporting infrastructure and inability to index images by content. To circumvent these problems, the evolutionary medical image database system being implemented at UCSF is based on a three-tiered client-server architecture: client medical workstations, database application servers, and a hospital-integrated picture archiving and communication system (HI-PACS). The approach used is to integrate content-based retrieval and knowledge base techniques within the existing HI-PACS to make the whole database system useful in medicine.

## **1. Introduction**

Medical images are at the heart of the patient's diagnosis, therapy treatment, surgical planning, and long term follow-up for outcome assessment. Medical imaging is becoming increasingly digital and diverse in scope. In the past two decades, we have witnessed tremendous changes – new techniques include: ultrasound, X-ray computed tomography (CT), magnetic resonance imaging (MRI), digital subtraction angiography (DSA), magnetic source imaging (MSI), positron emission tomography (PET), and single photon emission computed tomography (SPECT), to name just a few. These digital imaging modalities, currently constitute about 30% of medical imaging examinations of the nation, have revolutionized the means to acquire patient images. They have provided noninvasive means to view anatomical cross sections and physiological states, and reduced patient radiation dose and examination trauma. The other 70% of imaging examinations on skull, chest, breast, abdomen, and bone are done in conventional x-rays or computed radiography (CR) [7]. Different kinds of film digitizers, such as laser scanner, solid-state camera, drum scanner, and video camera, can be used routinely to convert plain X-ray films into digital format for processing.

These advances of imaging technology will certainly continue. However, reorganizations and reengineering of the health care system is shifting the medical imaging focus from the generation and acquisition of images to the post-processing and management of images. The motive is to realize the greatest possible benefit from the data that already exist. The new megachanges of the next twenty years will center around gathering, managing, and using multimedia clinical information.

In the past, medical imaging researchers have been working on individual modalities in isolation while the medical informatics community focused primarily on textual information processing. Thus, in spite of their importance, medical images are poorly incorporated into the overall collection of patient data. Today, tools for archiving, communicating, classifying, organizing, and retrieving medical images to take full advantage of their rich information content are nonexistent [17,2]. This hinders the efficient use of digital modalities in routine medical practice and in improved patient care, such as

retrieving multimodal images in real time for use in telemedicine or in virtual reality environments for surgery planning.

Recent technological advances in computers and networks, the press for improved access to health care, the expectation of having health care provided at lower cost, and the presence of recognized and implemented standards, all bring forth the effort of integrating diverse imaging and data sources into a vertically integrated database system, available for querying and analysis. As the integrated database system expands, it will facilitate the correlation of massive multimedia clinical data to advance the biomedical knowledge. The creation of this framework and the assessment of its performance is considered to be the next important milestone in medical imaging [26]. The success of this endeavor will make significant contribution to the national information infrastructure for multimedia health care.

The purpose of this paper is to present the design approach and implementation experience of integrated image database system at the UCSF digital radiology environment. This paper is organized as follows: Section 2 summarizes the new challenging issues of medical image database systems. To address these issues systematically, in Section 3, we describe the design approach and general architecture of the integrated medical image database system at UCSF. Section 4 discusses computing methods and tools that are found to be essential during our development. These application-specific tools are illustrated with implemented examples. Finally, in Section 5, we summarize our presentation and future work.

## 2. Definitions and Challenges

A medical image database system (MIDS) manages a large amount of heterogeneous, changing, pictorial, and symbolic data and provides a means to process, query, and manipulate these medical data in an efficient manner to facilitate optimal decision making in a health care environment. A properly organized imaging database not only can compensate for limitations in human memory and in film recording, but also open up many new vistas to improve patient care, research, teaching, and administration.

Considerable effort has been spent on developing textual medical databases in the past three decades [2]. On the other hand, the integration of medical images with textual patient data in database systems has hardly begun. This section summarizes technical issues facing this new form of health care information management. While some of them are common to other image database applications [5,6,12,13,14], the different nature and acquisition of medical images and the legal-medical implications in using medical images for patient care suggest that specific solutions are needed.

**2.1. Multimodality.** About two decades ago, two-dimensional projectional radiography was the only imaging modality used in clinical examination. Today, a plethora of sectional imaging modalities generate volumetric medical images containing a great variety of previously unattainable biomedical information, e.g., anatomical, geometrical, and spatial, biochemical, physiological, of different body organs. Features and information contained in these images are diverse and interrelated in complex ways that are not always easy to interpret and to represent. Geometric considerations, such as shape and location [14,20], are as important as organ functionality, as in the heart and liver. Physiological changes, such as the rate of cerebral blood flow, may have subsequent effect on body biochemical activities, such as the concentration of neural metabolites.

**2.2 Data Heterogeneity.** Medical image data are heterogeneous in how they are collected, distributed, and displayed. Images are acquired from scanners of different modalities and in different positions, represented in internal data formats that vary by modality and manufacturer, and displayed in distributed remote workstations with divergent appearance, orientation, and spatial resolution. PET and CT images appear entirely different from one another and are also distinct from other modalities, such as MRI, and ultrasound. Even within the same modality and for the same anatomy, two sets of medical images can vary greatly in slice thickness, dataset orientation, scanning range, and data representation. It is worth noting that diagnostic images are acquired and displayed in gray scale. Thus, color

indexing, a popular technique used in other application domains [3], is of little use for biomedical images, except for color microscopy, doppler ultrasound, and pseudo color enhancement in nuclear medicine modalities.

**2.3. Structural and Functional Contexts.** Structural information in a medical image contributes essential knowledge of the disease state and clinical decisions. For example, the location of a tumor and its adjacent anatomical structures have profound implications in treatment planning (spatial context); whereas monitoring of growth or shrinkage of that tumor is an important indicator of the patient's progress in therapy (geometric context). What distinguishes medical images most from other types of digital images, however, is the ability to represent functional information, e.g., biochemistry and physiology, of body parts. As an example, temporal lobe hypometabolism exhibited in SPECT scans, if not coincide with the tissue atrophy of the suspected region shown in MRI images, may indicate that a different epileptic focus as otherwise would be concluded.

**2.4. Imprecision.** Images contain information that is often imprecise and can only be expressed informally. This phenomenon applies to entire data sets as well as features within individual images. The boundary of a tumor, for example, may be fuzzy, depending on the resolution of the imaging modality that is used. Yet, the accurate delineation of a tumor is an essential first step for curative therapy. Worst, the imprecision of image resolution is often compounded by the vague description of features extracted from the images or ambiguous coding of disease classification in diagnostic reports. For instance, the diagnosis "acute myocardial infarction, anterior wall" imparts a sense of certainty and specificity, and can be linked to ICD-9 (International Classification of Diseases, ninth revision) code 410.10. On the other hand, for the diagnosis "left ventricular aneurysm," an agreeable definition has yet to reach within the medical community. New, more expressive medical data dictionary and representation models that can cope with such imprecision are needed.

**2.5. Temporal.** Staging of the disease state and the monitoring of patient progress over time are fundamental to diagnostic and therapeutic decisions and to outcome assessment in long term follow up. The ability to define and track temporal relation in the medical image sets of a patient taken at different periods, together with the medical history of that patient, is an essential component of medical image databases. As the database expands to incorporate a large mass of similar patient data, tools developed for the intra-subject temporal monitoring can also be extended to inter-subject temporal comparison to improve prognostic and diagnostic outcomes.

**2.6. Large Datasets.** Like other imaging applications, medical images are distinct from textual records in their large data sizes. Images acquired in one examination can range from one or two megabytes in nuclear medicine to 32 megabytes or more in mammograms and digital radiographs. Digital medical images accumulated in a 1000 bed hospital are estimated on the order of a terabyte a year [15]. These imaging modalities generate such a large volume of image data that traditional methods for information management are no longer adequate. More advanced tools and schemes to manage multimodal images and their complementary textual information for database processing are needed.

**2.7. Infrastructure Support.** The accessibility to heterogeneous images, patient data, and medical reference materials from diverse sources in the clinical decision making process presupposes the existence of a communication infrastructure. The reality, however, is that most databases and information systems in hospitals today are disjoint, under the control of different medical sections and departments, and often do not communicate with external reference database networks well. The isolation creates many technological and administrative barriers that make the gathering of information cohesively from different medical images extremely difficult [8,2]. This lack of supporting infrastructure causes the concept of multimedia data integration far from practical deployment.

**2.8. Security.** Along with the multimedia data integration and communication comes a new issue: How do we establish integrity and privacy in medical records that exist only in the easily altered digital storage of computers. This issue is especially prominent when an image database framework is connecting to national and international health care networks. An open yet insecure system is unlikely to be acceptable by the clinicians due to legal-medical complications. The existing means of security

measures, such as checking for personal ID or authorization password, enforcing administrative policy, and prioritizing user privileges, all have known problems in ubiquitous networking environment [9,18].

### 3. UCSF Image Database System Architecture

Among the new issues confounding the building of medical image database systems, two fundamental ones stand out: Lack of supporting infrastructure and inability to index information by image features. These two issues form the core of our research of integrated medical image database system (MIDS). A MIDS is a repository of integrated information, i.e., various kinds of medical images and clinical text, available for querying and analysis. As relevant information becomes available or is updated, the information is extracted from its source, translated into a common data model, and integrated with existing data at the MIDS. Within this integrated medical system, queries can be answered and data analysis can be performed quickly and efficiently since the information is directly available, with model and semantic differences already resolved. Furthermore, MIDS data can be accessed without slowing down the operations of original information sources, such as a radiology information system or an MR imaging system, and is available even when these information sources are inaccessible.

The concept of integrated MIDS is similar to the notion of data warehouse discussed in general database literature [10,11]. The advantage of radiologic imaging is that the raw image and patient reports, once generated, are rarely modified. The update is limited to a subset of processed and extracted data for decision support queries. Hence, MIDS avoids much of the difficulty in providing a consistent database source compared to business or engineering applications. In radiology, this information framework is a logical progression from the primary digital modalities and the stand-alone medical information systems towards the networked multimedia database systems. Building a comprehensive, enterprise framework that encompasses generation, acquisition, processing, management, and communication of radiologic images is an evolutionary process. System design must provide the capability to phase out old components and incorporate new devices and technology as the system evolves.

**3.1. Hospital Integrated Picture Archiving and Communication System.** The MIDS being developed at UCSF is based on a three-tiered client-server architecture. This architecture divides applications into parts that run on different types of computers: a picture archiving and communication system (PACS), an image database server (IDBS), and client workstations for specific medical applications. A PACS is a system integration of many components, including image acquisition devices, computers, networks, display stations, medical databases, and legacy information systems. Figure 1 shows the hospital-integrated picture archiving and communication system (HI-PACS) implemented at UCSF. The system contains a centralized, dynamic, and massive storage of multimedia clinical data acquired from various imaging scanners and databases across the UC hospitals.

-- Figure 1 --

The implementation of this HI-PACS emphasizes recognized and implemented standards, open systems connectivity, hierarchy memory management, database integration, and security. The UCSF PACS connects not only various imaging scanners directly, but also smaller scale PACS of subspecialty sections and textual information sources, such as hospital information system (HIS), radiology information system (RIS), and library information system (LIS). The PACS fiber-optic backbone has been installed interconnecting four major facilities in UCSF, i.e., Moffitt and Long hospitals, ambulatory care clinic, and campus library, to provide a data warehouse integrating information from multiple imaging and data sources.

Wide area connections between UCSF and affiliated hospitals in the San Francisco Bay area using asynchronous transmission mode (ATM - OC-3c, 155 Megabits per second (Mbps) for real-time

telemedicine) are being implemented. Image communications between UCSF hospitals using frame relay technology (T1 link - 1.54 Mbps) for non real time telemanagement have been in operation. The data warehouse at the departmental PACS cluster can store up to 1.3 terabytes of image and text data. These PACS data can be retrieved and transmitted over the networks to high resolution, multi-monitor display stations located in intensive care units and radiology reading rooms for clinical review.

The intent of picture archiving and communication systems is to provide management and fast review of vast volumes of image and text files in a digital radiology department. The PACS file management system can only handle query by artificial keys, such as patient name or a hospital ID. It lacks the means to classifying and indexing these data files by their information contents. Thus much of the rich, useful patient information stored in PACS has been not tapped for broader medical practice, research, and teaching. To provide genuine database management capabilities, we develop an image database server that can extract relevant features from PACS images and textual reports and organize the extracted data, together with annotated image slices, into a common data model. The image database server can also process various computer-aided medical applications, such as fast indexing of medical records, interactive visualization of images, and activate rules for decision making. It consists of a set of sharable, multitasking components that interact with client workstations and PACS. Section 3.2 discusses this service tier in more detail.

The client workstation provides a graphical user interface (GUI) to display multimedia information and communicate with the image database server for specific biomedical application processing, such as computer-aided diagnosis and surgical planning [22,23]. Screens and GUI routines are managed in a repository as objects to permit a modular approach to development and foster reuse of objects. In addition, tools that are not necessary for traditional textual databases but are perceived as essential will then be developed and integrated into the overall framework. This three-tiered system development is a bottom up endeavor, progressing from building the rudimentary level of image acquisition, archiving, and communication to providing sophisticated image database management. It lets the developers and system administrators to separate the specific application logic from general processing logic. This modularity allows new services and clinical changes to be more rapidly incorporated into client applications.

**3.2. Image Database Server.** The image database server (IDBS) is a centrally located core of software objects or application programs for accessing, processing, managing, and manipulating data from distributed information sources in the HI-PACS environment as if they were from a single database. This middleware database server, resided in between HI-PACS and client medical workstations, consists of three cooperating processes or engines, running in parallel. It also serves to filter all client requests through various, centrally controlled levels of security, such as IP (Internet Protocol) firewalls, access control, and image authentication. Figure 2 shows the key components in our medical image database framework.

-- Figure 2 --

(1) **Database Engine** that contains a collection of objects and a set of operations that they can perform. Data models are devised to classify, organize, and represent images and text to support information retrieval by image content for various biomedical applications. The database engine treats system components and multimedia data of the entire database framework as a logical conglomeration of distributed, interacting objects of various levels of granularity.

For rapid prototyping purpose, we decided to use a commercially available relational database management system and then to develop customized object software modules on top of the commercial system for special applications. The current database engine used is based on the object relational technology – Sybase database management system (Sybase Inc., Emeryville, CA) for the underlying persistent data storage.

- (2) **Image Processing Engine** that provides image processing services for various medical applications. These services include computationally intensive operations, such as image registration and feature extraction. The decision to locate heavy duty imaging application routines in the powerful image database server is to reduce the hardware and software requirements on the client workstations. The medical image processing and visualization software packages used include VIDA (University of Iowa College of Medicine, Iowa City, IA), 3DVIEWNIX (Hospital of the University of Pennsylvania, Philadelphia, PA), and several in-house developed programs [16,19,21,22,23].
- (3) **Knowledge Base Engine** that consists of three kinds of knowledge base modules. Rule-based modules that encode general and specific medical rules for decision making, case-based modules that aid the navigation of database through similar or reference imaging cases, and model-based modules that provide interactive work-up models of different disease categories closely involving physicians making joint decisions. Our current knowledge base engine includes a mix of C routines, embedded SQL (Structured Query Language) statements, and GEL (Gain Extension Language) scripts of Gain Momentum.
- (4) **GUI Client Workstations** each provides a graphical user interface (GUI) for the local user to request applications of the image database server to perform information retrieval, manipulation, analysis, and update. Each client workstation also contains simple image processing capabilities, such as thresholding and filtering, but accesses the IDBS for more sophisticated database services.

Our current database front-end development tools include Gain Momentum and several Web browsers, such as Mosaic and Netscape. Gain Momentum is a Motif (UNIX) application development environment with an underlying object-oriented database (Gain Technology, Sybase Inc., Palo Alto, CA). The underlying object-oriented database contains the objects that allow the client applications to access, manipulate, or present the underlying multimedia data in the server. Web browsers are efficient hypermedia navigation tools for "public domain" file documents. To allow complex, hierarchical representation of database information, however, we have to incorporate new software tools with extensive multimedia querying capabilities into the web browsers. Due to the complex programming effort involved, our current web-based database application is limited to simple teaching file queries [24].

- (5) **HI-PACS** consists of the PACS multimedia file storehouse and a federation of distributed, heterogeneous medical databases, imaging acquisition devices, and hospital information systems, providing the environment in which the above system components run.

The IDBS prototype is being developed on a 4-processor SUN 690 MP SPARC computer. Security mechanisms are set up to allow the database server to access the multimedia information of HI-PACS with fine-grained control over individual transactions and data items, yet without running the risk of contaminating daily clinical services.

Figure 2 provides a generic architecture of integrated medical image database system. Since medical technology research is driven by well-defined hypotheses, our development focuses on particular clinical applications that have specific kinds of representation and queries applied to medical images. The applications currently under development are: surgical planning of epilepsy [23], neurological diagnosis of dysmyelination and multiple sclerosis, and the age assessment of hand bone [19]. By restricting the scope of applications and the types of images, we will provide some useful though initially limited functionality for content-based retrieval of medical images. The formation of these medical image databases follows knowledge and feature-based indexing methods [1,3], but adds sophisticated services in automatic image registration, image feature extraction, multimodal image presentation, and multimedia client server application development in HI-PACS environment.

## 4. Medical Imaging Database Tools

Section 3 describes the general architecture and design of our integrated image information system. This section discusses specific tools being developed within this framework for various medical applications. The key idea behind the MIDS approach is to extract, filter, and integrate relevant information in advance of queries. Extraction of keywords and phrases from textual reports into medical databases can be done automatically via string pattern recognition. On the other hand, current medical image databases allow queries to be formed based on the qualitative features of images that have been manually annotated.

For content-based information query, the next generation of medical image databases must have capabilities to recognize and quantitate image content and merge the quantitated image data with textual patient data into a common data model. The content of a medical image includes three levels of abstraction: (i) features or attributes, such as the texture pattern of an organ, the shape of a breast tumor, and the signal strength of brain metabolites; (ii) structural objects, such as image segments corresponding to anatomical regions; and (iii) semantic information, such as the types of relationships between two objects.

-- Figure 3 --

Figure 3 illustrates the operational steps that the MIDS taken to extract image and textual features from the on line PACS data. The feature extraction is based on the *a priori* approach, rather than the dynamic and automatic feature extraction during user query as proposed in many non-medical image database applications [4,6,12].

The extraction and composition of textual data from the diagnostic reports and the DICOM (Digital Imaging and Communications in Medicine) header fields can be automatic, whereas the segmentation and extraction of medical images often are done interactively. Further, we divided the image features into primitive and logical. Primitive features are directly obtained from the medical images, such as volume, shape, and texture of certain organs in CT images or metabolic activities of brain tissue in PET scans. Logical features are abstract representations of images at various levels of detail and deeper domain semantics. For example, whether the volume of an anatomic structure is normal or whether certain brain tissue is hypometabolic in reference to certain established data. These logical features are synthesized from primitive ones and additional domain knowledge. All features extracted are entered into appropriate attributes of the objects defined in a data model to facilitate subsequent information query by content.

In what follows, we describe the application-specific tools that are found to be essential in the MIDS development. Our emphasis at this stage will be on those which have direct impact on the building of a working prototype, rather than on query performance and system administration. Few of these tools are currently supported in commercial systems, and many of them are still awaiting efficient solutions from the database community. In addition, examples are taken from our applications to better illustrate the methods we used, as well as to highlight the unresolved problem areas.

**4.1. Image Data Modeling.** *An image database system must be able to transform pixel based representation of medical images and their contents into object based representation for symbolic manipulation. This transformation is a crucial step for building any content-based pictorial query system.*

Once medical images can be efficiently represented in a content-based manner, by objects rather than pixels, then a wealth of established database theory, methods, and tools can be readily applied for content-based access, creation, and manipulation. In addition, clinical decisions are based on not only information obtained from medical images, but also from textual and graphical data of the patients, including medical history, physical examination, laboratory tests, diagnostic workups, and pathological results. These images and data from diverse sources must be the integrated with the physicians' medical knowledge and available reference materials to deal with therapeutic or diagnostic decision at

hand. More expressive data models than those used in textual databases are needed in modeling and indexing such multimedia medical data. Relational approach is perceived to be ineffective in modeling multimedia medical data. We have taken the object oriented approach to represent anatomical structures.

Figure 4 shows an object data model for hand bone structures [19,1] and a computed radiography (CR) image instance of left hand in posterior-anterior (PA) view. Often, textual data indices are extracted automatically from associated patient reports to complement the descriptors of image object features.

-- Figure 4 --

Object data models also provide flexible means of indexing and linking both images and text and can incorporate biomedical production rules for decision aids. Let us consider the following query for determining bone normality of hand computed radiographs.

*(Q1) Assess the skeletal age of a pediatric patient based on the given CR left hand image and decide whether the skeletal maturity is normal.*

The assessment of skeletal maturity is an important subject in pediatric radiology. Database solution to such a problem is thus of significant clinical value. This query, however, is vague; only the information about the patient's CR hand images is given and "normal" is an ambiguous term in medicine. A computerized method has been developed in our research [19] to assess bone age, based on measurements from different structure of the hand bone. This method consists of the following steps: (1) separate the third finger image from the given hand image pattern, (2) measure the lengths of the distal, middle, and proximal phalanx, (3) convert the measurement into age using the standard phalangeal length table in medicine, and (4) compare the estimated age with the patient's age. The skeletal maturity is judged to be normal if the patient's age is within the limit of standard deviation.

These computational steps can be modeled as production rules in third finger object class of the skeletal model as shown in Figure 4. The image database system thus directs an ambiguous query of the type similar to query (Q1) to the third finger object. If there is no pre-calculated value in the patient finger object's attributes, that object instance would invoke the encapsulated production rules to extract the features and analyze the results.

The hand bone application is an example of content-based medical image database using a single imaging modality. The tools and data models developed, though optimized for this specific research, are restricted because they are derived from digital radiography only. Daily medical decisions often desire the gathering and correlation of information from several types of patient images as well as from various data sources. In addition, the correlation of multimodal images of the same body organ may be more effectively modeled by biomedical relationships between objects of different modalities. Object-relational approach, such as those advocated in UniSQL™ (UniSQL Inc., Austin, Texas) and Illustra™ (Illustra Technology, Oakland, CA), may offer alternate solutions in multimodal image modeling than pure object oriented approach.

**4.2. Image Fusion.** *If the pixel-object transformation is the crucial step for realizing content-based representation, the fusion of volumetric image pairs is the key to enabling multimodal image databases.*

The fusion of images consists of two major steps: (i) the accurate registration of images from different kinds of imaging modalities, different periods of the same patient, or different patients, and (ii) the correlation and extraction of biomedical information represented in registered images to derive new, previously unavailable medical knowledge.

-- Figure 5 --

**4.2.1. Medical Image Registration.** Figure 5 provides a case of the use of image registration to align two sets of MRI images taken from a 30 year old patient suffering multiple sclerosis; a disease condition marked by patches of hardened tissue in the brain or the spinal cord and associated especially with partial or complete paralysis and jerking muscle tremor. The top row shows the patient's brain scan taken a year ago while the middle row is the recent one. The objective is to look for differences between the two scans to evaluate changes related to disease progression (marked by the sizes of high intensity dark regions around the ventricular regions of the brain). Using conventional films, the radiologist has to rely on imagination to correct for discrepancies in head position that result from different positioning in the two scans. Even so, he or she cannot compare changes of lesions in top scan to middle scan.

The IDBS provides automatic image retrieval and registration services for neuro images. A client neurological workstation requests the database server to retrieve the specified image pair stored in PACS in real time and automatically align the first image based on its 3D dataset with the recent one. The result is displayed in the bottom row, which indicates that the disease condition of the patient has not worsen, as the size of high intensity regions, i.e., lesions, around the ventricular areas remain the same. This example also illustrates the use of the temporal relation of medical images in monitoring the patient's progress or response to treatment.

We adopted the clinically validated Woods' algorithms for automatic image registration [25] and since then have extended the registration routines in the image processing engine to cover the inter- and intra-registration of four different kinds of brain images, i.e., MRI, PET, magnetic resonance spectroscopy (MRS), and magnetoencephalography (MEG) [23]. Automated image registration with on-line access to a vast mass of image data makes possible many value-added functions that cannot be done in the conventional medical record, such as looking for subtle changes in signal across time, combining different types of functional (such as PET and SPECT images) and structural information (such as CT and MRI images), and comparing different subjects or patients to one another for better clinical decision making. What is more important, it prepares the stage for the second step of image fusion, i.e., providing a means of correlating and quantitating content information of different images.

-- Figure 6 --

**4.2.2. Image Correlation and Features Extraction.** The correlated image datasets provide diagnostic insights not readily apparent when examining the imaging modalities independently. They are encoded into the targeted data model to serve as definitive indexing in image query. For instance, registration of functional images obtained using nuclear medicine modalities, such as PET or SPECT, with anatomic MRI images is of particular important in neurological diagnosis, since it allows the intrinsically better spatial resolution of the MRI image to be used in interpreting the functional activities of the brain.

Figure 6 illustrates such a case of MRI defined regions of interest (ROI) registered with PET. In this image the left and right amygdala are delineated. The clear definition of anatomical volumes for MRI allows quantitative analysis of functional information (metabolic count of glucose consumption) captured in PET scans. Various methods for specifying and manipulating three dimensional regions for quantitative analysis of functional brain images, such as PET, MRS, and MEG, based on the structural MRI images have been implemented in our MIDS testbed [16,22].

The extracted quantitative data will be important to understand the clinical implementation of a particular therapeutic or diagnostic decision while providing exact numerical indices for ad hoc queries by image features, e.g., MRI anatomical volume or PET glucose uptake count. Similar structural mapping with other functional images, such as the concentration of energy metabolites measured in magnetic resonance spectroscopy, will reveal a different aspect of biomedical information and will serve yet another kind of indexing for image queries.

-- Figure 7 --

Figure 7 shows a multimedia user interface developed in our client neurological workstation. This user interface enables real time retrieval of PACS images for comparison and reference during diagnostic workup of epilepsy surgery candidates. Query of the image database can be by image content, e.g., MRI hippocampal volume or PET glucose concentration, at the lower left window, by patient attributes, e.g., name, age, and sex, at the upper left window, or a combination of these features. The sliders in the "Query by Image Attributes" window can be used to specify the range of the image attributes for data retrieval. The neuroimaging database returns with a list of seven patients satisfying the combined image and patient attribute constraints. The user can select any patient name to retrieve and display MRI and PET images (in this case, the third patient). By default, only two pairs of representative, co-registered image slices are displayed. Once the images of interest are retrieved, the user can invoke the 3D image visualization routines of the image database server for more sophisticated image rendering and manipulation.

-- Figure 8 --

Figure 8 shows yet another kind of content-based information query. The user first specifies the structural, functional, and textual attributes of the MRI studies of interest. The image database server returns with a list of patients satisfying the query constraints, and, in addition, a set of representative picture of interests of these patient cases in thumbnail form. The user then clicks on the particular thumbnail image to zoom it up to the full size or to directly retrieve the complete 3-D MR dataset for further study.

Image fusion being an integrated part of medical image databases is just in the beginning. Tools for rapid implementation of automated image registration procedures are still lacking. In particular, tools for the accurate, automatic image registration of non-rigid body organs, such as prostate gland and lung, are non-existent. Owing to this limitation, our current applications are restricted to rigid body organs or organ encapsulated by non-moving solid boundary, such as hands and brain. On the other hand, medical researchers are still trying to understand many biomedical implications and opportunities of the new, powerful information derived in the image fusion process. The establishment of medical image databases for large scale statistical analysis or "data mining" will in turn contribute to the advances of biomedical knowledge and clinical practice.

**4.3. Case Based Reasoning.** *Pixel-object transformation and multimodal image fusion form the foundation of content-based image retrieval, but they must be integrated with case-based reasoning techniques for medical decision making to make the entire image database systems useful in medicine.*

Reasoning by previous patient cases is an important characteristic in medical diagnosis. A senior physician accumulates a growing set of past medical cases experienced over years of practice and can match and retrieve relevant ones quickly and accurately for the clinical problems at hand. To model this pattern of decision making, case-based reasoning techniques can be incorporated to improve the efficiency of searching and retrieving relevant information from a vast storehouse of patient cases. This can be as simple as providing "normal" examples of a feature or structure for purposes of comparing the one under investigation to learn whether it falls within normal limits. Tools that support flexible shape and pattern matching will be needed.

On the other hand, the user can develop possible hypotheses to explore and navigate through the image database. Thus, general tools that allow for "show me one like this shape but is smaller" or "find me something of size in between the two" will provide powerful means of synthesizing new concepts and knowledge. This "changing the field of view" approach is perceived as an important attribute of a medical image database [26]. Sometimes, fundamentally different pathologic processes may exhibit similar image configurations. This phenomenon is known to give rise to what is called the "differential diagnosis," that is, narrowing down the disease diagnosis by systematically comparing with known cases and images of suspected similarity. For example, consider the following two queries to help analyzing a multiple sclerosis case:

*(Q2) Find all MR head examinations of patients with multiple sclerosis which show white spots in the MR scans, categorize all the white spots found by size and locations.*

*(Q3) Locate and find the size of white spots in this current MR examination and match them to the cases found in query (Q2) as close as possible.*

MR images of T2-weighted sequences are normally used to identify white spots or fiber deposits in the brains of multiple sclerosis patients. The white spots are usually found within white matter, particularly the periventricular area around the lateral ventricles, the optic nerves, brain stem, and spinal cord. This application requires the extraction and quantitation of white spot volumes in these brain regions.

-- Figure 9 --

Figure 9 illustrates a magnetic resonance imaging case of a patient with severe multiple sclerosis caused by abnormal amount of large fiber deposits in the periventricular region. The segmented objects are white spots found in the region and their volumetric sizes are quantitated using the similar set of feature extraction services discussed in Section 5.2. A query such as (Q2) thus indexes pre-processed MR examination cases through these symbolic objects. The following (Q3) then filters the findings with the current images under investigation. The added difficulty of this query is the fuzziness of the phrase "as close as possible." In this case, the heuristic rule used is to sum up the volumes of all white spots found in the periventricular area and then use this sum to match the first five cases with similar white spot volumes in the periventricular region.

**4.4. Customized Schema and Languages.** *New tools will be needed to facilitate ad hoc and customized schema for a broad range of disease applications so that the schema can be edited, modified, and adapted to new queries.*

In DBMS the schema represents the information model of the application. In object oriented databases the schema represents the various classes, the attributes (fields, instance variables), and the relationships between the classes. In relational databases the schema defines the tables, both base and derived; various integrity constraints on the tables; and other meta-information of the database extensions. The schema therefore represents the structure of the information such that a data manipulation language can be used to query the information stored.

Database languages for multimedia applications has been discussed widely in the literature [6,12,13,14]. Here we highlight some of the schema issues in medical imaging domain. Clinicians view and use the same medical image database in a vastly different manner. For example, a radiation oncologist may ask the database system about the physical shape and spatial position of tumor regions targeted for therapy, whereas the diagnostician may query the functional and morphological features of the tumors for staging of the tumor. In general, image queries in medicine can be categorized as follows:

- **Clinical practice** concerns the daily delivery of health care such as diagnosis and therapy. Some typical image queries include: (1) retrieve the current imaging studies and/or previous imaging studies (may be of different modalities) of a patient; (2) retrieve all imaging studies which show common (or different) image features of the current case; and (3) inquire about biomedical attributes of an object or a group of objects of interests, such as size of a lung nodule and glucose concentration in a brain region.
- **Education** concerns the teaching of graduate and medical students, residents, and fellows. Some desired query capabilities include: (1) retrieve pathologically proven medical images according to a particular type of disease; (2) step through a sequence of examination questions and medical images dynamically based on the user response to the questions; and (3) ask the image database to provide

another set of example images by modifying certain descriptors, such as severity of disease and age of patient population, of a teaching case.

- **Research** uses databases to support large scale, systematic study of diseases and medical procedures. Typical queries involve mass screening a large number of clinical imaging studies from various patient groups, i.e., sex, age, race, and medical history. Given the size and region of a brain tumor, for example, queries may be issued to provide statistics on the improvement of patients operated under various treatment methods. On the other hand, a hypothesis (e.g., a certain drug is effective in controlling multiple sclerosis) can be composed and then be validated by a series of queries sent to the database system (e.g., measure the variation in the sizes and numbers of white spots shown in MR head images of all multiple sclerosis patients). Such queries lead to new biomedical knowledge.

Thus, the schema of an integrated database system should be able to evolve over time to serve the expanding scope of clinical applications. Management tools should be provided to cope with such changes. The description of a query language will be dependent on the underlying data model used. Standard query languages are encouraged to ensure the interoperability of medical image databases with medical information systems using similar data models, e.g., SQL (structured query language) for relational DBMS and OQL (object query language) for object oriented DBMS. Iconic or pictorial queries built on top of these standard languages are essential to reap the full benefit of medical image databases, and this is the direction of our GUI development effort. The use of object icons and graphical relations provides an efficient means to develop customized queries and better conceptualization for the problem at hand.

**4.5. Image Manipulation and Communication.** *A medical image database system must incorporate image processing functions into its symbolic processing capabilities so that, ultimately, image database engines manage the data or information stored and manipulated by these application tools.*

Unless image database systems can acquire and manipulate image pixels as easy as they have done for alphanumeric symbols, the systems developed will not be of any practical use. These new functions include: (1) manipulate multimodal images, such as zoom, pan, rotation, contrast enhancement, region of interest contours; (2) recognize and segment images, such as edge detection, morphologically filtering, and feature extraction; (3) visualize images, such as surface and volume rendering, texture discrimination, mosaic, cine, and multiple and multimodal image display; and (4) communicate images, such as compression for speedy transmission and real time protocols.

Simple image processing tools, such as edge enhancement and filtering, used in our client medical workstations are inherited from the existing PACS display stations. More computational intensive tasks, such as visualization and volume rendering, are delegated to the image processing engine of the IDB server. However, existing image segmentation techniques become useless when the boundary between the object of interest and the background is fuzzy. The segmentation of the brain objects in our applications is being done interactively. This constitutes the most time consuming operation in our development. Efficient tools for automatic feature recognition of fuzzy images, image registration of non rigid objects, and real time transmission of volumetric image datasets over broadband networks are particularly lacking.

**4.6. Data Integrity and Privacy.** *Integrity and privacy of electronic medical records in the integrated image information environment can be divided into the database server level and file transfer level.*

PACS file servers can provide data security, but at the file and directory level. This means various groups and users can be given or denied access to certain patient files or directories. Some directories might not be "visible" to certain users. Image database servers, on the other hand, can provide security at the object level. Users can be granted various privileges on logical data objects and collections of objects. The granted privilege is also more diverse in scope. It might be for read or update operations, create operations, or execution of methods associated with persistent objects. With a

file-server architecture such as PACS, the PACS controller software checks and imposes the file- and directory-level security that executes on the PACS controller node. With the MIDS database-server architecture, the software that stores and executes database-server security resides and is executed on the middleware image database server node. This provides more flexibility and fine grain control on the different levels of user-access privilege.

Another level of security concerns the transmission of queried image data to the requested client workstation via the local or wide area networks. Cryptography is considered as one of the strongest methods to ensure data security and authenticity. MIDS can provide software encryption using the public-key algorithms and a combination of public- and private-key algorithms for data privacy. The key length used ranges from 256 bits to 1024 bits. Software encryption, however, adds considerable computation time for 3-D image datasets or computed radiographs [21]. It thus does not support real time delivery of medical images. For example, encrypting a volumetric brain MR image of 8 Mbytes takes about 2 minutes in a Sun SPARC LX workstation (transmission time is not a factor with our 155 Mbps ATM links). Nevertheless, encrypting a few selected MR image slices (256 x 256 x 2 bytes/slice) for the referring physician spends over one or two seconds. The latter thus is amenable for real time clinical practice.

In certain clinical scenarios, such as in emergency room or during surgery operation, the integrated information framework should provide real time transmission of full volume datasets as well as the checking of data integrity or authenticity [21]. One approach adapted in MIDS is to create a fingerprint of the image. The fingerprint is a 128-bit integer unique to that image created by an one-way hashing function. The calculation of fingerprint takes less than a fraction of seconds in our Unix workstations. The fingerprint can be either embedded in a selected attribute field of the image header or sent together as a separate data file. The receiving site can verify the authenticity of information received by calculating the fingerprint value of the received image and compare that with the received fingerprint. Any discrepancy between the two indicates the tampering of the image during the transmission. Note that the fingerprint method is a tradeoff between timing performance and data privacy. It detects unauthorized modification to the image data but does not prevent the image from unauthorized access.

## 5. Conclusions

Image database management with content-based retrieval will provide an environment for effective clinical decision making. It will open up many opportunities in advancing medical services and reducing health care costs. There, however, lacks a clear definition of medical image database systems; let alone developing much needed tools that cannot be found in traditional databases.

The purpose of this paper is to present our design approach and implementation experience of image data warehouse at UCSF. We identified challenging issues and necessary tools for this new class of database systems. We provided a generic architecture of the vertically integrated medical image database system being developed at UCSF. We described the design methods and application-specific tools developed within this information framework and illustrated them with implemented examples.

Among the new database issues, the infrastructure support for medical data integration and content-based indexing of medical images, are the most immediate ones that must be first resolved. For the former issue, HI-PACS provides a genuine testbed for medical image database prototyping by removing the administrative and technological barriers of gathering information from diverse imaging and data sources. For the latter issue, the *a priori* approach is used to extract relevant features of PACS images and reports. These processed data then are translated and integrated into a common data model in the IDBS for enabling content-based information retrieval. The integration of content-based retrieval of relevant images and knowledge base techniques for decision making in the existing PACS installations will make the whole system much more useful in medical domain. Currently, tools for automatic mapping and correlating images of non-rigid body organs are still lacking.

In addition to current applications in epilepsy surgical planning, bone age assessment, and myelination disorders diagnosis, we will continue to identify and develop new MIDS applications driven by sound hypotheses. By gradually enlarging the scope of applications, we will learn more functionality for content-based query processing. The application-specific tools developed in these application can also be extracted and composed into a more generic package for medical image database development.

The research into medical image database systems has just begun. Many fundamental challenges in image information processing have yet to be resolved. Research must proceed in multidisciplinary teams that combine the expertise of database and imaging disciplines with the knowledge of the medical community. The successful implementation of medical image database systems in the health care system will have a direct impact on the cost and quality of patient care, contributing to improved disease diagnosis, therapy treatment, biomedical knowledge, and access to health care.

## References

- [1] AF Cardenas, IT Jeong, et al. The knowledge-based object oriented PICQUERY+ language, *IEEE Trans Knowledge & Data Engineering*, 5(4), Aug., 1993:644-656.
- [2] RS Dick, EB Steen (ed.) *The computer-based patient record - an essential technology for health care*. Institute of Medicine. National Academy Press, Washington D.C., 1991.
- [3] M Flickner, H Sawhney, W Niblack, J Ashley, Q Huang, BDom, M Gorkani, J Hafner, D Lee, D Petkovic, D Steele, and P Yanker. Query by Image and Video Content: The QBIC System. *IEEE Computer*, Sept. 1995:23-32.
- [4] WI Grosky. Multimedia Information Systems. *IEEE Multimedia*, 1(1), 1994:12-24.
- [5] WI Grosky and R Mehrotra. Index-based object recognition in pictorial data management. *Computer Vision, Graphics, and Image Processing*, vol. 52, 1990:416-436.
- [6] V Guidivada and V Raghavan. Content-Based Image Retrieval Systems. *IEEE Computer*, Sept., 1995:18-22.
- [7] HK Huang. *Picture Archiving and Communication Systems in Biomedical Imaging*. VCH Publishers, New York, 1996.
- [8] HK Huang, O Ratib, et al. (eds.) Picture archiving and communication system (PACS). *J. Digital Imaging*, 5(1), 1992:22-25.
- [9] Information Week/Ernst & Young Security Survey. *Information Week*, Nov 27, 1995:32-40.
- [10] WH Inmon. EIS and the data warehouse: a simple approach to building an effective foundation for EIS. *Database Programming & Design*, 5(11):70-73, Nov. 1992.
- [11] WH Inmon and C Kelley. *Rdb/VMS: Developing the Data Warehouse*. QED Publishing Group, Boston, MA, 1993.
- [12] SS Iyengar and RL Kashyap (eds.) Special Section on Image Databases, *IEEE Trans Software Engineering*, 14(5), 1988.
- [13] R Jain (ed.) *NSF Workshop report on visual information management systems*. February 1992.
- [14] R Jain and AP Pentland (eds.) *NSF-ARPA Workshop report on visual information management systems*. June 1995.
- [15] Y Kim, HW Park, DR Haynor. Requirements for PACS workstations. *Proc. 2nd Int. Conference on Image Management and Communication in Patient Care*, Kyoto, IEEE Computer Society Press; 1991:36-41.
- [16] GJ Klein, BW Reutter, RH Huesman, PG Coxson, TF Budinger. Volume-of-Interest specification on arbitrarily resliced volume datasets. *SPIE Medical Imaging Proceedings: Physiology and Function from Multidimensional Images*, San Diego, CA, March, 1995.
- [17] NM Lorenzi, RM Garder, TA Pryor, WW Stead. Medical Informatics and the New Health Care Environment, *J American Medical Informatics Association*, 2(6), 1995:391-392.
- [18] A Pfitzmann and B Pfitzmann. Technical Aspects of Data Protection in Health Care Informatics. In *Advances in Medical Informatics*, IOS Press, 1992:368-386.
- [19] E Pietka, F McNitt-Gray, ML Kuo, HK Huang. Computer-assisted phalangeal analysis in skeletal age assessment. *IEEE Trans Medical Imaging*, 10(4), 1991:616-620.

- [20] HD Tagare, FM Vos, CC Gaffe, and JS Duncan. Arrangement: A Spatial Relation Between Parts for Evaluation Similarity of Tomographic Section. *IEEE Trans Pattern Recognition and Machine Intelligence*, 17(9), 1995:880-893.
- [21] STC Wong, HK Huang. Authenticity techniques for PACS images and records. *SPIE 2435: Medical Imaging Proceedings: PACS Design and Evaluation*. San Diego, March, 1995:68-79.
- [22] STC Wong, RC Knowlton, K Soo Hoo, HK Huang. Use of multidimensional, multimodal imaging and PACS to support neurological diagnosis. *SPIE 2433: Medical Imaging — Function and Physiology of Multidimensional Medical Images*, Feb., 1995:236-247.
- [23] STC Wong, RC Knowlton, RA Hawkin, KL Laxer. Multimodal image fusion for noninvasive surgical planning of epilepsy. Special January 1996 issue of *IEEE Computer Graphics and Applications* on Computers in Surgery and Therapeutic Procedures, 1996, to appear.
- [24] STC Wong, PS Whaley, A Cheong, K SooHoo, J Wang. Interactive Query and Visualization of Medical Image Data on the World Wide Web, Proc. *SPIE Medical Imaging 96*, Newport Beach, CA, Feb., 1996, in press.
- [25] RP Woods, JC Mazziotta, and SR Cherry. MRI-PET registration with automated algorithm. *J Computer Assisted Tomography*, 17(4), 1993:536-546.
- [26] S Zink and CC Jaffe. Medical image databases, *Investigative Radiology*, 28(4), 1993:366-372.

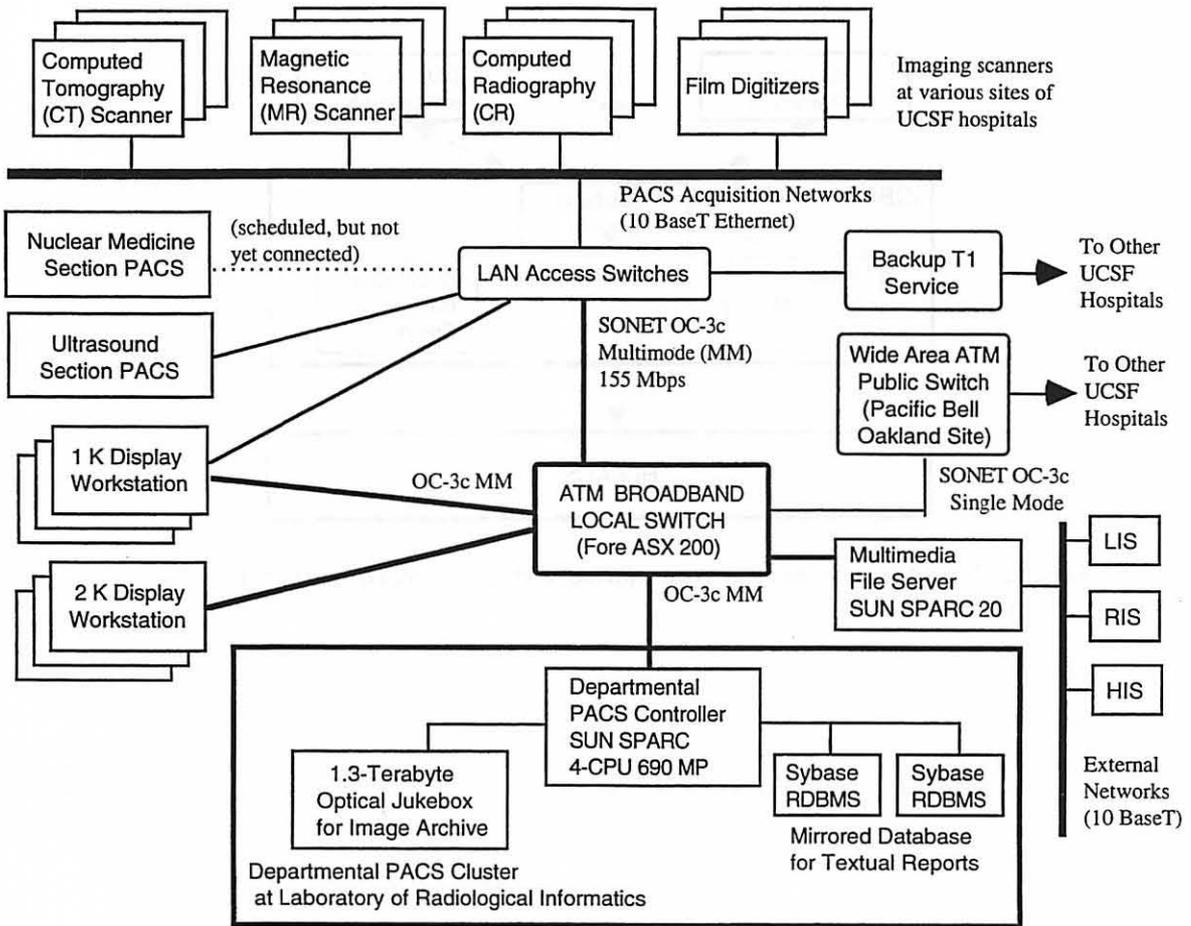


Figure 1. Schematic diagram of the hospital integrated PACS at UCSF.

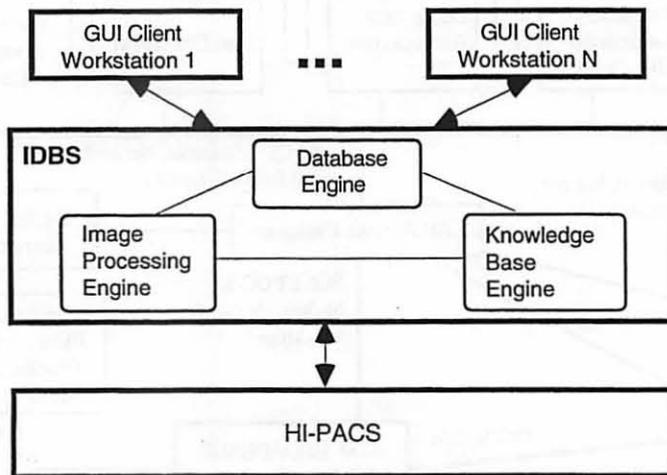


Figure 2. The medical image base system (MIDS) architecture with IDBS over HI-PACS.

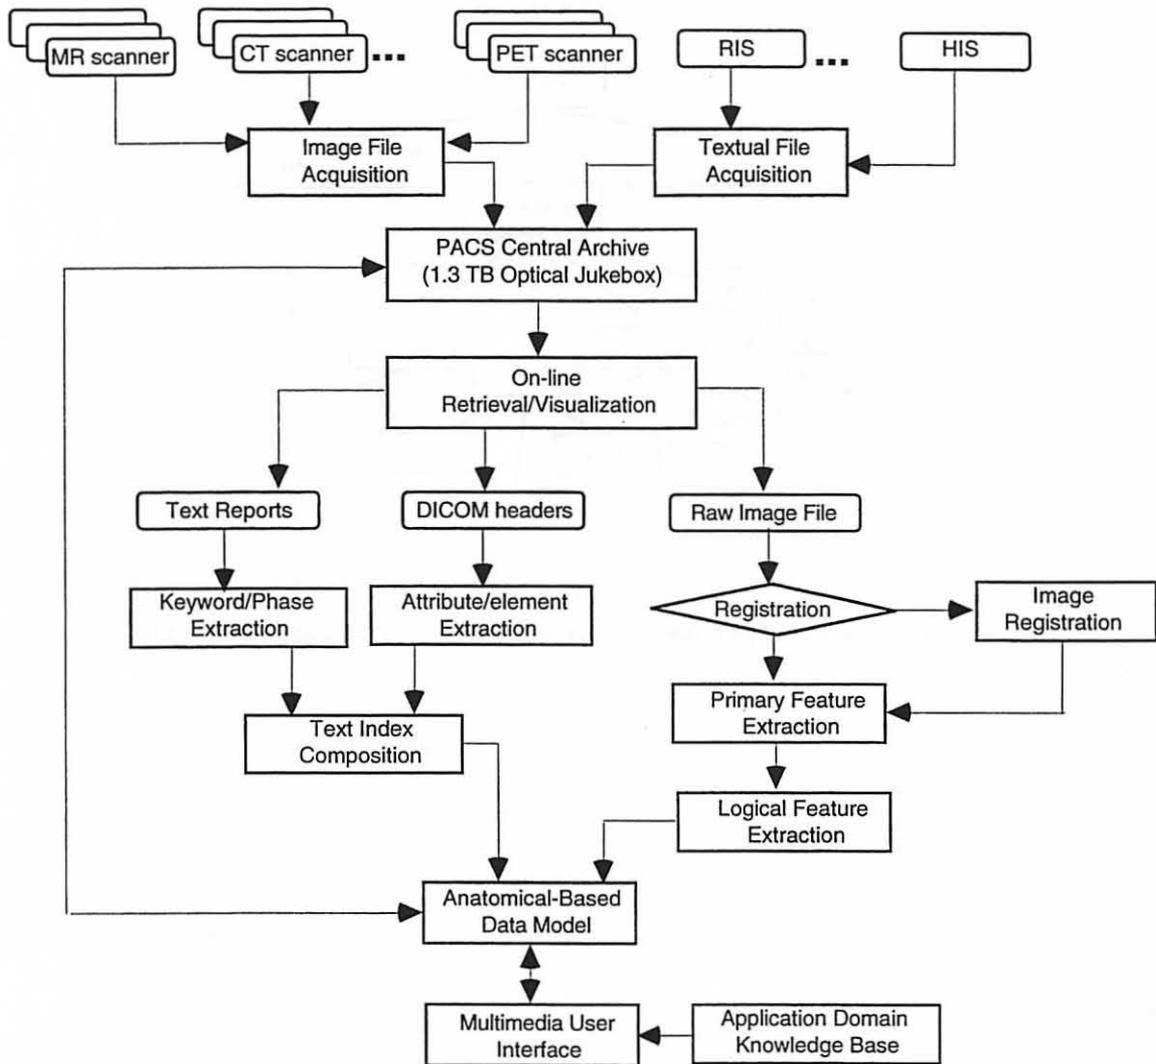


Figure 3. The operational flow of extracting image and text features into an anatomical based data model for subsequent content-based image indexing in MIDS. Specific knowledge or heuristics is triggered to aid the query and navigation of the medical image database.

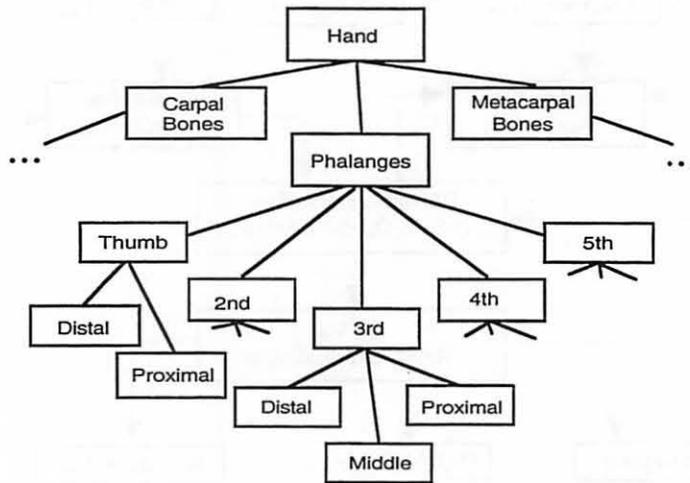


Figure 4.a The object-oriented, anatomic data model of hand bone structure.

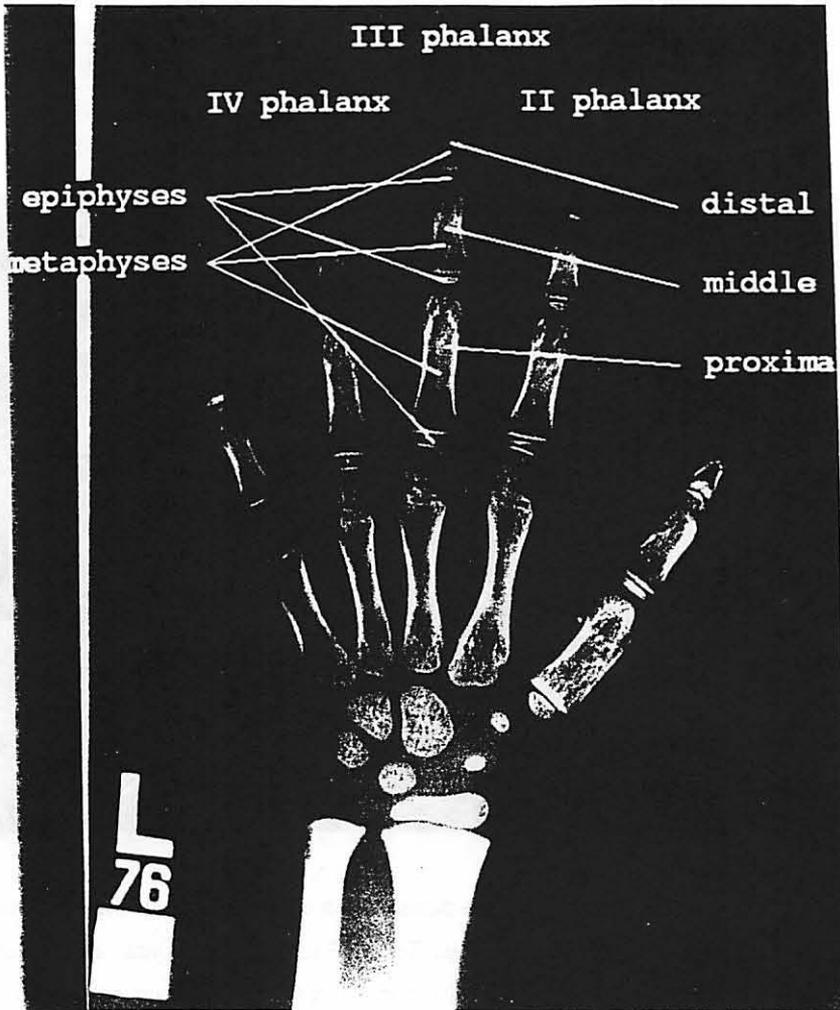


Figure 4.b. Digital representation of hand bone structure in computed radiograph (Courtesy of Dr. Ewa Pietka).

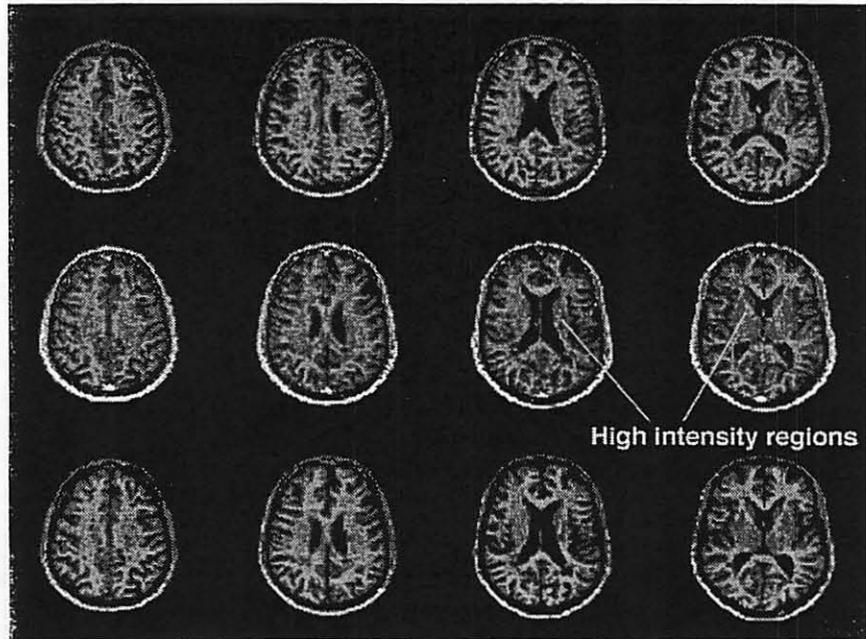


Figure 5. Automated Image registration of two MRI image datasets of a multiple sclerosis patient. Dark X-shaped regions in the center of the brain are left and right ventricles. Top: MR scans taken a year ago; middle: recent MR scans of different orientation and position; bottom: registered image of the top row with respect to the second row.

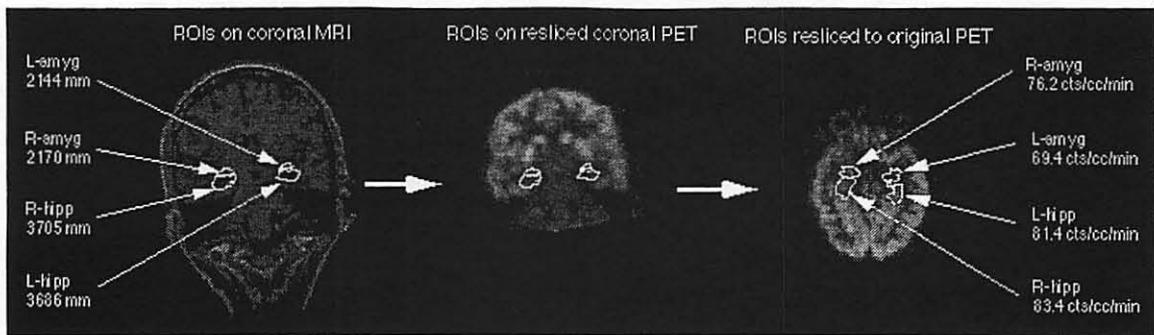


Figure 6: Mapping regions of interest from MRI to PET. For simplicity, we show only two dimensional case here. The middle one shows the aligned position of PET image with that of MRI image while the right one is the orientation of original acquired PET image. The outlined areas, right and left amygdala and hippocampus, are the region of interest of all three images in the common space.

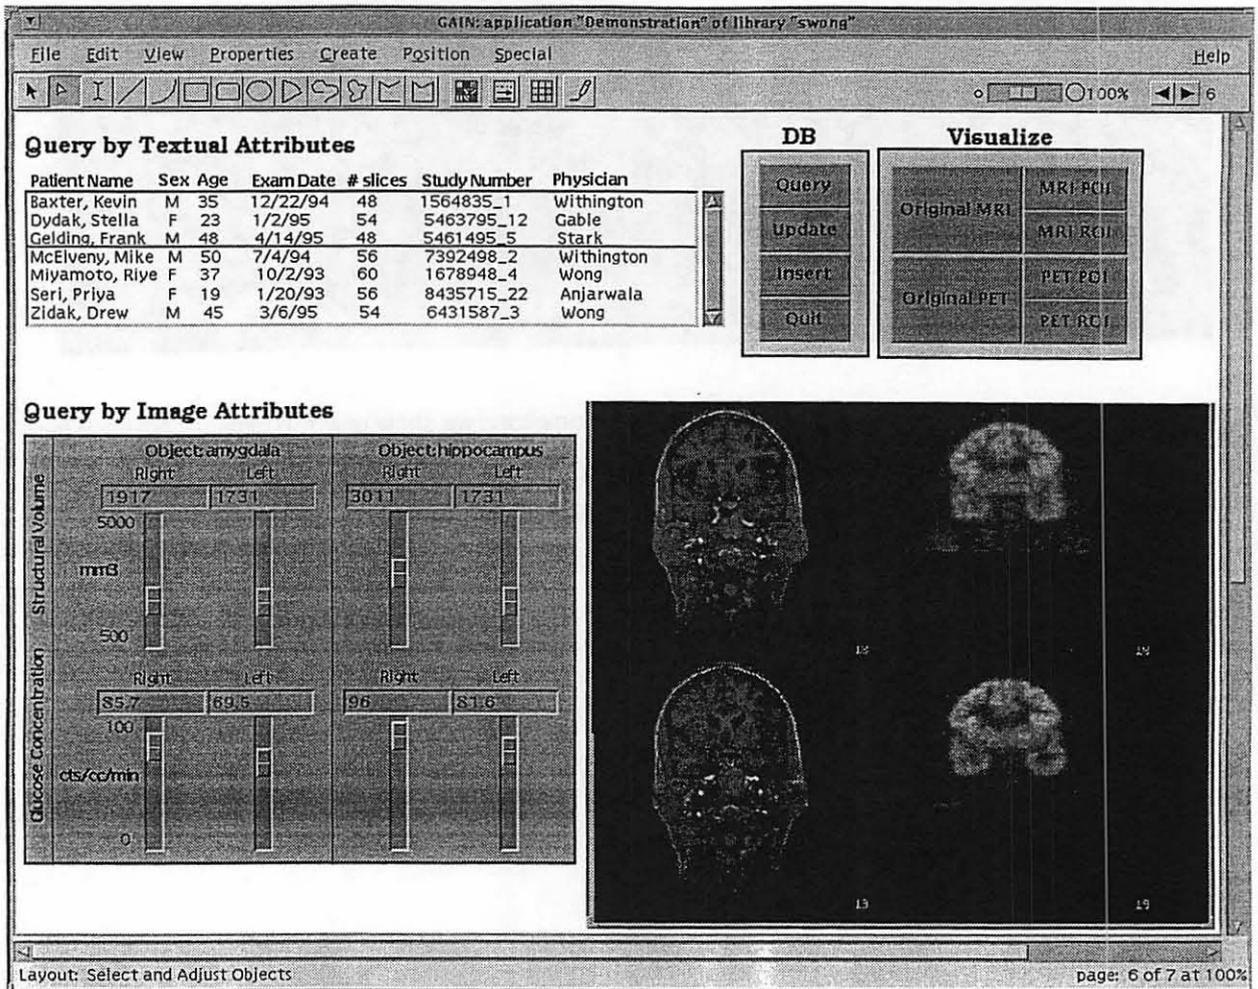


Figure 7. A content-based GUI developed for retrieving patient PET and MRI images for reference and comparison during surgical planning and diagnostic workup of a new epilepsy case. For simplicity, MRS and MEG image features of anatomical objects are not shown. The images shown in the right are the third patient in the list selected by the clinician. For confidentiality reason, all patient and physician names have been changed.

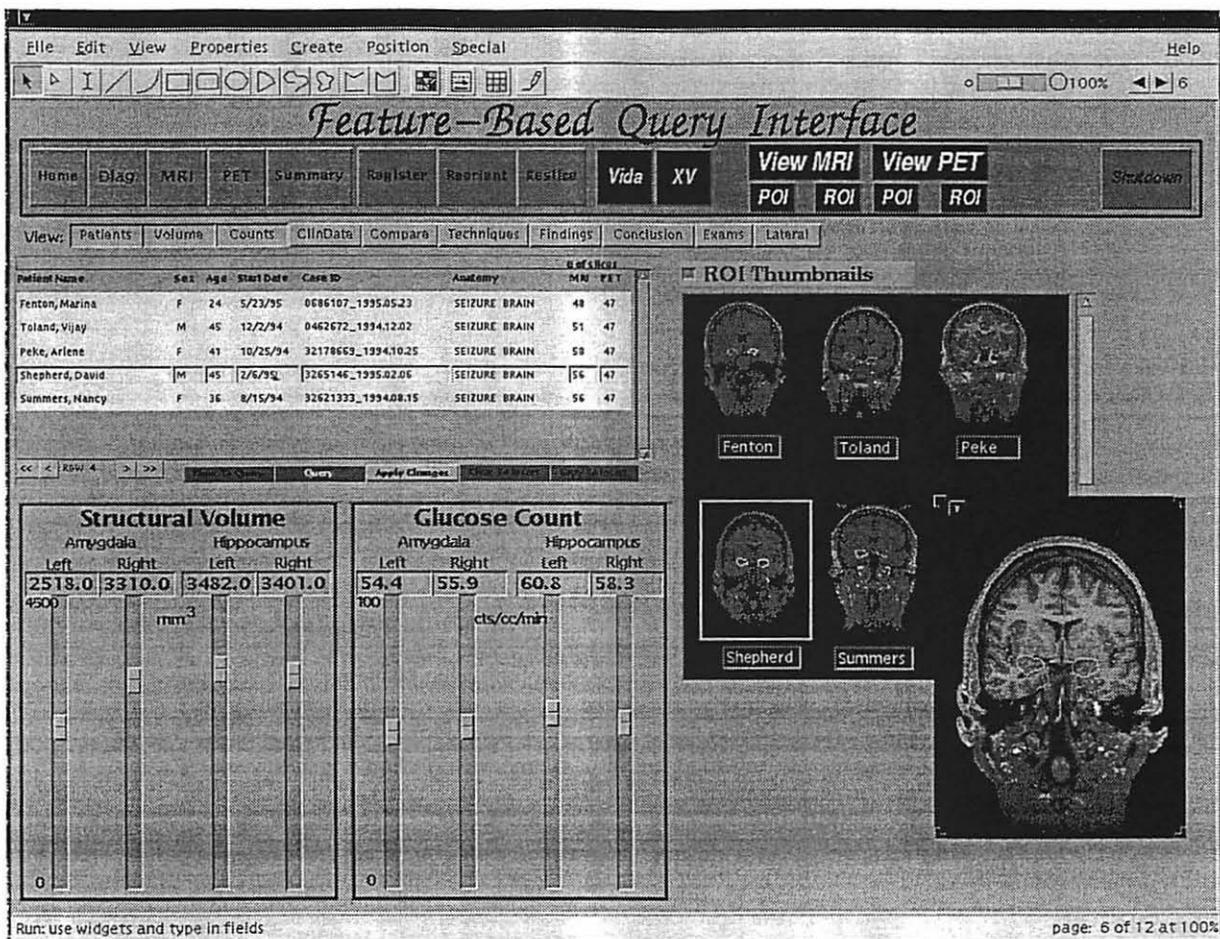
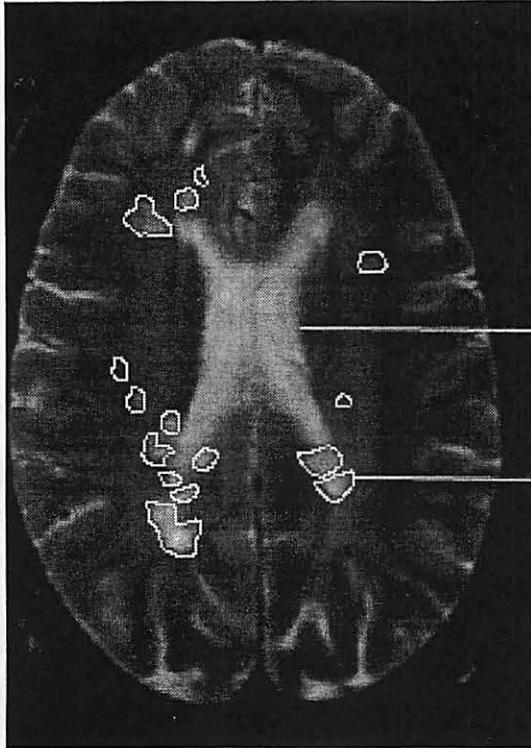


Figure 8. Content-based information retrieval based on the picture of interests. The full size MR image slice corresponds to the highlighted thumbnail image at the bottom left of ROI picture window. Both Figures 7 and 8 are programmed within the same hypermedia GUI such that the user can pick any GUI page for query according to individual preferences. For confidentiality reason, all patient and physician names have been changed.

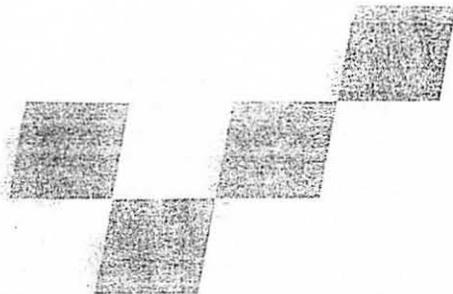


Lateral ventricles

White spots or fiber deposits

Figure 9. A cross-section slice of a T2 weighted brain MR images of a 40 year old multiple sclerosis patient. The X-shaped bright regions in the center brain are the left and right lateral ventricles. The objects of interest, i.e., white spots around the periventricular area can be segmented out interactively and their sizes can be quantified by using the image processing engine of IDBS .

# Multimodal Image Fusion for Noninvasive Epilepsy Surgery Planning



Stephen T.C. Wong, Robert C. Knowlton, Randy A. Hawkins, and Kenneth D. Laxer  
University of California at San Francisco

Epilepsy refers to a tendency toward recurrent seizures—that is, paroxysmal derangements of cerebral functions caused by uncontrolled, excessive discharges from an aggregate of neurons. Despite advances in anti-epileptic drugs, about 25 percent of epilepsy patients have seizures that cannot be controlled by drug treatment. Such cases are termed *medically refractory*. Approximately 125,000—that is, half—of medically refractory patients in the United States are candidates for surgical treatment.<sup>1</sup> Planning for surgery requires accurate identification either of a localized site of seizure onset or of a disconnection between epileptogenic zones so that palliative surgical procedures can interrupt the seizure spread.

The current “gold standard” for evaluating surgical candidates suffering from medically refractory epilepsy is intracranial electroencephalography (EEG) recording. An intracranial EEG study requires invasive surgery to place electrodes for recording brain electrical activity during spontaneous seizures. Figure 1 is an example of this procedure. The subdural grid of electrodes will remain in the patient for a week to a few weeks to monitor seizure onset. Intracranial EEGs impose definite risks (6 to 8 percent morbidity), and the costs associated with such an invasive procedure, which include EEG monitoring in the hospital, are considerable.

This invasive and expensive procedure is not performed unless the results of noninvasive studies support a reasonable guess about the possible lateralization or location of a resectable epileptogenic region. Usually, the noninvasive presurgical evaluation must identify lateralization

choices before invasive testing can be planned.

Technological advances in medical imaging have produced many new noninvasive techniques that not only help to lateralize the epileptogenic region with reasonable certainty before invasive EEG recording is performed but can also eliminate the need for such invasive testing. The technological revolution began with positron emission tomography (PET) scanning and the ability to display regional metabolic dysfunctions associated with epileptogenic zones. Magnetic resonance imaging (MRI) followed, providing unsurpassed anatomic detail to detect structural abnormalities. Research studies using PET and magnetic resonance spectroscopy (MRS) separately have demonstrated noticeable metabolic abnormalities in epileptogenic zones. Experiments with magnetoencephalography (MEG) provide real time, three-dimensional location of dipole sources from neurophysiologic epileptiform abnormalities. The sidebar briefly describes these noninvasive diagnostic techniques.

Ideally, the various preoperative localizing examinations produce concordant results that implicate a single brain region as the epileptogenic zone. Unfortunately many patients have either discordant or inconclusive imaging data. In this article, we describe an image fusion application that combines biomedical information from multiple independent neuroimaging modalities. Its purpose is twofold: to provide accurate, noninvasive localization of epileptogenic tissue for surgical resection and to resolve discordant or inconclusive imaging evaluation.

## Diagnostic workup of epilepsy

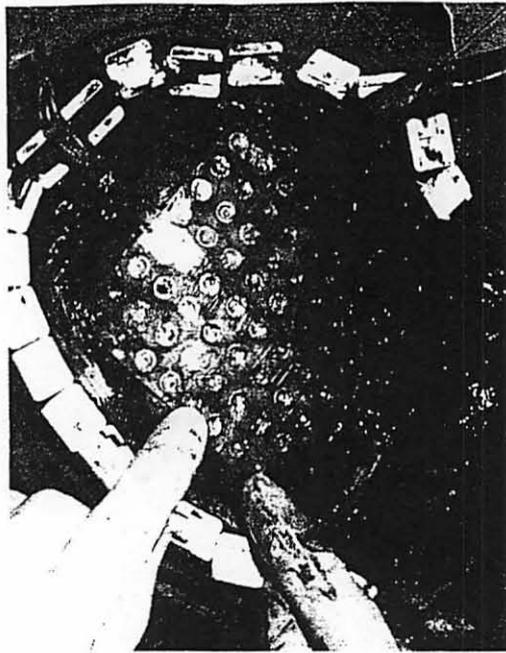
Figure 2 outlines the general diagnostic workup for presurgical epilepsy diagnosis.<sup>2</sup> A patient who enters this workup is medically refractory as determined by patient history that suggests partial epilepsy and indicates failure of drug treatment. EEG and MRI are routinely used to classify seizures and localize the epileptogenic zone. The detail in MRI visualizations

A neurodiagnostic workstation integrated into a hospital picture archiving and communication system supports low-cost, low-risk epilepsy surgery planning by combining biomedical information from four neuroimaging modalities.

brain anatomy greatly surpasses that of any other imaging modality, including x-ray computed tomography (CT), and the spatial resolution has greatly improved the detection of most mass lesions or tumors. Detection of structural lesions in epilepsy makes surgical evaluation relatively straightforward. Many epilepsy cases, however, do not exhibit structural abnormalities in MR images. Further, paroxysmal abnormalities occur intermittently and might not be captured in EEG recordings. Prolonged inpatient recording with video EEG telemetry (VET) is necessary to capture spontaneous seizures.

Positive results from MRI and VET in localizing the epileptogenic zone typically indicate that surgery can eliminate the patient's seizures. Nonconcordance between these two modalities almost invariably means that invasive EEG recordings are required to localize the epileptogenic zone. These recordings use one or more of the following: subdural strips, arrays, and stereotactically inserted depth electrodes. Approximately 35 to 40 percent of medically refractory epilepsy patients in the Northern California Comprehensive Epilepsy Center operated on at the University of California at San Francisco (UCSF) require invasive intracranial EEG monitoring. The average cost of intracranial EEG is about \$60,000 per evaluation, and the totals for this procedure are about \$2 million to \$3 million each year in the UCSF Epilepsy Center alone.

The UCSF Epilepsy Center is currently performing PET, MRS, and MEG imaging examinations on selected epilepsy patients before deciding on invasive evaluation. These additional imaging tests cost about \$3,000 to \$4,000 per patient, less than 10 percent of the cost of an intracranial EEG study. All neuroimaging studies of these patients are collected for computer analysis. The use of combined noninvasive tests to reduce the number and improve the accuracy of intracranial EEG recordings will help reduce patient risks and presurgical evaluation costs and improve surgical outcomes. In addition, multimodal imaging analysis should increase our understanding of decision thresholds and thereby increase the identification of patients with previously

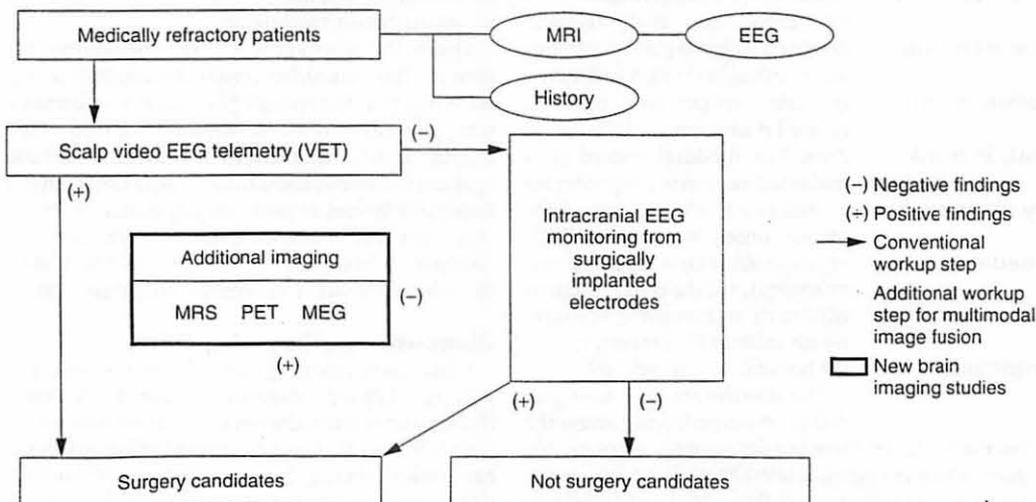


1 Intracranial EEG recording, the "gold standard" for evaluating surgical candidates for epilepsy surgery, requires its own surgical procedure—in this case, to place a subdural grid of electrodes over the frontotemporal cortex.

unlocalized seizure origination suitable for surgical treatment. We have already performed these additional imaging studies in 25 epilepsy surgery cases.

Nevertheless, an effective means to access, manipulate, and correlate large, heterogeneous multidimensional medical images was lacking. We therefore developed a computer-aided neurodiagnostic workstation system to fulfill this need. This medical workstation aims to provide an environment for optimal interpretation of multiple noninvasive image modalities by combining functional PET, MRS, and MEG information with structural MRI anatomy, both qualitatively and quantitatively. Since there is currently little knowledge relating to the positive and negative findings of various imaging and intracranial EEG studies on patient popu-

2 Diagnostic workup of epilepsy surgical candidates. Multimodal image fusion aims both to eliminate invasive intracranial EEG tests and to identify more medically refractory patients as candidates for surgical treatment.



### Diagnostic modalities for neurological disorders

Note that medical diagnostic images are mostly represented in gray-scale form, with the exception of Doppler ultrasound and color microscopy (which do not apply to neurological disorders). The pseudocolor images seen in the neurological literature are post-processed for research or presentation purposes. Clinicians and radiologists rarely use color images in daily operations. The diagnostic modalities for neurological disorders are the following:

**Electroencephalography (EEG)** is broadly classified into *extracranial*, which is noninvasive measurement of cerebral voltage potentials by attaching electrodes to the scalp, and *intracranial*, which is an invasive operation that opens the skull to put electrodes and grids of electrodes directly on the cortex surface or into the brain.

**Magnetic resonance imaging (MRI)** is the predominant way to check for structural brain-tissue abnormalities that do not show up in x-ray or computed tomography (CT). MR images are based on the signals from hydrogen nuclei contained in hydrogen-rich compounds in the body: water and lipids. Image contrast is based primarily on inherent properties of different tissues. These inherent tissue properties are hydrogen nuclei (proton) spin density and hydrogen nuclei relaxation times (T1 and T2). MRI alters image contrast by selectively altering the hydrogen nuclei times of tissues or the pulse-echo sequence of radio frequency.

**Magnetic resonance spectroscopy (MRS)** measures in-vivo chemicals within the body. It exploits the principle that every chemically distinct nucleus in a compound resonates at a slightly different frequency, allowing the detection of a wide variety of proton, phosphate, and carbon signals in organic molecules. The same instrument can be used for both MRI and MRS. The major difference important to clinicians is that MRI produces a high-resolution visual image whereas MRS obtains chemical information, typically expressed as numerical values.

**Magnetoencephalography (MEG)** involves three-dimensional dipole localization methods using multichannel ultrasensitive antenna to

detect magnetic fields from the brain. By analyzing weak magnetic signals, MEG can help noninvasively pinpoint the source and time sequence of electrical activities in the brain. By transforming the coordinates from the biomagnetic measurement system to MRI, it is easy to project all or parts of the reconstructed dipole pathway into a multislice or 3D MR image of the patient. Source locations are then overlaid onto MR images to create a composite picture that aids in identifying the relationship between function and anatomy.

**Positron emission tomography (PET)** measures in-vivo substances labeled with positron-emitting isotopes or radiopharmaceuticals and then constructs tomographic images reflecting tissue biochemistry and physiology from these measurements. The most common clinical PET procedure uses the radiopharmaceutical (F-18) fluorodeoxyglucose ( $^{18}\text{F}$ FDG) and results in images representing tissue glucose utilization. Another nuclear medicine modality, single photon emission computed tomography (SPECT), can be considered a lower resolution but less expensive version of PET.

**Video electroencephalography telemetry (VET)** is a form of extracranial EEG that measures cerebral voltage potentials on epilepsy patients while the seizures occur. This procedure normally requires in-patient hospital stay, recorded by a video monitor and an on-line computerized spike controlled by seizure detection programs. VET is sometimes referred to as ictal EEG recording.

### Further reading

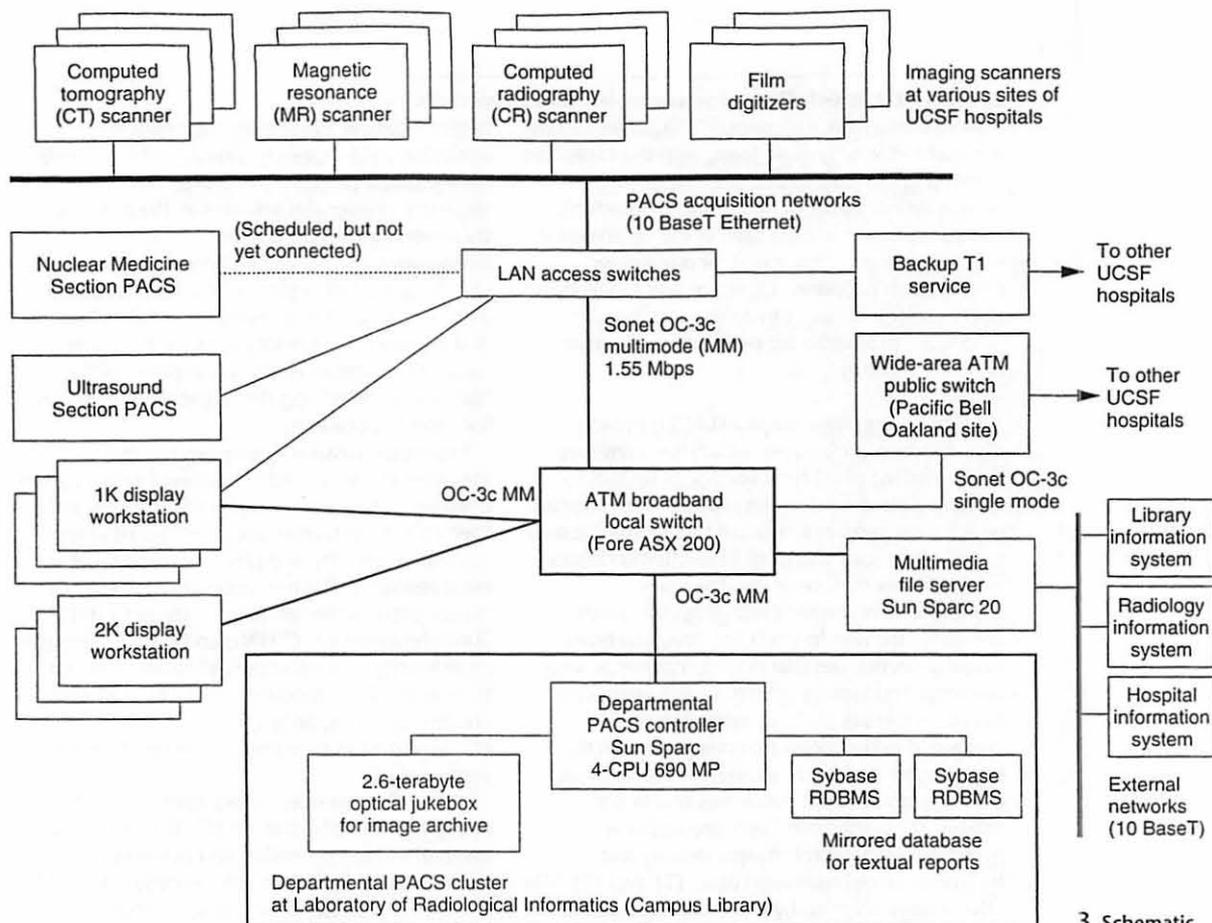
- T.S. Curry III, J.E. Dowdey, and R.C. Murry, *Christensen's Physics of Diagnostic Radiology*, 4th ed., Lea & Febiger, Philadelphia, 1990.
- J.O. Greenberg, ed., *Neuroimaging: A Companion to Adams and Victor's Principles of Neurology*, McGraw-Hill, New York, 1995.
- K.D. Laxer et al., "Experimental Technologies" in *Surgical Treatment of the Epilepsies*, J. Engel, ed., Raven Press, New York, 1993, pp. 291-308.

lations, the data extracted will also help in outcome analyses to identify inefficient or duplicative evaluation procedures.

### Infrastructure for multimedia data management

The current manual recording system in hospitals stores diagnostic images on film and EEGs and other textual records on paper. This approach poses several obstacles to data management. First, it contributes to

data communication problems, such as accuracy, coordination, accessibility, and timeliness. Second, it depends on the human visual system to mentally estimate corresponding points of interest in sets of tomographic images—a difficult process when the images are out of alignment with one another. Third, it is tedious, labor-intensive, and sometimes impossible to retrieve patient history and similar clinical cases manually from disparate information sources for comparison study. Fourth, neurodiagnosticians can miss subtle



3 Schematic diagram of hospital-integrated PACS at UCSF.

abnormalities or asymmetries on film-based systems; worse, there is no record of how frequently this happens. Finally, without co-registering PET to MRI, interpretation of functional PET scans is limited by the inability to normalize metabolic changes to structural volume and by the presence of volume-averaging artifacts. The interpretation of MRS biochemical spectra is similarly limited.

Research in multimodal image fusion must first overcome the administrative and technological barriers to gathering large-volume, heterogeneous, and fragmented medical images and patient records from distributed imaging and information sources. Thus, we built our computer-aided neurodiagnostic workstation system on top of the existing picture archiving and communication system (PACS) infrastructure at UCSF.

In the past decade, PACS infrastructures have been the prevalent means of acquiring, storing, communicating, and displaying digital images related to radiologic practice.<sup>3</sup> A PACS integrates image acquisition devices, computers, communication networks, image display workstations, and database management systems. Its operations can be summarized as follows:

1. Digital images generated in an examination are captured by an acquisition computer and sent immedi-

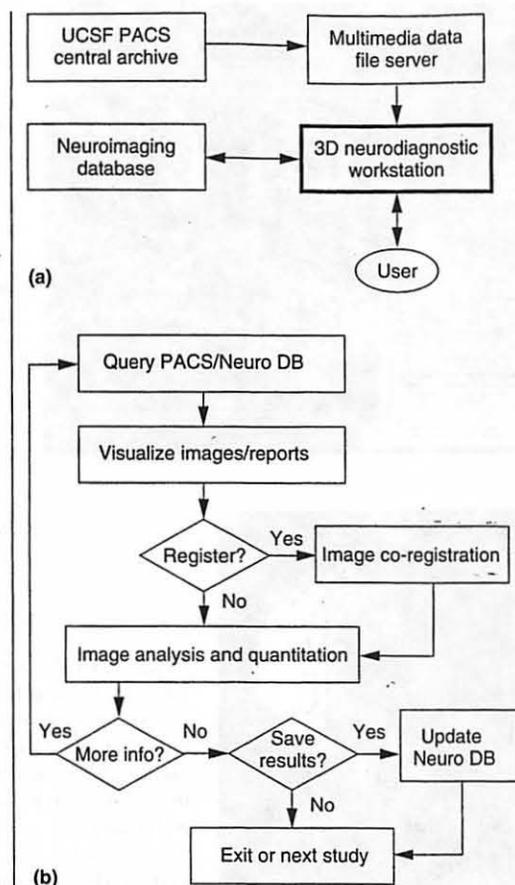
ately to a central database for long-term archive.

2. These images are then retrieved from the central archive and transmitted for display at a remote medical display station, either automatically or upon the user's request.
3. Diagnostic findings from viewing these images are appended to the images at the central archive.

The complexity of a PACS is application-dependent. A PACS for an intensive care unit can consist of a video camera to digitize radiographs, a broadband video system to transmit medical images, and a video monitor to receive and display images. On the other hand, a department-wide PACS is comprehensive and requires careful planning and a large capital investment. During the past decade, several such large-scale systems have been developed for clinical trial and use.<sup>3</sup>

Figure 3 is a schematic diagram of the hospital-integrated PACS at UCSF.<sup>4</sup> Its implementation emphasizes recognized standards, open systems connectivity, hierarchy memory management, database integration, and security. It connects not only directly to various imaging scanners, but also to PACS of subspecialty sections and to other textual information systems, such as the hospital, radiology, and library information systems. The PACS fiber-optic backbone connects four major facil-

4 (a) Distributed system architecture of the neurodiagnostic workstation, and (b) computer-aided workup for epilepsy surgical planning.



ities in UCSF—Moffitt and Long hospitals, the ambulatory care clinic, and the campus library—and a data bank that links various clinical databases, imaging devices, and medical reference sources.

Wide-area connections between UCSF and affiliated hospitals in the San Francisco Bay area are being implemented using asynchronous transmission mode, optical carrier level 3 (ATM OC-3c), 155 megabits per second for real-time telemedicine). Image communications between UCSF hospitals are in operation using frame relay technology (T1 link, 1.54 Mbps) for non-real-time telemanagement. The centralized archive at the departmental PACS cluster can store up to 2.6 terabytes of image and text files, which can be retrieved and retransmitted over the networks to high-resolution, multiterminal display stations for diagnostic review.

PACSs are intended for management and fast review of vast volumes of image and text files in a digital radiology department. The file management system can only handle query by artificial keys, such as patient name or hospital ID, and it lacks the means to organize, synthesize, and present complex, heterogeneous, multimedia medical data to clinicians in a structured way for sophisticated computer-aided diagnosis, therapy, and surgery.<sup>5</sup> To provide genuine database management capabilities requires an image database server that can extract relevant features and attributes from

images and textual reports stored in a PACS archive and organize the extracted data, together with selected image segments, into a separate image database. Such an image database system supports more flexible indexing of PACS medical images and text and thus querying by content or features. At UCSF, we have developed a multimodal neuroimaging database for epilepsy surgery planning.<sup>6</sup>

### Computer-assisted surgery planning

Figure 4 shows the distributed system architecture of the computer-aided neurodiagnostic workstation together with the computer-aided workup steps that the workstation supports. The medical workstation provides a multimedia user interface to (1) access medical data stored in the PACS file archive and the neuroimaging database and (2) analyze the retrieved data for noninvasive planning of epilepsy surgery.<sup>6</sup> For security reasons, all user queries from nondaily clinical operations pass through a multimedia data file server, which acts as a "firewall" against unauthorized access to the PACS central archive (see Figure 4a). All computers are Sun Sparc workstations running SunOS except the neuroimaging database server, which resides in a four-processor Sun Sparc 690 MP running the Solaris operating system. The workstation's multimedia user interface is built on top of the Gain Momentum object-oriented GUI builder (Gain Technology Division, Sybase, Palo Alto, Calif.) and the X Windows programming environment. Currently, not all of the four brain imaging modalities involved in epilepsy surgical planning have been archived into UCSF PACS (see Figure 3). Nuclear medicine PET image data sets are manually archived into the neuroimaging database using the File Transfer Protocol (FTP). The neurodiagnostic workstation includes the routing table on which to retrieve these PET image files upon user query.

### Information retrieval and visualization

As noted above, information retrieval from PACS is based on artificial keys, such as patient name or hospital ID, but the neuroimaging database supports user queries by patient history such as age and sex, by image features such as volumes of certain brain regions and sizes of tumors, or by a combination of text and image features.

Figure 5 illustrates the retrieval of patient studies, diagnostic reports, and medical images from PACS for on-line epilepsy surgical planning (query to the neuroimaging database is illustrated later in Figure 8). The user must enter a password before accessing the PACS, and information retrieved is filtered by the image server before it goes to the neurodiagnostic workstation for display (see Figure 4a). The upper portion of the left scroll window lists the imaging studies done on a patient (the patient name has been changed for confidentiality), while the lower scroll window displays the diagnostic report of one of the studies. Retrieval is accomplished by selecting the desired study from the upper window. After reviewing the diagnostic report, the user can retrieve the specific image file for visualization and manipulation (see the right image-display window).

## Image registration and combination

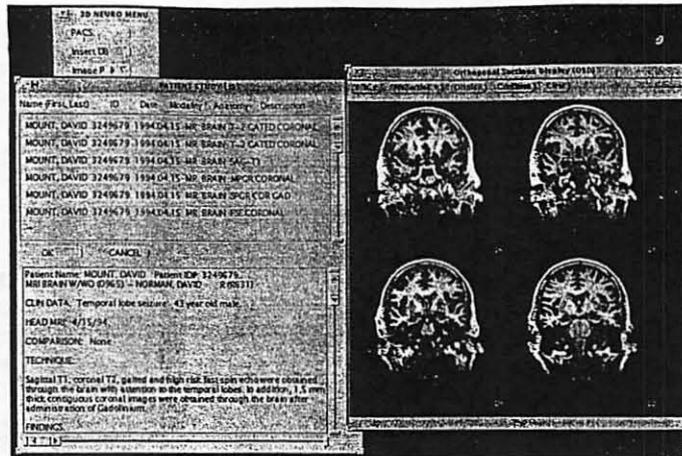
Before medical data retrieved from PACS or the neuroimaging database can be combined and rendered into 2D and 3D images for analysis, the medical images must often be registered so that voxels representing the same underlying anatomical structure can be superimposed or directly compared. A pair of image volumes might be of the same subject and the same modality but different positions or different scanning instruments. They might also be of the same subject but different modalities, or the same modality but different subjects.

The registration strategy for the multimodal image fusion application is to use the high-resolution, structural MRI data set as the reference point for mapping across different imaging modalities. We perform the MRI-MEG registration using a three-point fiducial marking technique to project all or parts of the reconstructed dipole pathway into a multislice or 3D MR image of the patient.<sup>7</sup> The MRI-MRS registration is straightforward: MRS scans are performed immediately after the MRI examination of the patient using the same MR scanner and no change of the patient's head position. For the MRI-MRI, PET-PET, and MRI-PET image pairs, we implemented the ratio-image uniformity (RIU) technique described by Woods.<sup>8</sup> First, we determine the translation and rotation parameters such that they minimize the voxel-by-voxel variance of the "ratio" image. Prior to calculating the ratio image for MRI-PET registration, a semi-automatic process strips the MRI volumes to remove the skull and scalp. RIU algorithms then iteratively minimize the ratio-image variance using a modified Newton-Raphson method based on explicitly computed derivatives of the variance with respect to each registration parameter.<sup>8</sup> Recent study confirms that optimizing image registration on the basis of measured similarities in the spatial distributions of voxel values is superior to techniques that do not use such information.<sup>9</sup>

Figure 6 illustrates an example of volumetric visualization of MRI-PET registration. The top image of the figure presents a patient case of left mesial temporal lobe epilepsy with the registered PET slice superimposed on the MRI volume, while the corresponding 2D MRI slice at the bottom shows very subtle hippocampal atrophy on the left temporal lobe. The initial image interpretation missed this subtle change of anatomy, but the abnormal metabolism detected in the PET slice within the same anatomic region helped to identify the atrophy.

## Extraction of structural and functional information

After image pairs are registered into common space, the multimodal image fusion application can combine, extract, and quantitate biomedical information to identify the location and nature of the epileptogenic zone. For example, the high-resolution MR image provides a



5 Retrieving multimedia medical data (patient studies, diagnostic reports, and MR image slices) for on-line epilepsy surgical planning.



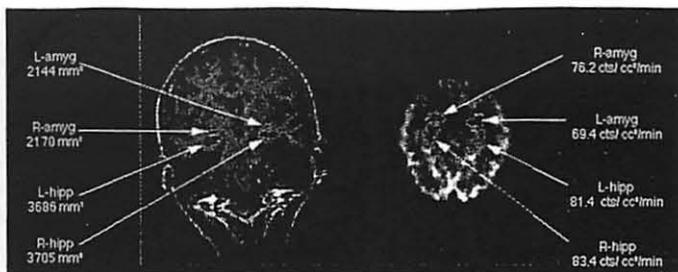
6 Top: <sup>18</sup>F<sub>2</sub>DG-PET viewed on the cut surface of registered 3D MRI volume for a patient with left mesial temporal lobe epilepsy. Abnormal metabolism appears at the mesial temporal lobe region on the superimposed PET image slice (arrow 1).

Bottom: The corresponding MRI coronal slice shows subtle structural atrophy of hippocampus within the left mesial temporal lobe (arrow 2), which was missed originally but later identified with the aid of the 3D image.

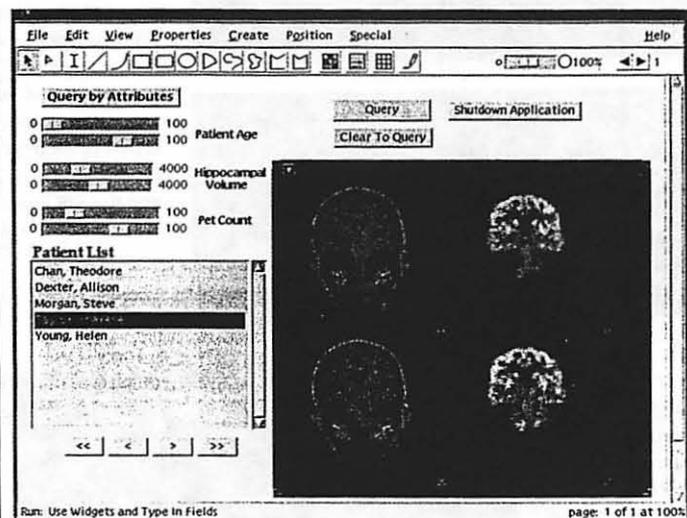
basis to anatomically map glucose metabolism. This supports the identification of metabolic rates in specific brain regions of interest.

The neurodiagnostic workstation incorporates a program developed at Lawrence Berkeley Laboratory's Center of Functional Imaging in conjunction with the Volumetric Image Display and Analysis (VIDA) project at the University of Iowa Medical College. The program maps regions of interest defined on the MRI to the PET sinograms for quantitation of glucose uptake. Figure 7 (next page) illustrates an example taken from another partial epilepsy case. It shows the original MRI and PET images on the left and right, respectively. The values called out on the left were obtained from structural/functional image analysis on volumes in the hippocampus and amygdala. The values on the right represent the <sup>18</sup>F<sub>2</sub>DG glucose uptake per unit volume. Such quantita-

7 2D original slices from MRI (left) and PET (right). Quantitation of hippocampus and amygdala regions of the MRI can be used with the co-registered MRI-PET images to obtain fluorodeoxyglucose concentration per unit volume in  $^{18}\text{F}$ FDG-PET.



8 Extracted image features, such as those in Figure 7, can be used for indexing the neuroimaging database to support comparison and reference during surgical planning and diagnostic workup of a new case.



tive data, which are not obtained in conventional workups, can define precise localization information.

#### Image query by contents

The clinician can save the analyzed results of the workup into the neuroimaging database for future reference. For example, structural and functional information extracted as shown in Figure 7 can be stored in the neuroimaging database to enable flexible user query of multimedia medical data. Figure 8 illustrates a graphical user interface for accessing information contained in the neuroimaging database. Query is by image content—in this case, MRI hippocampal volume or PET glucose concentration—or by patient attributes such as age. In Figure 8, three pairs of sliders specify the range of these three patient attributes. The image database returns a list of patients satisfying the combined attribute ranges. The user can then select a name from the patient list and retrieve MRI and PET images for display (only two slices are shown here, and the PET image is registered to the MRI image).

The multimedia user interface is object-oriented and compositional, so that users can quickly implement another interface format, tailored to a different application but using the same set of graphical icons and tools. In addition, as the image database grows, the analyzed data will help in performing outcome analyses and in evaluating the cost-effectiveness of various imaging and surgical protocols.

#### Clinical case example

We now describe the clinical use of image fusion techniques in a complex case of nonlesional temporal lobe epilepsy. (Note that MRI scans readily show the presence of gross lesions. Thus, there is no need to use additional functional information to localize the region of epileptogenic focus in these cases.)

The patient is a 29-year old right-handed female graduate student in performing arts. She was referred to the Northern California Comprehensive Epilepsy Center for evaluation of medically refractory, complex partial seizures. She was an otherwise healthy, young, and intelligent woman whose seizures began when she was five years old. She had been evaluated thoroughly but no

etiology was identified. The seizures were not clinically significant until her late teen years, especially after she began college. She was sequentially tried on three anti-epileptic medications during this period, none of which succeeded in controlling her seizures and all of which had adverse cognitive side effects.

For a few years prior to evaluation for surgery, her seizures became extremely problematic. They were characterized predominately by variable experiential phenomena and complex hallucinations associated with impairment of consciousness. The seizures she called "fire" dreams arose out of sleep and caused her to run from bed for safety, often resulting in falls and associated injuries before she would become aware of her surroundings and situation. Her daytime seizures occurred randomly and caused severe alteration in consciousness, exposing her to potentially lethal risks, for example, walking directly into a busy street full of traffic.

Several routine EEGs had been performed in the past. Most were normal, but at least one during sleep revealed an active disturbance of electrocerebral activity in the right temporal lobe. Inpatient monitoring with VET was nonlocalizing but suggested lateralization to the right hemisphere. Clinical semiology suggested the possible involvement of both frontal and temporal lobes. Repeat high-resolution imaging with MRI and PET were interpreted as normal (that is, of no help in localizing the epileptogenic region).

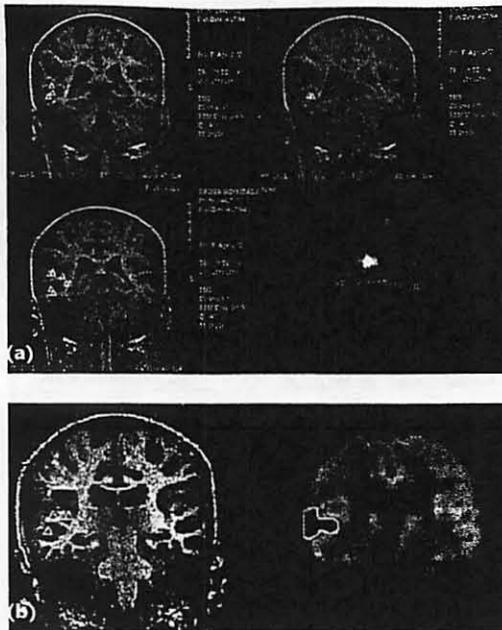
The patient underwent invasive seizure monitoring

with implanted subdural electrodes covering both the frontal and temporal lobes bilaterally. Several seizures were captured, but again no definite localization could be obtained. Seizures appeared to arise from the lateral temporal region on the right on some occasions; on others, they appeared to involve right mesial temporal, left lateral temporal, and both frontal brain regions simultaneously. These confusing findings indicated either poor electrode sampling and variable spread patterns or multifocal seizure onsets.<sup>2</sup> In essence, the invasive intracranial EEG evaluation failed.

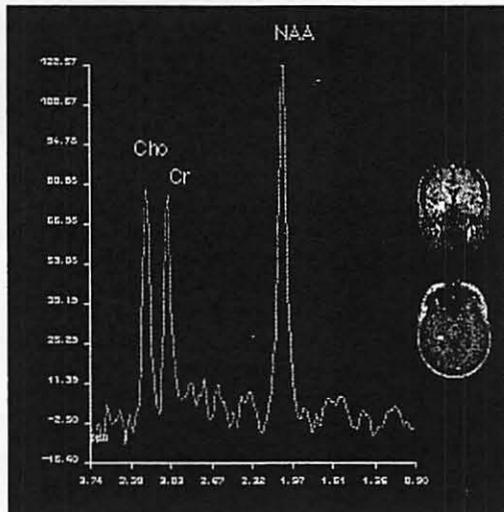
At this point, the patient was not considered to be a candidate for surgery. The only option for her was to repeat the monitoring with more intracranial electrodes, thus incurring non-negligible risks and considerable further expense. Instead she underwent further evaluation of her abnormal EEG with simultaneous recording and analysis by MEG. This study showed a strongly localized cluster of abnormal cerebral activity in the superior temporal gyrus of the right temporal lobe (see Figure 9a). Co-registration of the FDG-PET data revealed an unequivocal defect in glucose uptake in the same region of the right superior temporal gyrus (see Figure 9b).

An MRS study was also performed, limited to evaluation of the left and right mesial temporal lobe structures. Figure 10 illustrates the spectra of proton-based metabolites in one of the voxels used to sample the hippocampus of the right temporal lobe. The peaks measured with <sup>1</sup>H MRS in this case represent choline (Cho), creatine (Cr), and N-acetyl-aspartate (NAA). Recent published work suggests that metabolite ratio NAA/(Cr + Cho) is so far the most effective rule for lateralization in temporal lobe epilepsy.<sup>10</sup> Calculations performed on the workstation in this case indicated that the ratios of both sides were similar (nonlateralization) and only slightly below normal compared to the mean concentration ratios of the same defined regions of 20 normal control subjects stored in the neuroimaging database. Although the <sup>1</sup>H MRS provided no lateralization information, it was felt to represent a pertinent negative.

Based on the strong concordance of PET, MEG, and EEG results in one particular neocortical region within the temporal lobe area, a consensus panel of medical experts decided to go ahead and offer surgery to remove this focal region with the available information regarding localization, on the grounds that a focal resection in this region of the brain offered the least risk of neurologic injury to the patient. During surgery, the pathology of the normal appearing tissue by MRI revealed abnormally developed cortical tissue, confirming that epileptogenic tissue was present. The patient's surgery was completed successfully, and she is now free of seizures and suffers no neurological deficits.



9 (a) MRI-MEG superimposition with a strongly localized MEG dipole cluster at the right superior temporal region shown in the Cartesian coordinates of the brain. (b) MRI-PET-MEG co-registration and superimposition showing an unequivocal focal defect in glucose metabolism on PET in the same right superior temporal gyrus region (outlined area in the PET image) marked by abnormal MEG spikes (represented by triangles superimposed on MRI anatomy). For simplicity, 2D slices are shown; also the MR image is slightly posterior to the ones in (a).



10 MRI-MRS coregistration and spectral analysis. <sup>1</sup>H MRS spectra for the defined voxel region (white squares in MR images) in the temporal lobe structures as seen on the MRI.

## Conclusion

Planning for epilepsy surgery requires analyzing and combining structural and functional data from numerous sources to localize the seizure foci for resection. The multimodal image fusion work reported here is the first attempt to apply advanced computing technology to combine multiple neuroimaging modalities such as MRI, PET, MRS, and MEG for on-line, noninvasive surgery planning.

Our primary purpose is to reduce the frequency of, and ultimately to eliminate, invasive intracranial EEG tests and thereby to

- reduce costs from the \$40,000 to \$80,000 range per intracranial EEG study to the \$3,000 to \$4,000 range for multimodal image fusion study;

- eliminate the 6 to 8 percent patient morbidity risk associated with intracranial study to the risk-free imaging process; and
- increase the medical center's capacity to serve more patients by eliminating the prolonged hospital stay required for monitoring intracranial EEG signals.

In addition, as the clinical case described here shows, computer-aided image fusion can help push the surgical envelope by identifying more patients with intractable epilepsy for surgical relief.

Medical decision-making is an information-intensive endeavor, especially in today's medical environment involving large, heterogeneous, multimodal diagnostic images. Thus, infrastructural support is important. Building the computer-aided neurodiagnosis workstation on the hospital-integrated PACS and neuroimaging database at our institution makes the extracted data available for subsequent outcome analysis and performance evaluation.

In addition, once the computer architecture is in place, new image fusion techniques can be readily incorporated into the workstation for surgical planning of other neurological diseases. ■

#### References

1. "NIH Consensus Conference—Surgery for Epilepsy," *JAMA*, Vol. 264, No. 6, 1990, pp. 729-733.
2. J. Engel, ed., *Surgical Treatment of the Epilepsies*, Appendix II: Presurgical Evaluation Protocols, Raven Press, New York, 1993.
3. H.K. Huang et al., eds., *Picture Archiving and Communication System (PACS)*. *J. Digital Imaging*, Vol. 5, no. 1, 1992: 22-25.
4. S.T.C. Wong, H.U. Lemke, and H.K. Huang, "Applicability of ATM to Distributed PACS Environment," *SPIE Medical Imaging Conference: PACS Design and Evaluation*, Feb. 26-Mar. 2, San Diego, 1995. 430-439
5. S. Zink and C.C. Jaffe, "Medical Image Databases," *Investigative Radiology*, Vol. 28, No. 4, 1993, pp. 366-372.
6. S.T.C. Wong, et al., "Integrating Multidimensional Imaging, Multimodality Registration, and Multimedia Database for Epilepsy Diagnosis," *SPIE Medical Imaging Conference: Image Display*, Feb. 26-Mar. 2, San Diego, 1995.
7. K.D. Laxer, et al., "Experimental Technologies," in *Surgical Treatment of the Epilepsies*, J. Engel (ed.), Raven Press, 1993, pp. 291-308.
8. R.P. Woods, J.C. Mazziotta, and S.R. Cherry, "MRI-PET Registration with Automated Algorithm," *J Computer Assisted Tomography*, Vol. 17, No. 4, 1993, pp. 536-546.
9. S.C. Strother, et al., "Quantitative Comparisons of Image Registration Techniques Based on High-Resolution MRI of the Brain," *J Computer Assisted Tomography*, Vol. 18, No. 6, 1994, pp. 954-962.
10. A. Connelly et al., "Magnetic Resonance Spectroscopy and Temporal Lobe Epilepsy," *Neurology*, Vol. 44, 1994, pp. 1,411-1,417.

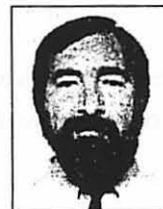


**Stephen T.C. Wong** is an assistant professor of radiology and bioengineering at University of California at San Francisco. He is also a licensed professional engineer in electrical engineering with several years of industrial experience with AT&T Bell Labs, Hewlett Packard, and the Japanese Fifth Generation Computer System Project. His research interests include biomedical informatics, imaging, and networking. Wong received the PhD degree in 1991 from EECS Department of Lehigh University, Bethlehem, Pennsylvania.



**Robert C. Knowlton** is an assistant professor of neurology and radiology at University of California at San Francisco. His research interests include clinical applications of epilepsy imaging and magnetic source imaging techniques. Knowlton is board certified in neurology. He obtained his MD from Louisiana State University School of Medicine, New Orleans in 1985 and completed his residency and fellowship at University of California at Los Angeles and UCSF.

**Randy A. Hawkins** is a professor of radiology, chief of the Nuclear Medicine Program and vice chair in the Department of Radiology at University of California at San Francisco. His research interests include image processing and modeling methods, positron emission tomography, and other imaging methods. Hawkins is board certified in diagnostic radiology and nuclear medicine. He received his MD from University of California at Irvine in 1975 and PhD in biomedical physics from University of California at Los Angeles in 1985.



**Kenneth D. Laxer** is an associate professor of clinical neurology and director of Northern California Comprehensive Epilepsy Center at University of California at San Francisco. Previously, he was director of the Department of Clinical Neurophysiology and the Oregon Comprehensive Epilepsy Program at Good Samaritan Hospital and Medical Center, Oregon. Laxer is board certified in adult neurology, electroencephalography, and clinical neurophysiology. He received his MD from University of California at Irvine in 1973.

Readers may contact Wong at Dept. of Radiology, Univ. of California at San Francisco, 530 Parnassus Avenue, Room 158 CL, San Francisco, CA 941430840, e-mail [steven\\_wong@rad-mac1.ucsf.edu](mailto:steven_wong@rad-mac1.ucsf.edu).

## **Coping With Conflict in Cooperative Knowledge Based Systems**

**Stephen T. C. Wong, *Member, IEEE***  
**The University of California, San Francisco**

### **Abstract**

Cooperative Problem Solving (CPS) is an important paradigm that will extend the power of current information systems to provide broader, cheaper services and to solve larger, and increasingly, more complex problems. In this paper, we address a critical issue of this new mode of computing: the existence of conflict among distributed agents. In particular, we focus our study on Cooperative Knowledge-Based Systems (CKBS). To obtain a better understanding and more balanced judgement of multiagent conflict, we provide a general scheme to study the logical structure of multiagent conflict and rational strategies of coping with it under different situations. Our research finding is that there is no grand unified theory of coping with conflict in performing complex real-world computer supported tasks. Instead, a library of alternative methods should be considered. We discuss four methods: inquiry, arbitration, persuasion, and accommodation. These methods can be combined in an order appropriate to the application domain such that if one method fails, the system will try the next. We point out merits and shortcomings of these methods and illustrate them using several high-level protocols and application examples from a prototype system, Building Design Network (BDN).

**Keywords and Phrases:** Cooperative problem solving, multiagent conflict resolution, knowledge-based systems, collaborative design, communication protocols.

## 1 Introduction

Cooperative problem solving is emerging as a prominent paradigm for the next generation of information processing systems. Broadly speaking, a cooperative problem solving (CPS) system is a set of several loosely connected and often heterogeneous agents that work together to solve complex problems beyond the capability of the individuals [14,20]. This paper concerns one important class of CPS systems: **cooperative knowledge-based systems** (CKBS), a merging field of distributed artificial intelligence (DAI) and distributed databases (DDB), and its applications for a class of important problems: collaborative preliminary design. In this context, a nodal knowledge-based system is an autonomous agent, i.e., it asserts control over its own programs, databases, and resources. Owing to its potential for wide-ranging applications, CKBS is a rapidly developing field involving researchers from many areas of artificial intelligence as well as from database, office automation, and distributed computing areas [13, 33, 11, 7, 5].

One of the central issues arising in CKBS is the existence of conflict among agents. An agent conflicts with another when both assert propositions that they know cannot both be true at the same time. The propositions can represent not only data and rules, but also resources, preferences, activities, or any information which can be encoded in symbolic form. Such a conflict causes a standstill, that is, it inhibits further action during computation. A task has to be done at once, but the agents stand in each other's way. This idea of conflict needs not be negative, however. Conflict is certainly bad in that it paralyzes action, and causes the entire computing operation to a halt. On the other hand, this conflict is good in that agents, when faced with such conflict, will discuss their differences and then search out alternatives which satisfy as many agents as possible. The result will be a high quality solution from the global perspective.

Conflicts in a CKBS can be distinguished into three types: **schema conflicts**, **data conflicts**, and **knowledge conflicts**. Schema conflicts result from the use of different schema definitions, such as tables, objects, and predicates, in the component systems. Data conflicts are due to inconsistent or incorrect data in the absence of schema conflicts. Even with consistent data content and schema representation among component systems, knowledge conflicts may still arise by deriving the conclusion using local inference rules and axioms (knowledge, in this sense, is an implicit form of data that must be derived through a set of nontrivial rules or methods). Schema and data conflicts have long been an active research subject in the DDB community [4, 37], whereas knowledge conflict is a central issue of DAI systems [6]; a representative set of research efforts in resolving conflicts through interagent communication is summarized in [31]. The technologies from both areas are needed to combine together to address the conflicts in CKBS. In addition, many of the current projects assume the availability of global knowledge about the application to resolve conflict or some super agent to assist agents in arriving at an agreement [30,40,18,32]. Nevertheless, our work in developing practical CKBS systems shows that there often is no universal way of coping with conflict in CKBS [43,48]. Global evaluation criteria are usually absent or are too costly to compute in such systems.

The purpose of this paper is to discuss a scheme and four methods of conflict resolution that are induced from our effort in developing cooperative knowledge systems for building design [43,48]. We call these methods: **inquiry**, **arbitration**, **persuasion**, and **accommodation**. Research in the field often notes inquiry and arbitration but often ignore the richness of persuasion and accommodation in achieving deeper

cooperative behavior. Even in our work, we can only explore limited aspects of persuasion and accommodation. Still, every method has its merits and shortcomings. These methods do not always work, and when they do they sometimes raise their own problems. The point in noting these problems is not to despair of solving multiagent conflict. Instead, it is a call for proper integration of these methods to overcoming individual shortcomings. Since agents (either human or machine) must communicate to cooperate, we illustrate these methods with a set of high-level protocols for a practical CKBS application. These protocols presume a global communication language and logic-based agents.

The organization of this paper is as follows. Section 2 introduces a concise scheme for conflict resolution in cooperative problem solving systems. Section 3 presents the particular class of problems that motivates our research: collaborative building design. It also provides an architectural overview of a CKBS prototype developed: Building Design Network (BDN). To focus on conflict issues, Section 4 sets up the basic notation for interagent communication while Section 5 discusses the notion of high-level protocols for use in the proposed resolution scheme. Following that, Section 6 proceeds to describe four modes of coping with knowledge conflicts using such high-level protocols, together with application examples taken from the BDN system. Finally, in the last section, we present the conclusion of this work ~~and discuss its extension to other related CPS fields, such as multidatabases [29, 37] and computer supported coordinated work [7].~~

## 2. Conflict Resolution Scheme

Conflict in CPS systems arises in many situations; for example, the agents are unwilling to release their resources to other agents (local autonomy), they disagree on who should do what (poor problem decomposition), or have different opinions on a choice (different perspectives). In a conflict the agents disagree, but, of course, that disagreement needs not continue. The agents might yet come to an agreement through a process of communication, thus ending the conflict. Our observation is that the conflict resolution process among agents in a cooperative computing environment can be organized in a reasonable, concise manner as sketched in Figure 2.1.

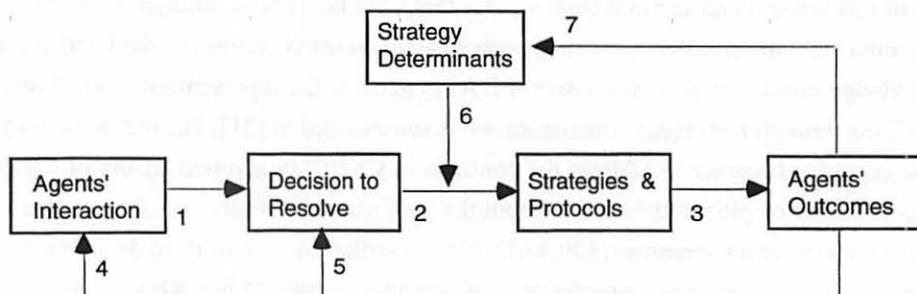


Figure 2.1. A conflict resolution scheme for CKBS.

In this scheme, the beginning point of the resolution process is an interaction between two or more information agents that runs into conflict. Depending on the type of conflict, the disputants decide whether to confront this conflict at once, deal with it later, or simply ignore it. In the case of direct or explicit conflict, the

agents must act and choose from a number of strategies and associated protocols (Arrow 2); otherwise, the distributed system operation will come to a halt. For better coordination, two or more disputants may also ask or allow a third party to assist in their selection of a conflict resolution strategy. This latter agent is called the **coordinator**. The choice of selecting an appropriate problem-solving strategy is influenced by what we call **strategy determinants** (Arrow 6), a set of defined conditions which considers various system parameters, including the rules and context of the application domain, the set of feasible outcomes, the number of disputants, the time constraint in reaching an agreement, the expectations about the probable success of different strategies, and so on. By applying the selected protocol, the agents would gather more information and produce intermediate results. These results, in turn, feed back to earlier segments. They influence future agents' interactions, which now follow the strict sequences defined in the committed protocol (Arrow 4), and the agents' future decision to resolve conflict (Arrow 5), since a settlement or improved communication may reduce the level of agents' dispute while a prolonged negative result may affect the agents' commitment to negotiation. The outcomes have also a significant effect on the choice of strategies and protocols (Arrow 7). For example, the agents are apt to choose protocols that have been successful and efficient in the past. The protocol efficiency can be gauged by the time taken to reach agreements and the agents' satisfaction of the outcomes. The following sections describe an application scenario that motivates this research and a communication framework that formalizes and extends the ideas advocated in this conflict resolution scheme into computational form.

### 3. CKBS for Building Design

One of the most crucial and time-consuming phases of building design is preliminary design, which involves layouts, selection of components (such as beams, columns, and connection plates) and materials, and specification of methods for connecting beams and columns. During this phase, feasible design alternatives are studied and the best alternatives chosen for quantitative analysis. For example, Figure 3.1 illustrates a schematic

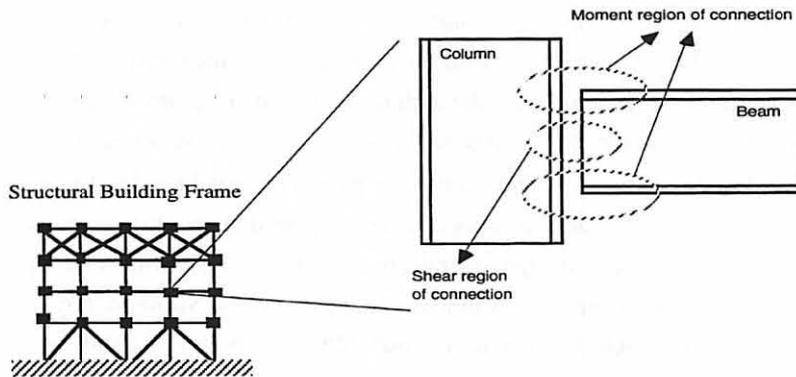


Figure 3.1. A schematic diagram showing the regions of a simple beam-column connection.

diagram of a beam-column steel connection. Basically, the problem of connection design can be decomposed into four smaller parts: the operation or connection method for the moment region and shear region, and the detail material used for those two regions.

Beam-column connection design is one of the most complex and critical parts of building design. Poor connections will have pernicious effects on the overall structural integrity and reliability of buildings. However, owing to a large number of alternatives and decision criteria involved, a full-scale analysis is infeasible. Current design practice is to select a small subset of alternatives for analysis, which runs the risk of missing many better designs. Worse, there is no effective means for the structural designer to include other participants' preferences and experience in his or her design. Figure 3.2a shows the conventional approach to building design, which is strongly sequential and iterative.

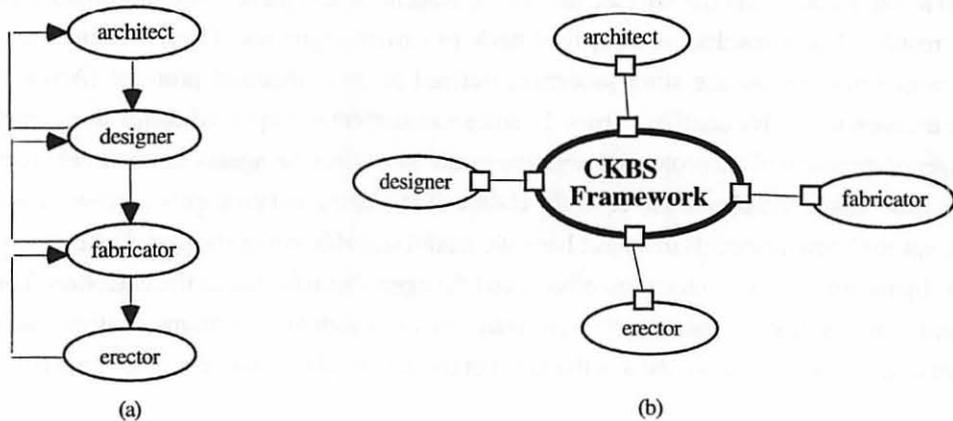


Figure 3.2: (a) the conventional mode and (b) the collaborative mode of design operation.

In Figure 3.2a, for example, the designer agent passes the preliminary design to the agent at the next stage (in this case, the fabricator). Should the latter find any disagreement, the design will be modified and passed back to the former for review. The designer agent may approve the change and return the change so that it can proceed for evaluation at the following stage, i.e., the building erection. Or, the designer agent may disagree with the change and start to negotiate with the fabricator over the conflict. Another possibility is that the designer may need further information from the architect before making any decision (who often represents the owner's view and interest), so, the draft design will pass further up the echelon. In any case, no more than two agents are involved in the discussion, and many such passes are conducted before an agreeable design is evolved. In other words, the building design defined in this manner does not fully consider the experience and resource of all participants in every design operation. It also involves many compromises among the participants to compensate for the time wasted in many iterative passes. Subsequently, costly corrections are frequently required in the field. Such deficiencies undoubtedly incur high costs and cause a time delay of the overall construction process.

By contrast, the ideal design environment would take into account all existing participants' expertise and resolve the potential discrepancies and conflicts together (see Figure 3.2b). A prototype system, **Building Design Network** (BDN), has been developed to provide such a cooperative design environment [48,43]. Different from the conventional mode of building design shown in Figure 3.2a, this distributed system

considers all feasible design alternatives and, at the same time, explores the CKBS technology to enable all participants of the design project to resolve possible conflicts together, derive a consensus of design solutions, and identify potential costs and timing problems at the early stage of the construction process. The current BDN prototype consists of three agents, the designer, the fabricator, and the erector, with the coding of relevant architectural specifications into these agents' knowledge bases. Each agent has different expertise and goals. Each can solve problems pertinent to its own, but must cooperate with others in solving three mutual design problems: connection design, beam design, and column design. A complete design problem cannot be solved by a single agent. Furthermore, there is no optimal solutions for any of these ill-defined design problems, that is, no algorithm is known to generate consistently an optimal solution of the building design problem. Satisfying solutions of these problems can only be obtained through intense communication and negotiation among agents and their users. When any conflict arises during the process of cooperation, the agents will adhere to the scheme of Figure 2.1 to resolve that conflict.

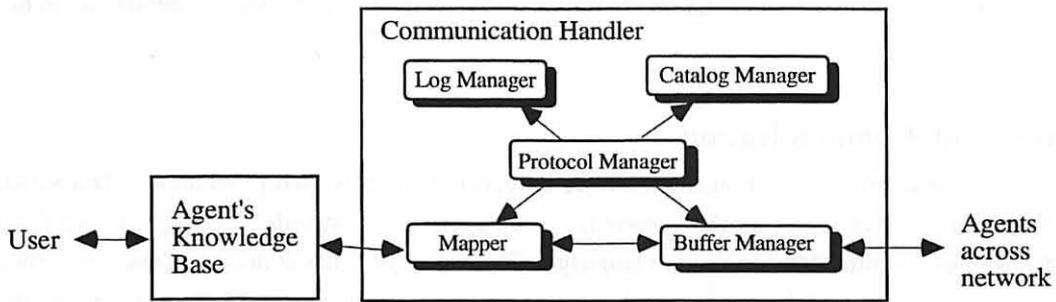


Figure 3.3: The system architecture of an agent in BDN.

In Figure 3.3, we show an overview of the system architecture of BDN agents [48,46]. Each agent resides in a separate workstation of a distributed network and communicates with the others through message passing. Each consists of two computational programs running concurrently. The **knowledge base program** in the foreground encodes domain rules and facts, problem-solving heuristics, and individual preferences of the participant of the design team, and draws conclusion from the encoded knowledge, as well as the additional information obtained during the interaction with other agents in the network. It may also query the user for any missing information during the inferencing process. Part of the agents' knowledge bases in BDN is extracted from a previous attempt to address the connection design problem [44]. The **communication handler program**, which is designed to be a background process but can be explicitly called by a user to view its records and files, manages the message transactions between agents, and maps the external representations into the internal form, or vice versa. The use of a **mapper** or translator program to resolve schematic conflicts among component systems is a predominant approach in mutidatabase systems [32,37]. The advantage of choosing a common language is to reduce the mapping complexity for different internal data and schematic representations. For  $n$  systems of different representations, we would require  $n(n-1)$  mappers without a common language, compared with  $2n$  mappers with a common language. The requirement, however, is that the semantics of the global representation must be rich enough to cover all local languages.

The **catalog manager** of a communication handler maintains a role table of every agent in any active problem solving session, specifies the type of messages permitted, and locks certain shared data which may be modified in an ongoing protocol session. The **buffer manager** manages the input and output message queues of various protocol sequences. The **log manager** keeps track of message transactions and updates, as well as other interaction statistics, in multiple log tables for performance and for possible rollback when a protocol session is aborted. The records and log tables in a communication handler are deleted automatically after the completion of a design session. The **protocol manager** coordinates various ongoing protocol sessions and regulates all the interaction between various components and the outside world. How this is done is described in [46]. In that paper we also have provided a detailed description of this handler architecture and compared it with other related work. The design of communication handler is influenced by the concept of **transaction processing monitors** in distributed database management [22,23]. Another paper has explained the strategy of problem decomposition and a scheme of cooperative decision making used in BDN [47]. The programming languages used for the BDN implementation are Prolog and C, while the underlying communication backbone is TCP.

#### 4. Interagent Communication

To focus on conflict resolution and not to be burdened with individual program and data structures of heterogeneous agents, we represent the knowledge of an agent in a first order language, L. Our prototype system, BDN, uses a similar communication language. This language L has connectives such as  $\neg$  (negation),  $\wedge$  (and),  $\vee$  (or), and  $\supset$  (material implication), but without explicit quantifiers. A sentence or lemma of L with free variables is treated as an existential sentence, as in most logic programming languages. The definition of a constant symbol also includes a list, where the empty list is denoted as  $[\ ]$ . The *set-of* predicate is used to obtain a set of objects in a local knowledge base which shares a certain property, thus, it functions as a localized quantifier. The knowledge, or the theory, of an agent can be viewed as a set of L-sentences at the global level. The logical closure of this set of sentences,  $\Gamma$ , is defined as  $Cn(\Gamma) = \{\alpha \mid \Gamma \vdash \alpha\}$ , which is often referred to as the belief state of that agent [19]. The belief state changes when the theory is revised. A global state can be considered abstractly as a union of theories in all agents. In CKBS applications, the belief state of an agent is normally taken to be consistent, that is, its underlying set of logical sentences is consistent. This, however, does not apply to the global state of the distributed system, as individual agents often differ in goals, beliefs, and inference abilities.

**Organizational hierarchies** are often used in cooperative problem solving systems to establish the problems to be solved, to decompose complex problems into simpler subproblems for different agents, and to provide effective means to coordinate activities and tasks [33, 21]. We adopt the notion of a role to indicate the position of an agent in the hierarchy. As to be discussed later in Sections 5 and 6, this role relationship affects the interactive behavior of agents and enables the modeling of various organizational problem solving strategies. Under certain operational situations, for example, a higher ranking agent may reject a request from a lower ranking one. This differs from the popular client-server mode of distributed processing.

A **message type**,  $\tau$ , denotes a procedure in the local handler which performs a certain simple computing function, such as query and commit. There are two types of messages: **action types** for initiating actions, and **response types** for replying to the former. Different message types can achieve the same function with various degrees of strength, e.g., suggesting that someone stop a task is weaker than ordering that person to stop, providing that the speaker is in the position to say such. Message types can be grouped into one class and their degrees of strength are indicated by the roles of communicating agents. The propositional content of a message,  $\rho$ , is a partial state description of the sender. It can be a term, a list, a predicate, or a sentence in  $L$ . Again, we are not concerned with the internal denotation of such content, which may be a database relation, an object attribute, or a Lisp function. We assume that the agent can translate such a content into its internal representation form. We use a first order language in the BDN prototype due to its expressive power and common usage. Our intent is not to design a language to standardize knowledge sharing, such as the goal of KIF [34].

Autonomous agents act in a cooperative manner through messages. The rule for a BDN agent to invoke a message is an expression of the form:

**[Message Invocation Rule]**  $\langle \text{Event} \rangle : \langle \text{PreCondition} \rangle \Rightarrow \langle \text{Message} \rangle : \langle \text{PostCondition} \rangle$

where  $\langle \text{Event} \rangle$  is an observable event, such as a goal, in the agent's state that triggers some action. Examples are queries, updates, and remote procedure calls.  $\langle \text{Message} \rangle$  is a message of the form  $\tau(\rho)$  with  $\tau$  as a member of the set of defined message types and  $\rho$  as one of the permissible message contents in this protocol.  $\langle \text{PreCondition} \rangle$  is a set of preparatory conditions that is evaluated in the state before the activation of  $\langle \text{Message} \rangle$ . Strategy determinants of the conflict resolution scheme in Figure 2.1 are encoded as preconditions in agents.  $\langle \text{PostCondition} \rangle$  is a set of conditions that is evaluated after the activation of  $\langle \text{Message} \rangle$ . For instance, one of the postconditions of an action message is to receive the particular type of response message.

**Example 4.1** (Message Invocation)

Suppose that a manager,  $Y$ , wants to inform employee  $X$  of an increase in  $X$ 's salary. The invocation rule of  $Y$  then is as follows (capital letters denote variables):

increase-salary( $X, \Delta$ ): work-for( $X, Y$ ),  $\Delta > 0$ , salary( $X, \text{Old}$ ),  $\text{New} = \text{Old} + \Delta$   
 $\Rightarrow \text{tell}(\text{salary}(X, \text{New})): \text{retract}(\text{salary}(X, \text{Old})), \text{assert}(\text{salary}(X, \text{New}))$ .

where *tell* is the message type,  $\langle \text{PostCondition} \rangle$  requires the updating of the salary information of  $X$  in  $Y$ , and the response type is ignored here. End of example.

This message invocation rule borrows the concept of trigger rules in *active database* [12]. It, however, further enhances that concept by adding the postcondition requirement advocated in Hoare's logic [26] to support intelligent interaction among knowledge-based agents. Using this rule, the postcondition of a message can trigger another event in the receiver agent; if the precondition of that event is satisfied, the agent will

respond with another message and execute any related action. This provides the basic mechanism to generate dialogue among autonomous agents. This invocation of a message is analogous to the utterance of an illocutionary act, which has a certain force and content and succeeds only when the hearer recognizes its intention or performative function [38].

## 5. High-Level Protocols

The language L defines a common message format and the message invocation rule provides the building blocks for advanced data communication. To support cooperative problem solving, however, we further need a set of high-level protocols that build on top of these components. Such CPS protocols are at a higher abstraction level than the conventional protocols of DDB for data transfer or deadlock prevention. These protocols are finitely describable sequences of *communication* and *computation steps*; where each step can be considered as an atomic knowledge base transaction. They convey more complex intentions than single messages and impose orderly interactions among agents. They are started and ended with communication steps. In addition, any protocol has a defined set of message types and contents.

In a communication step of a protocol, the sender encodes a specification into a message and transmits it to one or more receivers. The receiving agent in a computational step decodes and evaluates the content of an incoming message. The notation of these two transaction steps is as follows:

**[Communication Step]:**

<step\_no>: <local agent> → <remote agent> | <message type>, <message content>, <label>

**[Computation Step]:**

<step\_no>: <local agent> | <computing function>: <message content> ⇒ <result>

where <label> is a unique tag for the protocol sequence. Such a label is assigned by the initiator agent of a sequence and is included in all subsequent messages of the same sequence. It provides a means for the local agent to keep track of multiple protocol sequences occurring simultaneously. <computing function> dictates how an agent interprets an incoming <message content>. It also includes determining feasible events and checking the precondition of these events. Thus, the execution of an internal computation transaction may result in a change of state (that is, modification of a certain part of the local knowledge base) and in sending one or more subsequent messages, provided that the postcondition of these messages is satisfied.

Most speech act models in the current AI literature neither specify nor limit the number of illocutionary acts (message types) involved, because their primary interest is to produce a rich vocabulary for studying human communication and linguistics at an abstract level [25, 17, 28]. To carry such an approach to CKBS, the coding and managing of a large number of procedures that carry out the functions denoted by arbitrary message types, will simply generate too many degrees of freedom in design to arrive at good design haphazardly. Our design principle, however, is to keep the set of well-defined message types of an application as small as possible, but use them as the building blocks to form complex and powerful protocols to perform more complex kinds of intention, such as bargaining and negotiation. For example, the BDN system has three functional classes of

messages: *inquiry*, *inform*, and *complain*. Each class contains message types of the same function but with various degrees of strength depending on the role relationship between the sender and the receiver, as shown in Table 5.1.

Functional Class	Role(sender) < Role(receiver)		Role(sender) = Role(receiver)		Role(sender) > Role(receiver)	
	Action Type	Response Type	Action Type	Response Type	Action Type	Response Type
Inquiring	<i>request</i>	<i>assert/is_unknown</i>	<i>ask</i>	<i>answer/is_unknown</i>	<i>direct</i>	<i>answer/is_unknown</i>
Informing	<i>recommend</i>	<i>endorse/reject</i>	<i>offer</i>	<i>accept/refuse</i>	<i>tell</i>	<i>report/is_confused</i>
Complaining	<i>protest</i>	<i>approve/disapprove</i>	<i>complain</i>	<i>valid/ignore</i>		

Table 5.1: A list of the message types used in BDN.

where  $x/y$  refers to either type  $x$  or  $y$  in response. Moreover, types *notify* and *busy* can replace any of the response types for specific control purposes [46]. We briefly comment on the performative functions of these message types. Let  $\text{degree}(\tau)$  represent the characteristic degree of strength of type  $\tau$ .

**[Inquiring Communication Rule]:** To inquire is either to query for information or to request some action. The role relation between the sender and the receiver is used to indicate the degree of type strength. In Table 5.1,  $\text{degree}(\textit{direct}) > \text{degree}(\textit{ask}) > \text{degree}(\textit{request})$ . Such a comparison of message types has many implications. A *request* or *ask* type message lets the receiver know that the sender is either of the same rank or a lower rank. Thus, the receiver can grant or refuse the inquiry by returning messages with either an *assert* or *answer* type. However, in a *direct* act, such a refusal is precluded as it is coming from a higher authority. The receiver would send back an message of type *is\_unknown* when it does not know (i.e., cannot derive) the answer of a query.

**[Informing Communication Rule]:** To inform is either to give out information or instruct someone to do something. The content of an informative message must be in grounded form (a postcondition) while there is no such restriction for inquiry messages. An agent can accept or refuse a message of type *offer* or *recommend*, but must follow the instruction of a *tell* act. When the content of an informative message is not computable, a response act of type *is\_confused* would be sent.

**[Complaining Communication Rule]:** Complaint messages are used to express disagreement or dissatisfaction with others' choices or answers. They are used by the sender of a lower or equal rank to seek the approval of a receiver for starting negotiation over conflict. There is no complaining type defined for a sender of higher ranking. Such a sender can simply tell lower ranking receivers to start negotiation right away.

Several query strategies that make use of the notion of type strength have been described in [46]. It is worth noting that:

**Property 5.1:** Every high-level protocol involves an even number of communication steps and an odd number of computation steps.

*Reason:* The first part follows from the fact that each message of action type requires the matching of one response message; and the second part is due to every protocol starting and ending with communication steps.

The protocol manager of a communication handler uses Property 5.1 and the protocol labels to keep track of pending messages of various protocol sequences. Although this design application does not concern with low-level sensor signals, the same rule can still apply. A sensor input initiates an event of the local agent, which then checks the preconditions before sending out a message to a remote agent. If that remote agent is the same sensor agent, our scheme requires an acknowledgement in terms of a response message. The next example illustrates a simple query example based on BDN's communication framework.

**Example 5.1** (Simple Query)

The simplest form of protocols involves only two agents, two messages (action and response), and a single predicate or term as message content. Figure 5.1 shows such a protocol for single-stage query (SSQ) between two equal ranking agents.

[S1]: a → b	ask, p(x), l
[S2]: b	if compute: p(x) => q then reply q; else say nil.
[S3]: b → a	answer/unknown, q/nil, l

Figure 5.1: A single-stage query protocol, SSQ.

We elaborate the internal computation involved in SSQ in the following. The inquiry rule of agent a for a grounded predicate p(x) is (see the message invocation rule, inquiring communication rule, and Table 5.1):

goal(p(x)): not  $\Gamma_a \vdash p(x)$ , equal-role(a,b), on-line(b)  
 $\Rightarrow ask(p(x), l)$ : receive  $\tau(q)$ , check if  $\tau \in \{answer, unknown\}$ , if okay update  $\Gamma_a$

The belief state of agent b to generate an unknown response is:

receive  $ask(p(x), l)$ : equal-role(a,b),  $p \notin PRED_{SSQ}$ ,  
 $p(x)$  is ungrounded such that not  $\Gamma_b \vdash p(c/x)$  or not  $\Gamma_b \vdash \neg p(c/x)$ ,  
 $p(x)$  is grounded such that not  $\Gamma_b \vdash p(x)$  or not  $\Gamma_b \vdash \neg p(x)$   
 $\Rightarrow is\_unknown(nil, l)$ : store p(x)

where p(c/x) stands for an instantiation of p(x) (constant c substitutes x) when x is a variable, and PRED<sub>SSQ</sub> denotes the set of allowable propositional contents of SSQ. As the content of an unknown message carries no valuable information, the receiver does not need to spend time decoding and interpreting it and can go on to try some other means. The inquired agent, b, however, would store the unknown query in its database. If, in the view of future evidence, b can derive a solution to that query, it will then send the answer to the inquirer a. The log manager of the agent b's communication handler will notify the protocol manager when to check a pending unknown query again and to delete this entry after a number of retries. It also records and updates the communication transactions of the ongoing protocol sequences.

In case of a definite response, the state of b is:

receive  $ask(p(x), l)$ : equal-role(a,b),  $p \in \text{PRED}_{\text{SSQ}}, \Gamma_b \vdash q \Rightarrow \text{answer}(q, l)$

where  $q \in \{p(x), \neg p(x), p(c/x), \neg p(c/x)\}$ .

This ends the example.

**Property 5.2:** Agents are "sincere" in the sense that they are willing to share information and should not say things they believe to be false or lack evidence.

*Reason:* Follows from Example 5.1 that any agent k cannot assert a proposition  $\beta$  while  $\Gamma_k \vdash \neg\beta$  or not  $\Gamma_k \vdash \beta$ .

**Property 5.3:** Agents are "cooperative" in that they always respond according to the protocols.

Thus, the cooperating agents incorporate a three-valued logic system, i.e. they reserve their judgement on undefined propositions until more conclusive information are gathered later on. This echoes Austin's Maxims of Quality: the speaker should not say what he believes to be false and should not say that for which he lacks evidence [3]. This approach differs from predominant two-valued systems for reasoning about incomplete knowledge, such as default reasoning or closed world assumption [16]. The systems are not mutually exclusive, the former can subsume the latter [35]. We believe that this notion of "sincere" evaluation provides more leeway for reasoning about uncertain knowledge in a distributed systems environment.

## 6. Four Modes of Conflict Resolution

This section discusses four classes of strategies of the conflict resolution scheme shown in Figure 2.1. Each class of strategies will be illustrated with an application example taken from BDN using a high-level protocol. Our emphasis here is on the ideas behind each class of strategy and not on the specific details of protocols. For better presentation, at times, we will skip some obvious computational steps in listing a protocol. The conflict resolution methods can be combined in a pre-defined order such that if one fails the next is tried. This increases the flexibility and robustness of cooperating computing systems. When all methods fail, the system will inform the users about the deadlock state and request them for action, such as overwriting the conflict with a user-defined solution or re-selecting a previous tried method. Note that this previous method may not be optimal, since otherwise it would have succeeded. <sup>However,</sup> As the agents' knowledge bases may be changed during their interaction, a previous failed method may be workable in the second attempt in light of the new evidence.

### 6.1 Inquiry Mode

If the cause of conflict is due to opposing beliefs or judgments, focusing on background information relevant to them may be the thing to do. In the method of inquiry, the agents find and collect the underlying, more simple and basic, data and beliefs about the conflict. They, then, resolve the conflict by appealing to the retrieved

data and beliefs and some shared principles for interpretation. This method requires that agents structure their knowledge into lemmas or logical clauses, and use these lemmas for proof, and that they only *unwind* them when challenged. Not all CPS systems can support this method as the knowledge of their agents may not be transparent. Nevertheless, the inquiry in CKBS amounts to incorporating a certain mechanism into the agents to retrieve the inference steps that derive opposing beliefs.

Suppose that agent  $x$  deduces a conclusion  $\alpha_n$  in the  $n^{\text{th}}$  step such that  $\alpha_1, \dots, \alpha_{n-1} \vdash \alpha_n$  where  $\alpha_i$  is a belief deduced in the  $i^{\text{th}}$  step. An opposing agent  $y$  who disagrees with  $x$  can inquire what the premises of  $\alpha_n$  are in a step-by-step manner, i.e.,  $\alpha_{n-1}$  then  $\alpha_{n-2}$  then ..., until both agents pinpoint the source of their disagreement. Sometimes,  $x$  may want to retrieve all these steps in one inquiry, but this could incur a high computation overhead on  $y$ . The strategy to use would depend on the tradeoff between the communication cost and the retrieval cost. Retrieving knowledge represented in common representation schemes, such as production rules and frames, is rather straightforward. For instance, retrieving background knowledge in a rule-based agent simply requires the tracing back of the sequence of production rules being fired.

The following example considers a multiple-stage query (MSQ) protocol; the concept of a multiple-stage query is common among DAI systems based on contract nets [36] and partial global planning [15]. The MSQ protocol illustrated below, however, generalizes the SSQ protocol of Example 5.1. An agent keeps collecting premises about the conflict until it locates the problematic lemma, then it finds a way to change the use of that lemma, usually by filling in missing data or correcting wrong data. This example concerns a connection design scenario in which the erector agent disagrees with the designer agent's choice.

### Example 6.1 (Inquiry)

Referring to the building connection diagram of Figure 3.1, let us denote the method predicate for connecting a region of this connection as:

$$\text{connect-method}(\text{column}(\text{Place}, \text{Method}), \text{beam}(\text{Place}, \text{Method}), \text{Region}, \text{Status})$$

where  $\text{column}(\text{Place}, \text{Method})$  indicates the operation place and the joining method of a column section of this connection region,  $\text{beam}(\text{Place}, \text{Method})$  indicates the similar items for the beam section,  $\text{Region}$  indicates the region of the connection, and  $\text{Status}$  indicates the nature of this statement. In this design application, the status of certain domain data can be assigned as  $f$  (a fact that cannot be changed),  $o$  (an opinion that is subject to change in lieu of new evidence), or  $u$  (a sentence for which the status is unknown), such that facts are preferred to opinions and opinions are preferred to data of unknown certainty. This notion of ordering is inspired by the epistemic ordering proposed by Gardenfors for modeling logic-based agents' beliefs [19].

Let us consider the two choices about the moment region in the connection below:

$$\begin{aligned} \text{sw-fb-m} &=_{\text{def}} \text{connect-method}(\text{column}(\text{shop}, \text{welded}), \text{beam}(\text{floor}, \text{bolted}), \text{moment}, f) \\ \text{sw-fw-m} &=_{\text{def}} \text{connect-method}(\text{column}(\text{shop}, \text{welded}), \text{beam}(\text{floor}, \text{welded}), \text{moment}, f) \end{aligned}$$

where the term "shop" denotes operation in the factory, "floor" denotes on-site operation, and  $\text{sw-fb-m}$  and  $\text{sw-fw-m}$  are shorthand notations. Let us suppose that the designer agent,  $d$ , prefers  $\text{sw-fb-m}$  to  $\text{sw-fw-m}$  from the

viewpoint of reliability, i.e.  $\Gamma_d \vdash \text{choice}(>, \text{sw-fb-m}, \text{sw-fw-m}, \text{rel})$ , but the erector agent,  $e$ , disagrees. The reliability criterion,  $\text{rel}$ , addresses the quality and failure rate of the building using a particular connection method.

[S1]: $e \rightarrow d$		<i>ask</i> , $\text{why}(\text{choice}(>, \text{sw-fb-m}, \text{sw-fw-m}, \text{rel}), 1\text{-step}), l$
[S2]: $d$		$\text{why}:\text{choice}(\dots) \wedge \text{better-joints}(\dots) \wedge \text{same-quality-control}(\dots)$
[S3]: $d \rightarrow e$		<i>answer</i> , $\text{better-joints}(\text{sw-fb-m}, \text{sw-fw-m}, o) \wedge \text{same-quality-control}(\dots), l$
[S4]: $e$		$\text{Cn}(\Gamma_e^{S4}) = \text{Cn}(\Gamma_e^{S3} \cup \{\text{better-joints}(\text{sw-fb-m}, \text{sw-fw-m}, o), \text{same-quality-control}(\dots)\})$
[S5]: $e \rightarrow d$		<i>ask</i> , $\text{why}(\text{better-joints}(\text{sw-fb-m}, \text{sw-fw-m}, o), 2\text{-steps}), l$
[S6]: $d$		$\text{why}:\text{better-joints}(\text{sw-fb-m}, \text{sw-fw-m}, o) \Rightarrow \text{floor-joint}(>, \text{bolted}, \text{welded}, o)$ ; $\text{why}:\text{floor-joints}(>, \text{bolted}, \text{welded}, o) \Rightarrow \text{welding-quality}(\text{average}, o)$
[S7]: $d \rightarrow e$		<i>answer</i> , $\text{welding-quality}(\text{average}, o) \supset \text{floor-joint}(>, \text{bolted}, \text{welded}, o), l$
[S8]: $e$		$\text{Cn}(\Gamma_e^{S8}) = \text{Cn}(\Gamma_e^{S7} \cup \{\text{floor-joint}(>, \text{bolted}, \text{welded}, o), \text{welding-quality}(\text{average}, o)\})$
[S9]: $e \rightarrow d$		<i>ask</i> , $\text{why}(\text{welding-quality}(\text{average}, o), 1\text{-step}), l$
[S10]: $d$		$\text{why}:\text{welding-quality}(\text{average}, o) \Rightarrow \text{location}(\_, \text{building}, u)$
[S11]: $d \rightarrow e$		<i>answer</i> , $\text{location}(\_, \text{building}, u), l$
[S12]: $e$		$\text{location}(\_, \text{building}, u) \notin \text{Cn}(\Gamma_e^{S12})$ but $\text{location}(\_, \text{building}, f) \in \text{Cn}(\Gamma_e^{S12})$
[S13]: $e \rightarrow d$		<i>offer</i> , $\text{location}(\text{texas}, \text{building}, f), l$
[S14]: $d$		$\text{Cn}(\Gamma_d^{S14}) = \text{Cn}(\{\Gamma_d^{S13} \sim \{\text{location}(\_, \text{building}, u)\}\} \cup \{\text{location}(\text{texas}, \text{building}, f)\})$
[S15]: $d \rightarrow e$		<i>accept</i> , $\text{nil}, l$
[S16]: $e \rightarrow d$		<i>ask</i> , $\text{welding-quality}(X, S), l$
[S17]: $d \rightarrow e$		<i>answer</i> , $\text{welding-quality}(\text{good}, f), l$
[S18]: $e$		$\text{Cn}(\Gamma_e^{S18}) = \text{Cn}(\{\Gamma_d^{S17} \sim \{\text{floor-joint}(>, \text{bolted}, \text{welded}, o), \text{welding-quality}(\text{average}, o)\}\} \cup \{\text{welding-quality}(\text{good}, f)\})$
[S19]: $e \rightarrow d$		<i>ask</i> , $\text{choice}(C, \text{sw-fb-m}, \text{sw-fw-m}, \text{rel}), l$
[S20]: $d$		$\Gamma_e^{S20} \vdash \text{choice}(=, \text{sw-fb-m}, \text{sw-fw-m}, \text{rel})$
[S21]: $d \rightarrow e$		<i>answer</i> , $\text{choice}(=, \text{sw-fb-m}, \text{sw-fw-m}, \text{rel}), l$

Figure 6.1a: An instance of MSQ protocol.

Figure 6.1a shows a negotiation between the two agents, where  $\Gamma_x^N$  and  $\text{Cn}(\Gamma_x^N)$  denote the theory and the belief state, respectively, of agent  $x$  at step  $N$ . The high-level protocol makes explicit any change of belief state of agents. For example, the new states of the erector's knowledge base occur at [S4] and [S8]; on the other hand, we have  $\Gamma_e^{S4} = \Gamma_e^{S5} = \Gamma_e^{S6} = \Gamma_e^{S7}$ , if the interaction of the erector agent with the other two agents is restricted to this MSQ over the time frame of these four steps.

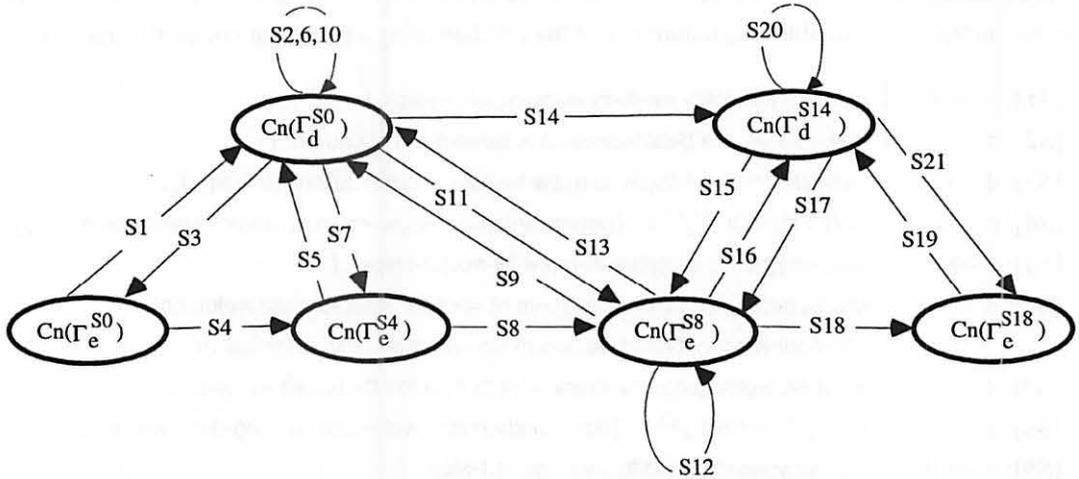


Figure 6.1b: The state transition graph of the MSQ protocol of Figure 6.1a.

Figure 6.1b illustrates a graphical representation of this inquiry process. For clarity, we skip some computational steps, but keep all communication steps (by Property 5.1, the number of communication steps in every protocol must be even) The "why" predicate is a meta-interpreter routine that can retrieve antecedent lemmas from n-steps back in the proof tree. In [S2], for instance, the designer agent retrieves the immediate antecedent lemmas that sw-fb-m has better joints than sw-fw-m while the quality control in both methods is the same, as their operations are done at the same place (column side at shop and beam side on field). Not satisfied with the answer, the erector agent probes further. In Steps [S6-7], the designer gives the reason that sw-fb-m has better quality joints than sw-fw-m two inference steps up the proof tree; because it believes the welding quality in this project is average, and thus bolting would be a better method than welding. Welding is generally a less reliable method for connection, as there are many ways that a welded connection can break compared to a bolted connection. Welding in the shop, however, is compensated by welding machines, thus, there is better control of the quality of welded products. In this case, the problem of sw-fw-m is at its beam section (welding site). The knowledge that determines the welding quality by the designer is represented as follows:

<Designer Knowledge Base>

...

location(texas,building,S)  $\supset$  welding-quality(good,S).

location(new-york,building,S)  $\supset$  welding-quality(bad,S).

location(X,building,S)  $\wedge$   $\neg$ member(X,[texas,new-york])  $\supset$  welding-quality(average,o).

The welders are of high calibre in the Texas area due to the experience gained in constructing offshore oil rigs, which have much more stringent welding specifications than buildings. It is certain that welding is as good as bolting if the building site is in Texas. The designer agent is not aware of the building site (see [S10]), and by default, it guesses that the welding quality is average (see the location rule given above). When the designer is told the fact that the site is in Texas ([S13] and the update operation at [S14]), it changes its impression that

welding quality is now good and becomes indifferent between *sw-fb-m* and *sw-fw-m*. The previous disputable result,  $\text{choice}(>, \text{sw-fb-m}, \text{sw-fw-m}, \text{rel})$ , now vanishes due to the new belief state of the designer agent,  $\text{Cn}(\Gamma_d^{S14})$ . This is the end of the example.

Example 6.1 illustrates the type of knowledge conflict that is caused by inconsistent data and missing information. The conflict resolution process is nonmonotonic as it involves modifying the agents' knowledge bases, e.g., the designer agent updates its belief about the building site at Step [S14]. Collecting more information or evidence via the inquiry could lead the agents to agreement and thus end the conflict, with the assumption that the agents, if not the users, have sufficient knowledge (often domain-specific) to evaluate the gathered evidence. Many methods in current CPS systems adopt the inquiry mode [30,45,27]. Even assuming that such a comparison is possible, many problems still remain. For example, a conflict may not derive from any difference in the agents' premises or plans. The conflict may persist no matter how much the agents have learned about each other's premises. In other words, no amount of inquiry will provide agreement. On the other hand, the agents may not have time for the inquiry. For instance, in real-time applications, fast action often has to be taken while the agents are still in conflict on the matter. A common mode of coping with these problems is **Arbitration**.

## 6.2 Arbitration Mode

An arbitration method is a procedure for selecting an outcome out of many competing alternatives. There are many such methods. In a personal conflict, we may toss a coin; this is arbitration. For explicit conflict among agents, consider two competing alternatives, *x* and *y*, in an issue. Two agents prefer *x* over *y*, while a third agent prefers *y* over *x*. Applying the simple majority rule, the outcome would be that *x* is preferred to *y*. Knowing how the vote came out, all the agents then accept the outcome. More sophisticated methods in selecting fair outcomes are the subject of social-choice theory [39]. These methods generally include an agenda that contains a list of criteria for each mutual problem. The agents would first form individual preferences (by qualitative means, such as preference relations, or quantitative means, such as utilities or probabilities) from the competing alternatives according to the problem criteria. They, then, apply some procedure to select an outcome, or a "best" choice, out of the individual preferences.

### Example 6.2 (Arbitration)

The agenda of this problem contains three major criteria: the cost of erecting the building, the cost of fabricating parts, and the cost of the joint material used, such as bolts, welds, and fillets. The first criterion is pertinent to the erector, while the latter two are concerned with the fabricator. Table 6.1 shows a case of individual preferences for five alternative connections. The same shorthand notation used in Example 6.1 is used for denoting these alternatives.

Rank	Erection cost	Material cost	Fabrication cost
1	sb-fb-m	fw-fw-m	fw-fw-m
2	sb-fw-m	sw-fw-m	fw-fb-m
3	sw-fw-m	sb-fw-m	sw-fw-m
4	fw-fb-m	fw-fb-m	sb-fw-m
5	fw-fw-m	sb-fb-m	sb-fb-m

Table 6.1 Preference orderings about direct cost.

The reader can get a feel for these rankings via the following reasoning. The erector agent prefers work to be done in the shop whenever possible, so it can spend fewer man-hours (thus, less labor cost) on the project, e.g., all floor-based methods are ranked lower. If the work has to be on site, the erector would favor bolting over welding, e.g.  $sb\text{-}fb\text{-}m > sb\text{-}fw\text{-}m$ . The bolting operation is faster and less hazardous, and thus affects labor cost and safety; the latter can be translated into the costs of insurance and worker compensation. In contrast, the fabricator prefers work to be done on site to reduce the load of fabricating parts. If the work must be done in shop, the agent would like to use less material, i.e., it prefers welding over bolting. Obviously, no amount of inquiry will help to alleviate the disagreement, as each agent has a different priority.

[S1]: $f \rightarrow d$	<i>protest</i> , protocol(arb,mc-method), <i>la</i>
[S2]: $d$	if approve then inform all agents;   else tell f to abort
[S3]: $d \rightarrow f$	<i>approve</i> , nil, <i>la</i> (assume approve)
[S4]: $d \rightarrow e$	<i>tell</i> , start(protocol(arb,mc-method)), <i>la</i>
[S5]: $e \rightarrow d$	<i>report</i> , okay, <i>la</i>
[S6]: $d \rightarrow e$	<i>direct</i> , order(ere-cost,L1,O1)
[S7]: $d \rightarrow f$	<i>direct</i> , order(fab-cost,L2,O2) $\wedge$ order(mat-cost,L3,O3), <i>la</i>
[S8]: $e$	order:ere-cost $\Rightarrow$ e-list;
[S9]: $f$	order:fab-cost a f-list; order:mat-cost $\Rightarrow$ m-list
[S10]: $e \rightarrow d$	<i>reply</i> , order(ere-cost,e-list,o), <i>la</i>
[S11]: $f \rightarrow d$	<i>reply</i> , order(fab-cost,f-list,o) $\wedge$ order(mat-cost,m-list,o), <i>la</i>
[S12]: $d$	borda:e-list,f-list,m-list $\Rightarrow$ d-list;   top:d-list $\Rightarrow$ fw-fw-m
[S13]: $d \rightarrow e$	<i>direct</i> , best(mc-method,fw-fw-m,o)?, <i>la</i>
[S14]: $d \rightarrow f$	<i>direct</i> , best(mc-method,fw-fw-m,o)?, <i>la</i>
[S15,16]: $e, f$	if okay then accept;   else protest
[S17]: $e \rightarrow d$	<i>reply</i> , accept/protest, <i>la</i>
[S18]: $f \rightarrow d$	<i>reply</i> , accept/protest, <i>la</i>

Figure 6.2: An arbitration protocol, arb.

In Figure 6.2, we show an arbitration protocol, arb, that uses the Borda position rule for selecting the best choice from a set of conflicting preferences [39]. This protocol assigns one agent as a high-ranking coordinator and the rest as peers. Since, by law, the designer is responsible for the entire building structure, the designer thus plays the role of the coordinator in this problem, the mc-method problem (the connection method for the moment region). Note, the primary function of the coordinator is to synchronize message transactions and to monitor negotiation among agents -- **not** to resolve the conflict. Its other functions depend on the problem solving methods used in the organization. In this case, for instance, it also acts as an arbitrator. As shown in Step 1 of Figure 6.2, the fabricator agent is the one who complains and initiates the arbitration. It has the motivation to do so as it has two criteria while the erector has only one. Moreover, one should also observe that the ordering of certain steps in the arbitration protocol can be interchanged, e.g. [S3] with [S4], [S6] with [S7], and so on, and the data consistency and performative intention of that protocol would still be preserved. Instead of the majority rule, the designer decides to use the Borda rule in Protocol arb to address a shortcoming of the preference relation of the majority rule, that is, the relation does not capture the notion of preference intensity. Group choice should depend not merely on individual orderings, but on their intensities of preference as well. Let the total number of choices in an ordering of criterion  $i$  be  $n(i)$ , the rank of a choice,  $x$ , in  $n(i)$  be  $r(i,x)$ , and the total number of criteria in the agenda be  $m$ . The weight of  $x$  in the total ordering is, thus, calculated as follows (we take each criterion to be of equal importance):

$$w(x) = \sum (n(i) + 1 - r(i,x)), \text{ for } 1 \leq i \leq m.$$

As an example, in Table 6.1,  $w(\text{sb-fb-m}) = 7$  and  $w(\text{sb-fw-m}) = 9$ . The choice obtained under this arbitration procedure is fw-fw-m, that is, this connection region should be welded on site. If any agent disagrees with the new choice, it can protest to the designer who may agree to start another round of arbitration under a different procedure suggested by the protestor. This ends Example 6.2.

This approach to arbitration has its own problems. Back in the Fifties, Arrow proved that fair or just methods for resolving social conflict are impossible to obtain, even for a few reasonable criteria [2]. In the case of the Borda rule, if some agent includes an irrelevant alternative in its preference ordering, it may upset the accumulated weight, or worse, the new alternative may be given the highest weight. This result has triggered an avalanche of research into this subject matter, but without any conclusive results. Today, any computer designer who wants to employ such arbitration in resolving conflict must be aware of its inadequacies.

Another popular approach is to include some measurement of the strength of the agents' knowledge, such as uncertainty factors, probabilities, epistemic postulates, and modal operators. Consider a proposition  $\alpha$ , say, an agent  $x$  asserts  $N\alpha$  ( $\alpha$  is necessarily true) while agent  $y$  claims  $P\neg\alpha$  (possibly false), where the epistemic operator  $N$  is stronger than  $P$ . Then,  $y$  would retract its claim and accept  $N\alpha$ . This agent might also have to revise some of the premises of the retracted belief in the process. This mode presumes the existence of a common reference for strength assignment. But such a reference generally would be hard to obtain in practice: it assumes the pre-existence of objective, global knowledge to do so. In BDN, for instance, we found limited use of facts (f-sentences) to overwrite opinions (o-sentences). With the exception of mathematical algorithms,

physical attributes of parts, and geometric constraints on parts, the agents' knowledge and experience in this design domain is largely opinionated. Thus, we see in both approaches the danger of arbitrarily ending conflicts through the improper use of arbitration.

Further, an arbitration method divides the agents in conflicts into winners and losers, and the opinions or beliefs of the losers are ignored. Arbitration settles conflicts by solving the problem of action so that the agents can be on the move, but it leaves the conflicts unresolved. The erector agent in Example 6.2 is certainly unhappy about the arbitrated choice, fw-fw-m, as this is the last thing it desires. Agents remain at odds with each other about their opposing beliefs on the matter as well as on their conflict. It may be asked, why bother? The agents are finally on the move. Why not just get on with it now? Why worry about the losers? The reason is clear. This runs against the consensus spirit of cooperative problem solving. The hidden sources of unresolved conflict will, sooner or later, resurface again, or worse, may cause a better solution to the global problems to be missed. Here, we ask more than arbitration can offer. Another method we can try is **Persuasion**.

### 6.3 Persuasion Mode

This way of coping with conflict is common among people, and, like cooperation, it is broad and hard to explain (just look at the amount of social science literature devoted to this topic with different interpretations). In the collaborative design of a structural building, we take a *narrower* view of persuasion and build it based on the way that arbitration works. Persuasion aims at agreement among agents, but not in the way that inquiry does. In inquiry, the goal is to settle the conflict by smoothing out the differences. Persuasion starts off from the premise that such agreement is not possible (or agreement can only be obtained if forced). Instead of figuring out acceptable solutions under the existing issues or agenda, this method attempts to **reshape** that agenda, or, rather, tries to replace the set of issues with different, but closely related, one.

#### Example 6.3 (Persuasion)

We illustrate this concept of persuasion with the same connection design example shown in Figure 6.2. This supposes that the erector agent disagrees strongly, and that no arbitration methods under the same set of criteria will provide for agreement. Since the erector agent cannot overwrite the designer agent's decision (being of lower rank) or alter the fabricator agent's preferences (same rank), it tries to persuade the designer (the coordinator of the project) to accept a new criterion. The erector can do so through sending the protest message at the last step of the arbitrary protocol at Figure 6.2, similar to a **distributed two-phase commit** protocol used in multidatabases [41], where a call for voting for all agents is issued in the **prepare phase** (Steps [S13 - 14]) before any change is committed in the **commit phase** (i.e. [S17 - 18]). Furthermore, the prepare phase has an associated timeout. If some participant does not vote within that time, the coordinator interprets the inaction as a no vote.

Figure 6.3 shows such a communication event of persuasion.

Steps [S1-5] similar to [S1-5] of Figure 6.2 but with the erector agent now as the initiator.

[S6]: e → d		<i>request</i> , agenda(mc-method,2nd(L),o), <i>lp</i>
[S7]: d → e		<i>assert</i> , agenda(mc-method,2nd(2-list),o), <i>lp</i>
[S8]: e		select:2-list => issue(rel); find support lemmas
[S9]: e → d		<i>request</i> , add(rel,1st(old-list)) ∧ support(location(new-york,building,f)), <i>lp</i>
[S10]: d		if $\Gamma_d^{S9} \cup \{\text{location}(\text{new-york,building,f})\} \vdash \text{upgrade}(\text{rel},[\dots],\text{o})$ then accept; else tell e to abort
[S11]: d → f		<i>direct</i> , agenda(mc-method,1st(new-list),o)?, <i>lp</i>
[S12]: f		if okay then accept; else protest
[S13]: f → d		<i>report</i> , okay, <i>lp</i> (assume okay)
[S14]: d → e		<i>assert</i> , agenda(mc-method,1st(n-list),o), <i>lp</i>

Figure 6.3: A protocol for persuasion.

In building design, there is usually a set of less important criteria in every problem agenda. The secondary criteria are considered only under exceptional conditions, such as when the agents cannot decide on a choice using the primary set of criteria. For the mc-method problem, reliability (rel) is a secondary issue by default. But the erector checks its database and finds that there is a high failure rate of buildings using the same connection method in the New York area ([S8]). Though it does not know the correlation between them, it suggests that the designer make reliability a primary criterion, add(rel,1st(old-list), based on this finding (a heuristic rule). The latter agent yields, as this request is consistent with the following clause in its knowledge base:

<Designer's Agenda Upgrade Rule>:  
 $\text{best}(\text{mc-method}, \text{fw-fw-m}, \text{Status}) \wedge \text{welding-quality}(\text{bad}, \text{Status1}) \wedge \text{agenda}(\text{mc-method}, 1\text{st}(\text{list}), \text{o})$   
 $\wedge \neg \text{member}(\text{rel}, \text{list}) \supset \text{upgrade}(\text{rel}, [\text{agenda}(\text{mc-method}), \text{agenda}(\text{sc-method})], \text{o}).$

The above clause can be read as: if the choice requires welding the connection on site only, i.e., fw-fw-m, and the welding quality is bad (derived from the location rule stated in Example 6.1), reliability becomes a concern (see Example 6.1) and should be made a primary criterion. At the same time, this rule also updates the agenda of a closely related problem, the sc-method problem, which is the operation method for the shear moment connection region (see Figure 3.1). Table 6.2 shows the designer's preference ordering with respect to reliability.

Rank	Erection cost	Material cost	Fabrication cost	Reliability	Overall
1	sb-fb-m	fw-fw-m	fw-fw-m	sb-fb-m	sw-fw-m
2	sb-fw-m	sw-fw-m	fw-fb-m	sw-fw-m	[sb-fw-m, sb-fb-m, fw-fw-m]
3	sw-fw-m	sb-fw-m	sw-fw-m	sb-fw-m	fw-fw-m
4	fw-fb-m	fw-fb-m	sb-fw-m	fw-fb-m	
5	fw-fw-m	sb-fb-m	sb-fb-m	fw-fw-m	

Table 6.2: New preference orderings after the persuasion process.

In general, for reliability, shop operation has better quality control and, for floor operation, bolting is preferred over bad welding. For example, we have  $sb-fb-m > sw-fw-m$ . Using the Borda weighting rule of Example 6.2, the overall ordering calculated table indicates that  $sw-fw-m$  is the new "best" choice for the revised agenda. A glance at individual orderings shows that this is a reasonable choice for all agents.

This ends the example.

One can see from this example that persuasion works by getting the agents to turn away from the existing criteria and to consider another set instead. What is the motivation for this? *Changing the agenda is changing the subject of the conflict*. Under the new criteria that the persuader brings up in the agenda, it is expected that other agents would likely agree with its proposal. By getting the opposing agents to accept its agenda, the persuader hopes to get them to join its side. Looking from another angle, we can say that persuasion is the attempt by one agent to change opposing agents understanding of something, or to get them to see things in some way that prompts them to act as they would not have done otherwise.

This kind of persuasion, that is, changing the problem agenda, enables us to incorporate deeper cooperative behaviors into autonomous agents. But, as with inquiry and arbitration, this strategy also has its shortcomings. For knowledge base applications, persuasion can sometimes lead to dead ends. The important question this raises is, when does this result in a dead end? This corresponds to some questions that other methods raise. The question regarding inquiry is: when is an inquiry conclusive? When have the agents collected information that should lead them to choose a certain option? This is a problem of epistemology and decision theory. The question for arbitration is: when is it fair? This is a topic of social choice theory. The question of persuasion boils down to this: when are the shared data and knowledge misrepresented? When do agents misunderstand each other's intentions and plans? Coupled with inquiry, persuasion can, to a certain extent, help agents clear up some of their misunderstandings, but many still would remain. This goes beyond the issue of conflict and steps into the subject of knowledge representation.

Let us shelve this difficult question and return to persuasion as an attempt to get the opposing agents to side with persuaders. An obvious problem is that the opposing agents might resist the new agenda proposed by the persuader. For instance, the designer in Example 6.3 might disregard the fabricator's request. If persuasion does not work, more persuasion might now be tried, but the prospects of this are not good. A different approach is in order. We come to a fourth mode of coping called **Accommodation**.

#### 6.4 Accommodation Mode

There are circumstances where agreement between agents cannot be reached, but action must be taken; otherwise, cooperation would have to be aborted. The aforementioned modes deal with isolated cases of conflict. They view conflict resolution as a single-shot process. Conversely, the method of accommodation takes account of the fact that a conflict is often part of history. A group of agents differ in what they now want, and some of them have to lose. But this might not be the last time they will meet nor will the circumstances remain unchanged. The accommodation mechanism records relevant information in each cycle of conflict resolution and looks at what might be done somewhere down the road. For example, in design, the winners may decide to

include a rejected criterion proposed by the losers into the agenda. They may add certain previously ignored propositions into their knowledge bases.

Is accommodation in CKBS a kind of payoff, such as a compromise or compensating losers for their losses by giving them something else? No, for this is a psychological matter that is of concern to people. At the computer level, the point is to let machine agents recognize the intentions and plans of one another so that they can better extract and synthesize decentralized information to solve hard problems. Accommodation relates conflicts with the dynamicity of agent knowledge. The winners retain the basic objectives and requests of the losers. When appropriate, they "accommodate" these propositions into their plans so as to move towards global coherence (the storing and checking of these unresolved propositions can be done by the local handler in the background). This mode of coping with conflict goes below the surface of isolated problem cases, and attends to the losers' views in the overall problem solving process.

**Example 6.4 (Accommodation)**

Figure 6.4 shows a communication event that exhibits the concept of accommodation. In this event, the fabricator agent disagrees with the designer's choice during the persuasion process shown in Figure 6.3. The protest is launched at Step [S18] of Figure 6.3 using the arbitrary protocol that is similar to two-phase commit protocol. In Step [S7], following the erector's strategy in the Example 6.3, the fabricator agent attempts to persuade the designer to upgrade another minor criterion in its favor. The criterion, stiffness (*sti*), concerns the rigidity of the connection, that is, moment resistance or shifting of the beam-column connection. Welding usually provides better rigidity than bolting. Stiffness and reliability are related to a certain extent. Thus, the argument of the fabricator agent is that since reliability and stiffness are related secondary criteria, the designer should upgrade them both at the same time to be fair.

Steps [S1-5] similar to [S1-5] of Figure 6.2 with acc now as the conflict resolution protocol.

[S6]: f		select:2-list => issue(sti);
		search for support lemmas
[S7]: f → d		request, add(sti,1st(new-list)) ∧ support(related-2nd(rel,sti,o)), la
[S8]: d		not approved as top:d-list => sw-fw-m;
		if connection-type:building => type(one)
		then remove rel from 1st(new-list)
[S9]: d → f		assert, reject ∧ not-in(rel,agenda(sc-method,1st(new-list),o),f), la

Figure 6.4: A protocol sequence for accommodation.

As listed in Table 6.2, the top choice is sw-fw-m, which exhibits strong stiffness of connection. Thus, there is no reason for the designer agent to change its plan. The agent, however, detects that the building connection is of Type I (type(one) in [S8]). The reliability and stiffness are both minor issues in the shear region of such a connection type. It, thus, removes the reliability criterion in the sc-method agenda (using the upgrading rule of

problem agenda stated in Example 6.3, this agenda was upgraded automatically in the previous persuasion) and informs the fabricator agent about the accommodation that it has made.

Example 6.4 is ended here.

As shown in Example 6.4, accommodation, unlike the other methods, can be considered as a **post-decision strategy**. This mode of resolving conflict has its share of problems too. There may be conditions that the winner agents cannot accommodate; for example, when the accommodation calls for them to give up their autonomy completely, or the winners have no plan to accommodate opposite views. What, then, is the tradeoff between accommodation and the agents' integrity, such as autonomy, goals, and private knowledge? What indeed are the states or conditions for accommodating opposing views? What is the test of global coherence? The use of accommodation must properly consider these questions. The knowledge-based agents must be equipped with sophisticated plans or with the ability to generate such plans dynamically to address these difficult issues. This is an area where CKBS can benefit from the work done in AI planning research [1].

From Examples 6.1-4, one should observe how different methods of conflict resolution complement each other using the notion of high-level protocols; where one fails, the other method picks up. For example, a protest at [S17] or [S18] of the arbitration protocol in Figure 6.2 triggers the event of persuasion in Figure 6.3; a protest at [S13] of the persuasion protocol in Figure 6.3 starts the mode of accommodation in Figure 6.4. This agrees with the concept of proposed conflict resolution scheme, which the agents' outcomes (in these cases, the configuration of beam-column connection of structural building) have a significant effect on the choice of resolution strategies and protocols. This is done mostly via the strategy determinants (see Arrows 6 and 7 of Figure 2.1) encoded in the pre- and postconditions of agents' message invocation rules. A conflict resolution session in BDN using the scheme of Figure 2.1 thus involves a combination of protocol sequences and communication modes. In principle, these modes can be combined in any order, but from the building design examples illustrated in this section, we can induce the following two properties of .

**Property 6.1:** During the conflict resolution process, inquiry should be the first method to try in a conflict resolution session, while arbitration should be the last method to use.

*Reason:* We suggest the mode of arbitration to be the last one to try as it guarantees a solution by enforcing it with a pre-defined decision rule, although it may not be fair and optimal (see the discussion of Example 6.2). On the other hand, the mode of inquiry should be first, as the cooperating agents may need to gather missing information, correct wrong data, and understand the capabilities of their cooperating partners before engaging a more sophisticated method of resolving the conflict at hand.

**Property 6.2:** None of the four modes of conflict resolution discussed involves compromise, that is, where an agent backs off and does not get what it wants.

Property 6.2 follows the discussion of the examples in this section. Even in the case of accommodation, the view and request of a "loser" at a single decision point will be incorporated into the overall problem solving

process. This property is essential in collaborative design, as the goal is to achieve an optimal solution that takes into account the expertise and knowledge of every participant in the entire design cycle.

## 7. Conclusions

This work has contributed to one important question -- the existence of conflict among cooperative knowledge based systems. Techniques discussed should also be applicable to other related fields of cooperative problem solving, such as multidatabases (MDB) and computer supported coordinated work (CSCW) [29, 8, 24, 7].

The MDB researchers focus on data and schema conflicts while the DAI researchers address the issues of conflicts at the knowledge level. Real-world problems, however, suffer from problems with all three kinds of conflicts. The use of a common language for communication in our study to alleviate schematic conflicts among agents follows the multidatabase approach [37,29]. The use of high-level protocols to resolve data and knowledge conflicts resembles the DAI way. On the other hand, there is a notion that *global domain knowledge* or use of a *super agent* is a sure way to resolve conflict in distributed knowledge systems. This carries forward the knowledge-intensive approach of building isolated systems to a distributed environment. Knowledge is there, we just need to work harder to acquire it. This is true when the problems and tasks of a single system are rather clear and the problem-solving knowledge comes from one source. In comparison, however, the problems faced in CPS and related fields are usually fuzzy and ill-defined. The knowledge involves multiple sources with different opinions, intentions, goals, and priorities. The agents have to negotiate their differences and coordinate tasks with incomplete information about the problems.

The focus of this paper is on the conflict that causes the **problem of action** during collaborative computation. We described a general scheme of conflict resolution and four viable, and complementary, strategies of conflict resolution, together with explanations and application examples from a working prototype,

BDN has been successfully evaluated to support collaborative design of beam-column connection of structural buildings and to resolve conflicts encountered during this new design process [43,49]. The conflict resolution methods drawn upon concepts from speech acts, database management, epistemic ordering, multiagent planning, knowledge representation, and social choice theory. Designing algorithms with that sort of holistic integration in mind requires an altogether different set of criteria than designing a resolution algorithm within the narrow confines of its functional specification. Furthermore, past studies have concentrated on inquiry and arbitration, but few of them notice the benefits of persuasion and accommodation. The latter two are just as important, if not more so, for coping with multiagent conflict. Understanding them will help to incorporate deeper and more powerful cooperative problem solving methods into machine agents. Still, every method has its merits and shortcomings. We have pointed out the problems of these methods and have, probably, raised more questions than we have settled (see Section 6).

The future CPS environment will consist of a geographically distributed network supporting a large number of autonomous, heterogeneous information systems. Can the high-level protocols deal with conflicts of agents at this scale? In addition to the complexity of writing translation code between the global language and each of the local representation languages, it certainly would be difficult for one coordinator to monitor

negotiation among distributed agents on this scale. One way is to organize the agents into groups and assign one agent in a group to represent that group in negotiation. Such an agent will discuss the immediate results with the local group. In a way, we envision a hierarchy of intelligent problem solvers using similar high-level protocols at various layers in the organization. Another related area of interest is concurrency control for multiple protocol sequences that are executed simultaneously to solve independent problems or multiple conflicts, e.g., two protocol sequences,  $S_1, S_2, \dots, S_n$  and  $T_1, T_2, \dots, T_m$ , that involve the same set of knowledge bases and overlap in their protocol steps.

Will the proposed methods work in other CPS fields? We believe so, but with some reservation. First, knowledge base agents have *communicative* and *planning* abilities that current database agents lack. To enhance multidatabases, we may need to add a set of mediators, as advocated in [45], in between local databases and the users. A mediator is an autonomous agent that embodies a specific planning, administrative, and communicative capability. A group of mediators would enable the user to freely access various information, explicit or embedded, in distributed local databases. The communication handler of an agent is, essentially, a localized mediator, but with little planning capabilities, as the inference and interaction are driven by the agent's knowledge base. If the mediator architecture is realizable, one may be hard pressed to tell such a multidatabase system apart from a cooperative knowledge-based system.

Second, the method of inquiry will be hard to implement in systems that solve problems based on game-theoretic analysis or state transition diagrams [9,10]. These systems use state-space or numerical representation of agents' choices, preferences, and beliefs rather than detailed description of the knowledge governing the information, which the method of inquiry tries to unwind. Without the assistance of in-depth inquiry, the efficacy of the other methods will be diminished. In addition, the long-standing focus of negotiation in this field is largely on the planning and search techniques that would satisfy various kinds of constraints, such as temporal and resource constraints, imposed by applications [31,40]. The ideas of persuasion and accommodation, we believe, will bring a new look and add dimensions to the negotiation study in CPS systems.

## References

- [1] J. F. Allen, C. R. Perrault, "Analyzing intention in utterances," *Artificial Intelligence*, Vol. 15, pp.143-178.
- [2] K. J. Arrow, *Social Choice and Individual Values*, Wiley, 1951.
- [3] J. L. Austin, *How to Do Things With Words*, Oxford University Press, Oxford, 1962.
- [4] C Batini, M Lenzerini, and S Navathe, "A comparative analysis of methodologies for database schema integration," *ACM Computing Surverys*, Vol. 18, No. 4, Dec. 1986, pp.323-364.
- [5] A. Barr, P. R. Cohen, E. A. Feigenbaum (eds.), *The Handbook of Artificial Intelligence*, Vol. 4, Morgan Kauffman, 1989.
- [6] A. H. Bond, L. Gasser, "An Analysis of Problems and Research in DAI," in A. H. Bond, L. Gasser, (eds.), *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, 1988, pp. 3-35.
- [7] J. M. Bowers, S. D. Benford, (eds.), *Studies in Computer Supported Cooperative Work - Theory, Practice and Design*, North-Holland, 1991.
- [8] M. L. Brodie, "Future Intelligent Information Systems: Artificial intelligence and database techniques working together," in J. Mylopolous, M. L. Brodie, (eds.), *Readings in Artificial Intelligence and Database*, Morgan Kaufmann, 1989, pp. 623-641.
- [9] S. E. Conry, R. A. Meyer, V. R. Lesser, "Multistage Negotiation in Distributed Planning," in A. Bond and L. Gasser (eds.) *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, 1988, pp. 367-384.
- [10] S. E. Conry, K. Kuwabara, V. R. Lesser, and R. A. Meyer, "Multistage Negotiation for Distributed Satisfaction," *IEEE Transaction on Systems, Man, and Cybernetics*, Nov/Dec 1991, pp. 1462-1477.
- [11] Daisy Working Group, "Distributed Aspects of Information Systems," *Proceedings of the European Teleinformatics Conference*, 1988, pp. 1029-1049.
- [12] U. Dayal, M. Hsu, and R. Ladin, "Organizing long-running activities with triggers and transactions," *Proc. ACM SIGMOD Int. Conf. on Management of Data*, 1990.
- [13] S. M. Deen (ed.), *Cooperative Knowledge Based Systems 1994*, Springer-Verlag, 1994.
- [14] E. H. Durfee, V. R. Lesser, D. D. Corkill, "Cooperative Distributed Problem Solving," in Barr, A., Cohen, P. R., Feigenbaum E. A. (eds.), *The Handbook of Artificial Intelligence*, Vol. 4, Morgan Kauffman, 1989, pp. 83-148.
- [15] E.H. Durfee, V.R. Lesser, "Partial global planning: a coordination framework for distributed hypothesis formation," *IEEE Transaction on Systems, Man, and Cybernetics*, Sept/Nov., 1991, pp. 1167-1183.
- [16] D. W. Etherington, *Reasoning with incomplete information*, Morgan Kaufmann, California, 1988.
- [17] F. Flores, M. Graves, B. Hartfield, and T. Winograd, "Computer Systems and the Design of Organization Interaction," *ACM Transactions on Office Information Systems*, 6(2), 1988, pp. 153-172.
- [18] J. R. Galliers, "Cooperative Interaction as Strategic Belief Revision," in S. M. Deen (ed), *Cooperative Knowledge Based Systems 1990*, Springer-Verlag, 1991, pp. 148-163.
- [19] P. Gardenfors, *Knowledge in Flux -- modeling the dynamics of epistemic states*, MIT Press, Cambridge, Massachusetts, 1988.

- [20] L. Gasser, R. W. Hill, "Coordinated Problem Solvers," in *Annual Reviews of Computer Science*, Vol. 4, (1989-90), pp. 203-254.
- [21] L. Gasser, "A Dynamic Organizational Architecture for Adaptive Problem Solving," in *Proceedings of AAAI-91*, July, 1991, pp. 185-190.
- [22] J. Gray (ed.), *The Benchmark Handbook for Database and Transaction Processing Systems*, Morgan Kaufmann Series in Data Management Systems, Morgan Kaufmann San Mateo, CA, 1990.
- [23] J. Gray, A. Reuter, *Transaction processing: Concepts and technologies*, Morgan Kaufmann Series in Data Management Systems, Morgan Kaufmann, San Mateo, CA, 1993.
- [24] I. Greif (ed.), *Computer-Supported Cooperative Work: A book of readings*, Morgan Kaufmann, CA, 1988.
- [25] B. J. Grosz, C. L. Sidner, "Plans for Discourse," in Cohen, Morgan, and Pollack (eds.) *Intentions in Communication*, MIT Press, 1990, pp.417-444.
- [26] C.A.R. Hoare, *Communicating sequential processes*, Prentice Hall, 1985.
- [27] M. N. Huhns, D. M. Bridgeland, "Multiagent Truth Maintenance," *IEEE Transaction on Systems, Man, and Cybernetics*, Nov/Dec 1991, pp. 1437-1445.
- [28] M. N. Huhns, D. M. Bridgeland, and N. V. Arni, "A DAI communication aide," *Proc. of 10th Int. Workshop on DAI*, Chapter 29, 1990.
- [29] A. R. Hurson, M. W. Bright, "Multidatabase Systems: An Advanced Concept in Handling Distributed Data," in M. Yovits, (ed.) *Advances in Computers*, Vol. 32, 1991. pp. 150-195.
- [30] M. Klein, "Supporting Conflict Resolution in Cooperative Design Systems," *IEEE Transaction on Systems, Man, and Cybernetics*, Nov/Dec 1991, pp. 1379-1390.
- [31] B. Laasri, H. Laasri, S. Lander, and V. Lesser, "A Generic Model for Intelligent Negotiating Agents," *International Journal of Intelligent and Cooperative Information Systems*, Vol. 1, No. 2, June 1992, pp. 291-318.
- [32] W Litwin, L Mark, and N Roussopoulos, "Interoperability of multiple autonomous databases," *ACM Computing Surveys*, 22(3), Sept. 1990, pp. 267-293.
- [33] T. W. Malone, "Organizing Information Processing Systems: Parallel between human Organizations and Computer Systems," in S. P. Robertson, W. W. Zachary, J. B. Black (eds.), *Introduction to Cognition, Computation, and Cooperation*, Ablex, 1990, pp. 56-83.
- [34] R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, W. Swartout, "Enabling Technology for Knowledge Sharing," *AI Magazine*, Fall, 1991, pp.36-56.
- [35] T. Przymusiński, "Three-valued nonmonotonic formalisms and semantics of logic programs," *Artificial Intelligence*, vol. 49, May 1991, pp. 309-343.
- [36] R. G. Smith, "The contract net protocol: High-level communication and control in a distributed problem solver," *IEEE Trans. Computers*, Vol. C-29, Dec. 1980, pp.1104-1113.
- [37] A Sheth, J Larson, "Federated database systems for managing distributed, heterogeneous, and autonomous databases," *ACM Computing Surveys*, 22(3), 1990, pp. 184-236.
- [38] J. Searle, *Speech Acts*, Cambridge University Press, 1969.

- [39] A. Sen, *Choice, Welfare, and Measurement*, MIT Press, 1982.
- [40] K. Sycara, "Resolving goal conflicts via negotiation," *Proceedings of AAAI-88*, August, 1988, pp. 245-250.
- [41] J. D. Ullman, Chapter 9, *Principles of Database and Knowledge-Base Systems*, Vol. 1, Computer Science Press, 1988.
- [42] F. von Martial, *Coordinating plans of autonomous agents*, Lectures Notes in Artificial Intelligence, Springer-Verlag, Heidelberg, FRG, 1991.
- [43] S. J. Wagaman, J. L. Wilson, "Design and Construction Information Systems at the ATLSS Center", NSF-ATLSS Center Report No. 92-08, June, 1992.
- [44] K Werkman, "Knowledge-based model of negotiation using shareable perspectives," *Proceedings of DAI-10*, Oct., 1990, Chapter 6.
- [45] G. Wiederhold, "Mediators in the Architecture of Future Information Systems," *IEEE Computer*, March 1992, pp.38-49.
- [46] S. T. C. Wong, "COSMO: A communication scheme for cooperative knowledge-based systems," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 23, No. 4, 1993.
- [47] S. T. C. Wong, "Preference-based decision making for cooperative knowledge-based systems," *ACM Transactions of Information Systems*, January, 1995.
- [48] S. T. C. Wong, *A Framework for Description and Design of Cooperative Knowledge-based Systems*, Ph.D. thesis, CSEE Department, Lehigh University, October, 1991.
- [49] G. Zlotkin, J. S. Rosenschein, "Cooperation and Conflict Resolution via Negotiation Among Autonomous Agents in Noncooperative Domains," *IEEE Transaction on Systems, Man, and Cybernetics*, Nov./Dec., 1991, pp. 1317-1324.

Manuscript received ...

**Affiliation of the authors:**

Stephen T. C. Wong is with the Department of Radiology, School of Medicine, University of California, San Francisco, CA 94143-0628, USA.

**Acknowledgement**

This research was funded by the NSF Engineering Research Center at Lehigh University: Advanced Technology for Large Structural Systems (ATLSS) Center and ICOT, The laboratory of Fifth Generation Computer Project of Japan.

# A Deductive Object-Oriented Database System for Situated Inference in Law

Stephen Wong Satoshi Tojo

## Abstract

Deductive Object-Oriented Databases and Situation Theory are two important areas of research in the fields of database and of linguistics. AI and Law is a new field attracting both AI researchers and legal practitioners. Our research brings together the former two fields with the aim of designing knowledge applications in the latter. This is achieved through a formal model for legal reasoning, *SM*, and a deductive object-oriented database system, *QUIXOTE*. The purpose of this paper is to introduce the key features of this formal model, based on situation theory, and to describe how this database system can implement this abstract model for complex legal reasoning applications. Concrete examples from legal precedents are used to illustrate these advanced features.

**Index Terms** deductive object-oriented databases, AI and Law, knowledge based management, situation theory.

## 1 Introduction

The research issues in Artificial Intelligence (AI) and Law include the interpretation of open-textured concepts, reasoning by cases and rules, creating computational decision making models that embody the norms of society, and drawing arguments under opposing viewpoints and different situations. Typically, a legal reasoning system draws arguments by interpreting judicial precedents (old cases) or statutes (legal rules) encoded in its knowledge base, and a more advanced system includes both kinds. Surveys on the leading projects can be found in [13, 14, 7, 1]. Most implementations are written in AI programming languages, such as Prolog or Lisp, and contain only small sets of cases and rules. They cannot access and manipulate large amounts of data and lack database management services such as concurrency control, nested transactions, and data persistence. Reasoning in law, however, is a knowledge-intensive endeavor. The lack of tools to scale up these legal reasoning prototypes is a major handicap to the growth and potential contributions of this interdisciplinary field. On the other hand, the database (DB) community has yet to develop tools which are expressive enough to satisfy the data modeling needs of the AI and Law researchers.

Legal reasoning systems has been a key research activity in the Fifth Generation Computer System (FGCS) project [11, 12]. This project devised a formal model of legal argumentation, *SM*, [16], based on situation theory [3, 2], and developed a Deductive Object-Oriented Database (DOOD) System, *QUIXOTE* [18], whose representation language can map the conceptual formulation into a computational form on the Parallel Inference Machines (PIM) [15]. The legal reasoning system developed includes a control program and a set of knowledge bases. The control program is written in the parallel logic programming language, KL1 [5]. The set of knowledge bases includes a dictionary of legal ontologies, a database of old cases, and a database of statutes. In this paper, we discuss the specific features of the *QUIXOTE* system, that can be used to support situated inference and to manage legal databases of various sorts. In addressing the complex issues of AI and Law, this study has brought together two previously unrelated fields, deductive object-oriented database and situation theory. This work, to our knowledge, is the first attempt to provide an advanced knowledge base management system tool to build large scale knowledge systems for legal reasoning applications. This paper is organized as follows. Section 2 describes the modeling of legal knowledge and reasoning at the abstraction level, using the theory of situations. Section 3 discusses the realization of this formulation at the database level using *QUIXOTE*. Section 4 illustrates situated inference mechanisms supported by this database system, and presents legal examples. We discuss the related work in Section 5 and conclude this paper in the last section.

## 2 Formal Representation of Legal Knowledge

As our formulation of legal inference is based on *situation theory*, we call it a *situation-theoretic model (SM)*. A legal concept exhibits *open texture* in that it is precisely defined *only* for those cases which have been decided by a court. The interpretation of such vague and discretionary legal concepts depends on the situations surrounding new cases. Many problems in natural language understanding are also ascribed to such situation dependency, and various semantics have been proposed, e.g., situations [2] and DRT [9]. One advantage of situation theory is its uniform way of representing various kinds of *situatedness*, i.e.,  $s \models \sigma$ , the interpretation of a phrase or sentence,  $\sigma$ , under the scope of a situation  $s$ . Our observation is that legal situations can be defined abstractly in terms of a set of infons or sentences about a case. The presumption is that abstract situations and the constraints between them would describe the *logical* flow of information in real situations [6] and would therefore be useful to the design of legal reasoning systems.

### 2.1 General Terms

The ontologies of *SM* include objects, parameters, relations, infons, and situations. An object designates an individuated part of the real world: a constant or an individual in the sense of classical logic. A parameter refers to an arbitrary object of a given type. An  $n$ -placed relation is a property of an  $n$ -tuple of argument roles,  $r_1, \dots, r_n$ , or slots into which appropriate objects of a certain type can be anchored or substituted. An infon  $\sigma$  is written as  $\langle\langle Rel, a_1, \dots, a_n, i \rangle\rangle$ , where *Rel* is a relation, each argument term  $a_k$  is a constant object or a parameter, and  $i$  is a polarity indicating 1 or 0 (true or false). If an infon contains an  $n$ -place relation and  $m$  argument terms such that  $m < n$ , we say that the infon is *unsaturated*. If  $m = n$ , it is *saturated*. Any object assigned to fill an argument role of the relation of that infon must be of the appropriate type or must be a parameter that can only anchor to objects of that type. Argument roles that must be filled to result in a saturated infon is dependent upon what the relation is [16].

An infon that has no free parameters is called a *parameter-free* infon; otherwise, it is a *parametric* infon. If  $\sigma$  is an infon and  $f$  is an *anchor* for some or all of the parameters that occur free in  $\sigma$ , we denote, by  $\sigma[f]$ , the infon that results by replacing each  $v$  in the domain of  $f$  that occurs free in  $\sigma$  with its value (object constant)  $f(v)$ . If  $I$  is a set of parametric infons and  $f$  is an anchor for some or all of the parameters that occur free in  $I$ , then  $I[f] = \{\sigma[f] \mid \sigma \in I\}$ . In addition, an abstract situation is said to be *coherent* if it does not support both an infon and its negation. If an infon is of polarity 1, its negation is of polarity 0. Two abstract situations  $s$  and  $s'$  are said to be *compatible* if their union is a coherent situation. The situations within a legal case are presumed to be compatible with one another, but no such presumption can be made across different cases.

A *SM* is a triplet  $\langle \mathcal{P}, \mathcal{A}, \models \rangle$ , where  $\mathcal{P}$  is a collection of abstract situations including judicial precedents, a new case,  $c_n$ , and a world,  $w$ , that is a unique maximal situation of which every other situation is a part;  $\mathcal{A}$  are the defendant and plaintiff agents; and  $\models$  is the support relation. The latter satisfies the following conditions [6]:

#### Condition 2.1 (Supports Relation)

- i. For any  $s \in \mathcal{P}$ , and any atomic infon  $\sigma$ ,  $s \models \sigma$  if and only if (iff)  $\sigma \in s$ .
- ii. For any  $s$ , any  $\sigma, \beta$ , (a) For any  $s$  that contains (as constituents) all members of  $u$ ,  $s \models (\exists x \in u)\sigma$  iff there is an anchor,  $f$ , of a parameter,  $\dot{x}$ , to an element of  $u$ , such that  $s \models \sigma[f]$ , and (b)  $s \models (\forall x \in u)\sigma$  iff for all anchors,  $f$ , of  $\dot{x}$  to an element of  $u$ , we have  $s \models \sigma[f]$ .
- iii. For any  $s \in \mathcal{P}$ , and any set of infons  $I$ ,  $s \models I$  if  $s \models \sigma$  for every infon  $\sigma$  in  $I$ .  $\square$

The notation  $s \models \sigma$  thus denotes a proposition about  $\sigma$  whose truth values are situation-dependent, whereas  $w \models \beta$  asserts that  $\beta$  is universally true. In addition, let  $v$  be a parameter. By a *condition* on  $v$  we mean any finite set of parametric infons. (At least one of these should involve  $v$ , otherwise, the definition is degenerate). We define a new parameter,  $v \parallel C$ , called a *restricted parameter*.  $v \parallel C$  will denote an object of the same basic type as  $v$ , that satisfies the requirements imposed by  $C$ . This amounts to our placing a more stringent requirement on anchors.

### 2.2 Concept Matching

We introduce certain specific terms, *relevance level*, *infony matching*, and *situation matching*, to extend the general *SM* terms into the legal domain. In a legal event, an agent would consider some facts (infons) to be more relevant than others

in reaching an argument. To estimate such weighting on facts,  $SM$  assigns every infon in an old case with a level of relevance. For example, the restricted parameter  $\dot{\sigma} = \sigma \parallel \ll \text{relevance-level}, \sigma, \lambda, 1 \gg$ , where  $\lambda$  denotes a certain weight of relevance. One distinction of legal reasoning is the matching of the new facts with those of precedents to generate similar arguments which may hold in the new case [11, 1]. No two events are exactly alike, but the idea of precedent-based matching presupposes that a prior decision will control subsequent facts that are like the first. Yet, given the lack of absolute identity, the decision-maker of the new case must evaluate the determinant of likeness. To this end,  $SM$  adopts a concept of structural matching. Since cases are composed of infons, the model first defines the matching relation between these basic units of information. A case infon is always parameter-free.

**Condition 2.2:** For  $c_n$  and an old case  $c_o$ ,  $\sigma_n = \ll Rel_1, a_1, \dots, a_n, i_1 \gg \in c_n$ ,  $\sigma_o = \ll Rel_2, b_1, \dots, b_m, i_2 \gg \in c_o$ , (a) (Exact Infon Matching):  $\sigma_n \simeq_{iem} \sigma_o$  iff (i)  $m = n$ ; (ii)  $i_1 = i_2$ ; (iii)  $Rel_1$  and  $Rel_2$  are of the same type; (iv) for every argument  $a_i$  of a non-infon type, there exists  $b_k$  which is of the same role or type and has not been matched with another argument; (v) for every  $a_j$  of an infon type, there exists  $b_g$  that satisfies the same set of conditions, and (b) (Partial Infon Matching):  $\sigma_n \simeq_{ipm} \sigma_o$  if  $m \leq n$  and all argument terms of  $\sigma_o$  are matched.  $\square$

where a  $Rel\ b$  intends to denote  $w \models \ll Rel, a, b, 1 \gg$ . Clearly,  $\simeq_{iem}$  is an equivalence relation while  $\simeq_{ipm}$  is an asymmetric relation. Infon-matching relations are the building blocks for defining situation-matching relations.

### Condition 2.3

(a) (Exact Situation Matching) For any  $s_n \subseteq c_n$ ,  $s_o \subseteq c_o$ ,  $s_n \simeq_{sem} s_o$  iff for every  $\sigma$  of  $s_o \models \sigma$ , there exists  $\rho$  of  $s_n \models \rho$  such that  $\sigma \simeq_{iem} \rho$ , and vice versa; (b) (Partial Situation Matching) For any  $s_n \simeq_{spm} s_o$  iff for every  $\sigma$  of  $s_o \models \sigma \parallel \ll \text{relevance-level}, \sigma, l, 1 \gg$  s.t.  $l \geq l_i$ , there exists  $\rho$  of  $s_n \models \rho$  s.t.  $\sigma \simeq_{ipm} \rho$ .  $\square$

When there is no confusion, we write  $\simeq_s$  to denote a matching relation between situations and  $\simeq_i$  between infons.

## 2.3 Situated Inference Rules

A legal reasoning process can be modeled as an inference tree of four layers. The bottom layer consists of a set of basic facts and hypotheses, the second involves case rules of individual precedents, the third involves case rules which are induced from several precedents or which are generated from certain legal theory, and the top layer concerns legal rules derived from statutes. An individual or local case rule is used by an agent in an old case to derive plausible legal concepts and propositions. These rules vary from case to case, and their interpretation depends on particular views and priorities. An induced case rule has a broader scope and is generalized from a set of precedents. Legal rules are general provisions and definitions of crimes. The applicability of these rules is independent of the view of either side (plaintiff or defendant) and every item of information (infon) included is of equal relevance. Though it rarely happens, it may be possible for an agent to skip one or two case rule layers in attaining a legal goal. Further, a local case rule is as follows:

**Rule 2.1 (Local Rule):** For  $c \in \mathcal{P}$ ,  $cr: c \models \sigma \leftarrow c \models I/B_{cr}$ .  $\square$

where  $I$  is called the antecedent of the rule,  $\sigma$  is the consequent infon, and  $cr$  is the label of the rule, which is not itself part of the rule but which serves to identify the rule. Sometimes, we simply write  $cr: c \models \sigma \leftarrow I/B_{cr}$ . Both  $\sigma$  and  $I$  are parameter-free. The reliability and the scope of application of a local rule will be subject to a set of *background conditions*,  $B_{cr}$ . The conditions include information such as an agent's goal and hypotheses; these are crucial in debate to establish the degree of certainty and the scope of applicability of that rule. Usually, it becomes necessary to take background conditions into account and investigate what they are when the conclusion drawn from the case rule leads to conflict with others or a change in circumstances that weakens the applicability of that rule.

Denote  $I'$  and  $\sigma'$  as two sets of parametric infons such that all parameters that occur in the latter also appear in the former. An induced rule and a legal rule are represented as:

**Rule 2.2 (Induced Rule):** For any  $c_1, \dots, c_k \in \mathcal{P}$ ,  $c = c_1 \cup c_2 \cup \dots \cup c_k$ ,  $ir: c \models \sigma' \leftarrow I'/B_{ir}$ .  $\square$

**Rule 2.3 (Legal Rule):**  $lr: w \models \sigma' \leftarrow I'/B_{lr}$ .  $\square$

where all cases are coherent and  $ir$  and  $lr$  are the rule labels.

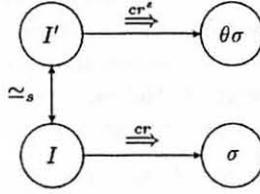


Figure 1: Case substitution.

## 2.4 Substitution and Anchoring

When a situation of a new case,  $c_n$ , supports a similar antecedent of a local rule of  $c_o$ , one can draw a conclusion about the new case similar to the consequent of that rule. That is,

**Rule 2.4** (Local Rule Substitution)

For  $c_n, c_o \in \mathcal{P}$ ,  $cr^s: c_n \models \sigma\theta$  if  $cr: c_o \models \sigma \Leftarrow I/B_{cr}$  and  $c_n \models I'/\{B_{cr}\theta \cup B_n\}$  such that  $I' \simeq_s I$ .  $\square$

where  $cr^s$  is the label of the new rule,  $B_n$  is the original background of the new case,  $I'$ , and the combined condition after the substitution,  $B = B_{cr}\theta \cup B_n$ , is coherent. The function  $\theta$  forms a *link* that connects  $c_n$  with  $c_o$  and replaces all terms (objects and relations) in  $\sigma$  and  $B_{cr}$  that also occur in  $I$  with their matched counterparts in  $I'$ . Figure 1 presents a substitution diagram that does not include the background conditions. Referring to Rule 2.4, the substitution merely replaces terms and does not change the polarities of infons. Also, the information of case matching  $B_n$  is not related to  $B\theta$  and thus does not create compatibility problems. It thus follows that  $\{c_n \cup B\}$  is coherent.

In a court case, both sides (plaintiff and defendant) are normally ignorant about the assumptions and hypotheses of each other's claims. An essential technique, used to reveal such 'hidden' information, is cross-examination. Incorporating the background conditions into legal constraints (case and legal rules) allows us to capture this essential feature of legal reasoning for knowledge-based applications. Rather than substitution, a consequent is derived from a legal rule.

**Rule 2.5:** (a) (Induced Rule Anchoring) For  $c_n, c_1, \dots, c_k \in \mathcal{P}$ , such that  $c = \{c_1, c_2, \dots, c_k\}$ ,  $ir^a: c_n \models \sigma[f]$  if  $ir: c \models \sigma \Leftarrow I/B_{ir}$  and  $c_n \models I[f]/\{B_{ir}[f] \cup B_n\}$ , (b) (Legal Rule Anchoring) for  $c_n \in \mathcal{P}$ ,  $lr^a: c_n \models \sigma[f]$  if  $lr: w \models \sigma \Leftarrow I/B_{lr}$  and  $c_n \models I[f]/\{B_{lr}[f] \cup B_n\}$   $\square$

where,  $c_n$ , again, is the new case and  $B_n$  is the background condition of this new case. Figure 2 shows a legal inference example of  $\mathcal{SM}$  in the forward reasoning manner. For simplicity, this inference involves only local and legal rules. The black circles,  $I'_1, I'_2$ , and  $\sigma$ , denote the situations of a new case,  $c_n$  while situations  $I_1$  and  $I_2$  are of old cases. Two immediate arguments,  $\beta_1$  and  $\beta_2$  are drawn using local rules  $cr_1$  and  $cr_2$ . Together with  $\sigma$ , the goal  $\gamma[f]$  is anchored or attained by applying the legal rule  $lr$ . From the case coherency condition, it follows that all concepts of a single goal tree must share the same legal perspective: the plaintiff's or the defendant's. This figure indicates that the matching relation of  $I_1$  is stronger than that of  $I_2$ . One can also probe into background conditions, linked by appropriate rule labels, of these arguments to retrieve the underlying hypotheses and legal theories.

## 3 Modeling of Legal Knowledge in *QUIXOTE*

This section introduces the *QUIXOTE* language and shows how it can be used to map the  $\mathcal{SM}$  concepts in computable form. The *QUIXOTE* language [18] is a hybrid of the deductive object-oriented database (DOOD) language [4] and constraint logic programming (CLP) language [8]. It resembles F-logic [10], in that it is a DOOD language that includes powerful extensions into logic programming such as subsumption, complex objects, and modules. Compared with conventional

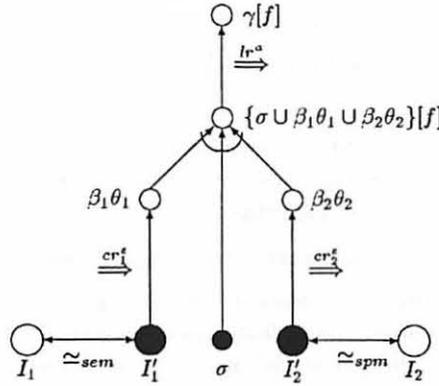


Figure 2: A visual representation of legal inference.

CLP language, *QUIXOTE* has a symbolic constraint domain which makes it suitable for describing the legal situations depicted in legal documents written in, albeit tight and formal, natural language. A typical *QUIXOTE* database includes the following: (i) the subsumption relations among basic objects, (ii) the submodule relations among modules, and (iii) rules. Our legal reasoning system consists of three distinct databases, namely, a dictionary, a case base, and a statute base. Accordingly, we first introduce the *objects* and *modules* of *QUIXOTE*, explain the data structure of a legal dictionary, and then describe the use of *QUIXOTE* rules to represent case-based rules and statutes. In *QUIXOTE*, the concepts of situation theory are rephrased as follows:

situation theory	<i>QUIXOTE</i>
situation	module
infon	attribute term
relation name	basic object
type	subsumption
role	label
supporting relation ( $\models$ )	membership in module ( $:$ )

### 3.1 Objects, Modules, and Matching

#### 3.1.1 Basic Objects, Complex Objects and Modules

The legal dictionary has two parts: the *concept lattice* and the definition of *relations* (viz. the tuples of roled slots with predicates). We first introduce the *QUIXOTE* notion of *objects* and *subsumption* relation for forming the concept lattice, and, thereafter, we describe *QUIXOTE*'s *attribute terms* to represent the relations and infons of situation theory. Object terms, *Obj*, in *QUIXOTE* consist of a set of *basic objects*, *Bobj*, a set of *complex objects*, *Cobj*, and a set of *variables*, *Var*. We denote the *subsumption* relation,  $\sqsubseteq$ , as a partial relation between basic objects such that for any  $a, b \in \text{Bobj}$ ,  $a \sqsubseteq b$  means that  $a$  is more specific than  $b$ , or, intuitively,  $a$  is *a*  $b$ . The following is an example of the subsumption relations among basic objects (In *QUIXOTE* syntax,  $\sqsubseteq$  is represented by ' $=<$ ').

infant  $\sqsubseteq$  person, baby  $\sqsubseteq$  person, person  $\sqsubseteq$  creature, lion  $\sqsubseteq$  creature,

Together with the basic objects  $\perp$  (bottom) and  $\top$  (top), we have  $\forall x \in \text{Bobj}, \perp \sqsubseteq x, x \sqsubseteq \top$ . Thus, a concept lattice of basic objects,  $\langle \text{Bobj}, \sqsubseteq \rangle$ , is a finite bounded complete partial order. A *complex object* is of the form  $o[l_1 = v_1, l_2 = v_2, \dots]$ , where  $o \in \text{Bobj}$  and for any  $l_i \in \text{Bobj}, v_i \in \text{Obj}$ ,  $l_i$  is also called a *label*. The order of labels is not strict, e.g.,  $o[l_1 = a, m = b]$

and  $o[m=b, l=a]$  are treated as being identical objects. The subsumption relation between basic objects is extended to the relation between complex objects, or between complex and basic objects, as follows:

$$h[l_1 = v_1, \dots] \sqsubseteq h'[l'_1 = v'_1, \dots] \text{ iff } h \sqsubseteq h', \forall i \exists j, l_j = l'_i, v_j \sqsubseteq v'_i$$

$$h[l_1 = v_1, \dots] \sqsubseteq o \text{ iff } h \sqsubseteq o$$

Similarly, the database operations, *meet* and *join*, between complex objects are defined as the greatest lower bound and least upper bound, since the basic objects compose a complete lattice. For example, the following relation holds when we have  $\text{strangle} \sqsubseteq \text{homicide}$  and  $\text{poison} \sqsubseteq \text{homicide}$ .

$$\text{strangle}[\text{agent} = \text{tom}] \sqsubseteq \text{homicide}, \quad \text{poison}[\text{agent} = \text{tom}, \text{coagent} = \text{mary}] \sqsubseteq \text{homicide}[\text{coagent} = \text{mary}]$$

while there is no subsumption relation between  $\text{poison}[\text{agent} = \text{tom}]$  and  $\text{homicide}[\text{coagent} = \text{mary}]$ . For object terms with variables, co-reference relation is considered in the definition as discussed here [17]. For example,  $o[l_1 = X, l_2 = X] \sqsubseteq o[l_1 = X, l_2 = Y]$ . An *attribute term* is an object term with attached property specifications, i.e., a set of ' $l=v$ '. Such a term for a complex object has the following form.

$$\underbrace{\text{basic obj. } o}_{\text{complex obj.}} \quad \underbrace{[l_1 = v_1, \dots]}_{\text{intrinsic}} / \underbrace{[l_2 = v_2, \dots]}_{\text{extrinsic}}$$

We distinguish the properties of a complex object from those of an attribute term. The former is called an *intrinsic* attribute while the latter is called an *extrinsic* attribute. The label-valued relations of attribute terms are:

$$o/[l = X] \text{ iff } o \mid \{o.l == X\}$$

$$o/[m \rightarrow U] \text{ iff } o \mid \{o.m \sqsubseteq U\}$$

$$o/[n \leftarrow V] \text{ iff } o \mid \{o.n \sqsupseteq V\}$$

where  $O \mid C$  denotes an object term  $O$  with constraint  $C$ . We introduce the *dotted term* notation,  $o.l$ , where  $o$  is an object term and  $l$  is a label, to specify the value of the  $l$  (extrinsic) attribute of  $o$ . By default, the properties of an object are inherited from related objects via the subsumption relation. If  $o \sqsubseteq p$ , then  $\forall l, o.l \sqsubseteq p.l$ . On the other hand, for complex objects, the values of intrinsic attributes override those of extrinsic ones, e.g., when  $\text{death}/[\text{cause}=\text{suicide}]$  holds,  $\text{death}[\text{cause}=\text{murder}].\text{cause}$  ( $= \text{murder}$ ) is not subsumed by  $\text{death}.\text{cause}$  ( $= \text{suicide}$ ) although we have  $\text{death}[\text{cause} = \text{murder}] \sqsubseteq \text{death}$ . These attribute terms can be used to represent *SM* infons. Let us consider the following relation (see Section 2.1):

$$\text{abandon}/[\text{agent}=\text{Agent}, \text{object}=\text{Coagent}, \text{place}=\text{Loc}]$$

$$\mid \{\text{abandon} \sqsubseteq \text{act}, \text{Agent} \sqsubseteq \text{human}, \text{Coagent} \sqsubseteq \text{human}, \text{Loc} \sqsubseteq \text{location}\}.$$

This is a *QUIXOTE* representation of the sentence, "Agent's act of abandoning Coagent at a certain place, where both Agent and Coagent are human." The subsumption relation stated in the constraints denotes the type specification in situation theory, such that the corresponding *SM* representation is:

$$\langle \text{abandon} : \text{action}, \hat{x}_{\text{agt}} : \text{human}, \hat{y}_{\text{obj}} : \text{human}, \hat{l}_{\text{obj}} : \text{location} \rangle$$

where *abandon* is of the *action* type, *agt* (agent) and *cgt* (coagent) are of the *human* type. The dictionary maintains legal relations of distinct names, and its object lattice includes the subsumption hierarchy between the relation names.

A *QUIXOTE* legal database consists of a hierarchy of modules. Each module is identified by an object term called a *module identifier* and consists of a set of rules. The rules of one module are inherited by its submodules. The submodule relation,  $\sqsupseteq_S$ , is a partial relation between module identifiers that specifies rule inheritance among modules. For example, if  $\text{case1} \sqsupseteq_S \text{case2}$ , all rules and facts in module  $\text{case2}$  are inherited by module  $\text{case1}$ . (In *QUIXOTE* syntax,  $\sqsupseteq_S$  is represented by '>-'.) We called  $\text{case1}$  a *submodule* of  $\text{case2}$  and  $\text{case2}$  a *supermodule* of  $\text{case1}$ . Module and rule inheritance are powerful devices for classifying and modeling situation-dependent knowledge. Identical objects must have equal properties within a module, but are allowed to have distinct properties between different unrelated modules. For instance, the following piece of code is consistent, provided that  $\text{sit}_1$  and  $\text{sit}_2$  are not related.

$$\text{sit}_1 :: \text{homicide}/[\text{agent}=\text{tom}];; \text{sit}_2 :: \text{homicide}/[\text{agent}=\text{mary}];;$$

### 3.1.2 Realization of Concept Matching

The concept of infon matching, stated in Condition 2.2, is realized in *QUIXOTE* as follows.

**Operation 3.1 (Infon Matching):** For any two attribute terms  $o_1$  and  $o_2$ , (1)  $o_3$  exists, such that  $o_1 \sqsubseteq o_3$ , and  $o_2 \sqsubseteq o_3$  in a given concept lattice, then  $o_1$  and  $o_2$  are interpreted as being *partially matched infons*, and (2) if the basic object parts of two attribute terms are found to be identical, the two attribute terms are interpreted as being *exactly matched infons*.  $\square$

Under Operation 3.1, for example, abandon and leave are *partially matched* if the legal dictionary contains: abandon  $\sqsubseteq$  act, leave  $\sqsubseteq$  act and abandon/[agent=jim] is *exactly matched* with abandon/[agent=tom].

**Operation 3.2 (Situation Matching):** For any  $m_1$  and  $m_2$ , (1) if, for every attribute term in  $m_1$ , there is one and only one attribute term in  $m_2$  that can match it exactly, and vice versa, then the two modules are interpreted as being *exactly matched situations*, and (2) if, for any  $o_1$  in  $m_1$  whose *relevance* value subsumes a given object (viz. the threshold level), there is an attribute term  $o_2$  in  $m_2$ , that can be partially matched with  $o_1$ , the two modules are interpreted as being *partially matched situations*.  $\square$

For example, if two modules contain:

```
m_n :: {abandon/[agent=mary, object=june], leave/[agent=mary, object=june]};;  
m_o :: {abandon/[agent=jim, object=tom], leave/[agent=jim, object=tom]};;
```

*QUIXOTE* would assert that  $m_n$  is *exactly matched* with  $m_o$ . Consider another pair of descriptions:

```
m_n :: {abandon/[agent=mary], leave/[agent=mary, object=june]};;  
m_o :: {abandon/[agent=jim, object=tom] | {abandon.relevance == 13},  
       leave/[agent=jim, object=tom] | {leave.relevance == 12},  
       poor/[agent=jim] | {poor.relevance == 11}};;
```

where  $11 \leq 12 \leq 13$ , we have: (i) if the threshold value is 12, then  $m_n$  is *partially matched* with  $m_o$ , and (ii) if the threshold value is 11, then  $m_n$  is *not partially matched* with  $m_o$ .

## 3.2 Situated Inference Rules

A *QUIXOTE* rule takes the following form.

$$\underbrace{m_0 :: H}_{\text{head}} \mid \underbrace{HC}_{\text{head\_constraints}} \Leftarrow \underbrace{m_1 : B_1, \dots, m_n : B_n}_{\text{body}} \parallel \underbrace{BC}_{\text{body\_constraints}} ;;$$

where  $H$  or  $B_i$  is a literal while  $HC$  and  $BC$  are sets of subsumption constraints. An object term,  $m_i$ , is called a *module identifier*. The above rule exists in the module  $m_0$ . Intuitively, this means that if every  $B_i$  holds in a module  $m_i$  under the constraints  $BC$ , then  $H$  and constraints  $HC$  hold in  $m_0$ , where  $H$  and  $B_i$ 's are *object terms* or *attribute terms*.  $HC$  works as constraints in the sense of conventional CLP language[8], while  $BC$  is processed abductively. Constraints in *QUIXOTE* are sets of formulas in terms of a subsumption relation among object terms and dotted terms. Each formula has the form  $\langle term \rangle, \langle op \rangle, \langle term \rangle$  where  $\langle term \rangle$  is an object term or a dotted term and  $\langle op \rangle$  is  $=$ ,  $\sqsubseteq$ , or  $\sqsupseteq$ . If the head constraints and module identifiers of a rule can be omitted, the body constraints,  $BC$ , of that rule then constitute the background conditions.

### 3.2.1 Case Representation

We give a sample case below, which is simplified from an actual legal precedent [11].

**Mary's Case:** On a cold winter's day, Mary abandoned her son Tom on the street because she was very poor. Tom was just 4 months old. Jim found Tom crying on the street and started to drive Tom by car to the police station. However, Jim caused an accident on the way to the police station. Tom was injured. Jim thought that Tom had died in the accident and left Tom on the street. Tom froze to death.

In *QUIXOTE* format, the aforementioned case contains objects, such as *mary*, *tom*, *jim*, *accident*, and *cold*, as well as several events, such as *abandon*, *find*, *make*, *injure*, *leave*, *death*, and *causes*. The relevance levels of these events are indicated through explicit attributes with ordering values.

```
&subsumption;; l1 =< l2 =< l3;;
&rule;; mary_case :: {mary, tom, jim, accident, cold,
    poor/[agent=mary, relevance=l1],
    abandon/[agent=mary, object=tom, relevance=l2],
    find/[agent=jim, object=tom/[state=crying], relevance=l1],
    accident/[agent=jim, relevance=l2],
    baby/[agent=tom, age=4months],
    injure/[agent=jim, object=tom, by=accident, relevance=l2],
    leave/[agent=jim, object=tom, relevance=l3],
    death/[agent=tom, cause=cold, relevance=l3]};;
```

There were many interpretations on the responsibilities of the actions of Mary and Jim. A lawyer might reason that: "If Mary hadn't abandoned Tom, Tom wouldn't have died. In addition, the cause of Tom's death was not injury but freezing. Therefore there exists a causality between Tom's death and Mary's abandonment." Another lawyer would, however, argue differently: "There is a crime of Jim, for his abandonment of Tom. And in addition, Tom's death is indirectly caused by Jim's abandonment. Therefore, there exists a causality between Tom's death and Jim's abandonment." These contradictory claims are documented, together with the final verdict as decided by the judge, as a judicial precedent. The next subsection, shows how to model these conflicting arguments using *case rules*.

### 3.2.2 Case, Induced, and Legal Rules

The deduction of legal arguments in *QUIXOTE* observes the following convention.

$$\underbrace{\text{result}}_{\text{Head}} \leftarrow \underbrace{\text{facts}}_{B_1, B_2, \dots, B_n} \parallel \underbrace{\text{lawyer's interpretation}}_{\text{Background conditions}}$$

Namely,  $B_i$ 's in the above are the facts that were accepted by both the plaintiff and defendant beforehand, and the set of *Background Conditions* is the interpretation of causal relations between events, scopes of an agent's intention, and so on. For example, we can represent the following case-based rules in Mary's case (see Rule 2.1).

```
c >- cr1;;
cr1 :: responsible_for_injury/[agent=jim, to=tom]
    <= accident/[agent=jim], injury/[agent=tom]
    || { injury.cause=< accident};;
```

$c >- cr1$ , claims that  $cr1$  is an extended case of  $c$ , including the case rule. This rule claims: when there existed jim's accident and tom's injury as facts, and if the injury's cause was ascribed to the accident, jim is responsible to tom for the injury. Similarly,  $cr2$  is another example of a case rule, again from Mary's case.

```
c >- cr2;;
cr2 :: responsible_for_protection_for_weak/[agent=jim, to=tom]
    <= accident/[agent=jim], baby/[agent=tom], injury/[agent=tom],
    leave/[agent=jim, object=tom]
    || { injury.cause=< accident};;
```

The idea of an *induced rule* is to abstract some of ground terms in case rules, either by common sense knowledge or by legal theories. For example, if there are several similar accident cases, the attorneys may draw the following generalization, because the causality between the accident and the injury is agreed by both sides (refer to Rule 2.2):

```
ir1 :: responsible_for_injury/[agent=X, to=Y]
    <= accident/[agent=X], injury/[agent=Y] || { X =< person, Y =< person};;
```

In the above rule, restrictions on variables X and Y are given in the background conditions, such that they have to satisfy certain roles. The following *ir2* is yet another, more abstract, form of *ir1*.

```
ir2 :: responsible/[agent=X, to=Y, for=Inj]
      <= Acc/[agent=X], Inj/[agent=Y, cause=Acc]
      || {Acc =< accident, Inj=<physical_damage, X =< person, Y =< person};;
```

In *ir2*, traffic accident and injury are abstracted to variable *Acc* and *Inj*, and are subsumed by their super concepts in the legal dictionary.

Legal rules are written in a form having free parameters. Consider the following penal code (Japanese Penal Code, Article 199): "In case an intentional action of person A causes the death of person B and the action is not presumed to be legal, A is responsible for the crime of homicide." Its *QUIXOTE* representation is (see Rule 2.3):

```
lr1 :: responsible_for_homicide/[agent=A, to=B]
      <= Action/[agent=A], illegal/[act->Action], death/[agent=B, cause->Action]
      || {Action =< intend, A =< person, B =< person};;
```

where *illegal/[agent=A, action -> Action]* claims that the action *Action* done by A, such as self-defence, is not found to be legal. The statute for the legality of self-defence is described as follows (Japanese Penal Code, Article 38):

```
lr2 :: illegal/[act->Action] <= Action || {Action =< intend};;
```

## 4 Query and Inference in *QUIXOTE*

### 4.1 Constraints and Answer with Assumption

*QUIXOTE* supports two kinds of constraints: *head constraints* and *body constraints*. During execution, the following transformation is performed first.

$$\begin{array}{c}
 m_0 :: H \mid HC \Leftarrow m_1 : B_1 \mid C_1, \dots, m_n : B_n \mid C_n \parallel \overline{Dot\_Cstr} \cup \overline{Oterm\_Cstr}; \\
 \Downarrow \\
 m_0 :: H \mid \overbrace{HC \cup \overline{Oterm\_Cstr}}^{\text{head constraints}} \Leftarrow m_1 : B_1, \dots, m_n : B_n \parallel \overbrace{\overline{Dot\_Cstr} \cup C_1 \cup \dots \cup C_n}_{\text{body constraints}};
 \end{array}$$

Constraints (subsumption relation) on object terms (*Oterm\_Cstr*) are merged to the head constraint (*HC*), and are used as background conditions for the applicability of the rule, while constraints that includes a dotted term (*Dot\_Cstr*) remain as the body constraint, and constraints on each object in the body ( $C_i$ ) are merged into the body constraint. To reply to a query, *QUIXOTE* often returns answer substitutions with a set of constraints among dotted terms called *assumptions*. An assumption is a set of unsatisfied constraints during derivation, such that they can be considered as being missing information. The control program or the user will then decide whether to fill in the missing information, or to invoke another query. Except for constraints among dotted terms, *QUIXOTE* works like a conventional CLP language [8]. However, dotted term constraints in the body constraints work as assumptions if they are not satisfied in the head constraints. In this respect, *QUIXOTE* supports abductive queries to partial information databases, and such partiality differs from *incompleteness* in databases represented as null values or Skolem constants.

The formal derivation in *QUIXOTE* is explained as follows. Let  $G_m$  be a set of goals in the  $m$ -th stage of an execution, the next set of goals is derived from the rule  $H \mid HC \Leftarrow B \parallel BC : G_{m+1} = (G_m - \{G\})\theta \cup B\theta$ , where there is a most general unifier  $\theta$  between  $H$  and  $G$ . Thus, the current goal,  $H\theta$  ( $= G\theta$ ), is removed from  $G_m$ , and new goals that are in the body part of the rule  $B\theta$  are added. When  $G_n = \phi$ , execution ends. The conclusion is the set of resolved head constraints:  $C_{m+1} = (C_m \cup HC)\theta$ , while a set of assumptions, or the remaining unsatisfied constraints,  $A_i$ , becomes:  $A_{m+1} = (A_m \cup BC)\theta - C_{m+1}$ .  $A_i$  is the accumulation of body constraints  $BC\theta$ , some of them being removed from this accumulation when they are satisfied in  $HC$  ( $= C_{m+1}$ ), and the final set of assumption,  $A_n$ , becomes the *abductive reason* for the conclusion. As an example, the following code says that there is a crime and the judgement result is guilty if self-defence is illegal, but innocent if self-defence is legal.

```

case::crime;;
case::judgement[result=guilty] <= crime/[self_defence->illegal];;
case::judgement[result=innocent] <= crime/[self_defence->legal];;

```

The first clause tells us of the existence of an object, `crime`, but nothing about the properties of its `self_defence` attribute. The second clause means that if `crime` exists in the case and the `self_defence` property is subsumed by `illegal`, `judgement[result=guilty]` holds. When we initiate a query `?-case:judgement[result=Result]`, that is, the judgement result of the crime, *QUIXOTE* returns the following two independent answers.

```

Result=guilty if case:crime.self_defence=<illegal
Result=innocent if case:crime.self_defence=<legal

```

Each answer assumes that the `self_defence` property of `crime` coming from the body of the second or third clause. Neither constraint is satisfied by the head constraint, which is empty in this example, so they are accumulated as assumptions.

## 4.2 Inference of Legal Knowledge

We list a small case example (a traffic accident) and use it to show how the induced rule `ir1` is invoked.

```

n_case :: injure/[agent=tom];; n_case :: traffic_accident/[agent=jim];;
&subsumption;; traffic_accident =< accident;; injure =< physical_damage;; person >= {jim, tom};;
&submodule;; ir2 -< n_case;;

```

Now, consider the following query, `?- n_case: responsible/[agent=jim, to=X, for=Y]`. This query may be read as “Is Jim responsible to someone, X, for something (represented variable Y)?” *QUIXOTE* returns the following answer: `IF n_case:injure.cause == traffic_accident THEN Y == injure, X == tom`. This answer says that if the cause of the injury is the traffic accident in this case, then Jim is responsible. Consider the following case, `hanako_case`, where *QUIXOTE* invokes a sequence of case and legal rules to draw a conclusion, as shown in Figure 2.

```

hanako_case ::{hanako, taro, jiro,
  death/[agent=taro, age=4months], baby/[agent=taro age=4months],
  injury/[agent=taro], abandon/[agent=hanako, object=taro],
  accident/[agent=jiro], leave/[agent=jiro, object=taro]};;

```

Using **Operation 3.2**, *QUIXOTE* would partially match `hanako_case` with `mary_case` (see Section 3.2.1) with the threshold relevance value, 12. That is, there is a rule substitution,  $\theta$ , on `cr2` (see Rule 2.4):  $\theta = \{hanako/mary, taro/tom, jiro/jim\}$ , where ‘ $x/y$ ’ stands for a substitution of  $y$  for  $x$ . `cr2_s`, as generated, is represented as follows:

```

cr2_s :: responsible_for_protection_for_weak/[agent=jiro, to=taro]
  <= accident/[agent=jiro], baby/[agent=taro], injury/[agent=taro],
  leave/[agent=jiro, object=taro] || { injury.cause=< accident};;

```

The concept of *anchoring*, mentioned in Section 2.4, is realized in *QUIXOTE* by invoking either induced case rules or statutes within a case description. Let us suppose the following submodule relation:

```

&submodule;; w >- hanako_case;; w >- lr3;; w >- cr2_s;;

```

with the following subsumption relations.

```

&subsumption;; leave =< abandon;; abandon =< intend;;

```

In addition, we need one more rule that is derived from common sense: `weak =< baby`; then with the query: `?-w:responsible_for_death_by_abandonment_of_weak/[agent=X, to=taro]`, meaning that “Is someone responsible for the death of Taro by abandoning the weak person?” *QUIXOTE* returns with two answers as follows.

```

** Answer 1 **
IF w:injury.cause =< accident
  w:leave.agent >= responsible_for_death_by_abandonment_of_weak.agent
  w:death.cause =< leave          THEN X =< jiro
** Answer 2 **
IF w:injury.cause =< accident
  w:abandon.agent >= responsible_for_death_by_abandonment_of_weak.agent
  w:death.cause =< abandon        THEN X =< hanako

```

The first answer interprets the causality in Hanako's case as: if the cause of Taro's death is some event under Jiro's leaving Taro, then Jiro is responsible for the homicide. The second answer states yet another interpretation, i.e., Hanako is responsible if Taro is killed by Hanako's intended abandonment. This rather confusing response arises from the fact that there were two deeds, leave and abandon, both of which can be regarded as being abandonment, i.e., both belong to the same class in the legal dictionary. To verify this, one can further query the database with new constraint:

```

?-w:D || {D =< abandon}.
  ** Answer 1 **
  D == leave
  ** Answer 2 **
  D == abandon

```

Thus, this section has shown that *QUIXOTE*: (i) returns answers with assumptions when there are unsatisfied background conditions for applying legal and case rules, (ii) proposes all the alternative solutions to the query program for unsatisfied background conditions, and (iii) accepts queries with additional information that has not yet been stored in its databases. These features confirm the knowledge processing capability of *QUIXOTE* in supporting situated inference within an OODB framework and in managing persistent legal data.

## 5 Conclusions

In this paper, we have outlined the motivation behind this study, presented the basic features of a formal model for legal reasoning, and a deductive object-oriented database for implementing this model. The foundation of this model, *SM*, is based on the theory of situations and clearly defines the notions of open-texture concepts and situated inference in the legal domain. The purpose of this model is to study the fundamental issues of AI and Law at the abstraction level, to help design better and more robust legal reasoning systems. In Section 2, we introduced the key features but leave a more detailed description in a future paper. In Sections 3 and 4, we described how *QUIXOTE*, a deductive object-oriented database system, is used to implement *SM* for our legal reasoning applications. In addition, we have illustrated the features of *QUIXOTE* with implemented legal examples. To the best of our knowledge, this is the first reported work that brings together two previously unrelated fields, namely, deductive object-oriented databases and situation theory, to design knowledge systems for solving complex problems and for modeling human intellectual behavior. It is also the first attempt to enhance the reasoning capability and application scale of the current generation of legal reasoning systems with an advanced database tool.

*QUIXOTE* provides a single language for both query and programming purposes, and it exhibits the inference features of deduction, object-oriented, and constraint logic programming. Most legal reasoning systems are small programs that lack the database management capability to access and store large volumes of data, presenting a stumbling block to the growth of this knowledge-intensive field. The DOOD approach is proposed here to satisfy such needs. In addition, research into legal reasoning systems is closely related to a broader and more complex field, natural language processing (NLP). The ability of DOOD systems, such as *QUIXOTE*, to model abstract concepts of situation theory in a database environment may pave the way for the natural language processing community to tackle concrete, demanding problems, such as building a comprehensive dictionary database of general linguistic concepts.

## Acknowledgements

This work was a part of the Fifth Generation Computer Systems (FGCS) project of Japan. The authors would like to thank the following ICOT members, Drs. K. Nitta, K. Yokota, and H. Yasukawa.

## Affiliation of Author

Stephen T. C. Wong, Dept. of Radiology, UCSF School of Medicine, 505 Parnassus Ave., S.F., CA 94143-0628  
Satoshi Tojo, Mitsubishi Research Institute, 8-1 Shimomoguro 1, Meguro-ku, Tokyo 153, Japan.

## References

- [1] K. Ashley, *Modeling legal argument*, MIT Press, Cambridge, MA, 1990.
- [2] J. Barwise, J. Perry, *Situations and attitudes*, Cambridge, MA, MIT Press, 1983.
- [3] J. Barwise, *The situation in logic*, CSLI Lecture Notes 17, CSLI, Stanford University, CA, 1988.
- [4] S. Ceri, G. Gottlob, L. Tanca, *Logic programming and databases*, Springer-Verlag, Berlin, 1990.
- [5] T. Chikayama, "Operating system PIMOS and kernel language KL1," *Proc. Int. Conf. of Fifth Generation Computer Systems*, ICOT, Tokyo, June, 1992, pp. 73-88.
- [6] K. Devlin, *Logic and information I*, Cambridge University Press, 1991.
- [7] A.v.d.L. Gardner, *An artificial intelligence approach to legal reasoning*, MIT Press, Cambridge, MA, 1987.
- [8] J. Jaffer, J-L. Lassez, "Constraint logic programming," *Proc. 4th IEEE Symp. Logic Programming*, 1987.
- [9] H. Kamp, "A theory of truth and semantic representation," in J. Groenendijk, T. Jansson, and M. Stockhof, (ed.), *Methods in the Study of Language Representation*. Math Carter, Amsterdam, 1981.
- [10] M. Kifer, G. Lausen, "F-logic - a higher order language for reasoning about objects, inheritance, and schema," *Proc. ACM SIGMOD Int. Conf. on Management of Data*, Portland, June, 1989, pp. 134-146.
- [11] K. Nitta, Y. Ohtake, S. Maeda, et. al "HELIC-II: A legal reasoning system on the parallel inference machine," *Proc. Int. Conf. of FGCS*, ICOT, Tokyo, June, 1992, pp. 1115-1124.
- [12] K. Nitta, S. Wong, Y. Ohtake, "A computational model for trial reasoning," *Proc. 4th Int. Conf. on AI and Law*, Amsterdam, June, 1993.
- [13] E.L. Rissland, (ed.), Special issue: AI and Legal Reasoning, Part 1, 2 *International Journal of Man-Machine Studies*, Vol. 34, No. 6, June 1991.
- [14] M. Sergot, "The representation of law in computer programs: a survey and comparison," in T.J.M. Bench-Capon (ed.), *Knowledge Based Systems and Legal Applications*, Academic Press, 1991, pp. 3-68.
- [15] K. Taki, "Parallel inference machine PIM," *Proc. Int. Conf. on FGCS*, ICOT, Tokyo, June, 1992, pp. 50-72.
- [16] S. Wong, "A situation-theoretic model for trial reasoning," *Proc. of the 6th Int. Symp. on Legal Knowledge and Legal Reasoning Systems*, Tokyo, Oct., 1992, pp. 32-54.
- [17] H. Yasukawa, K. Yokota, "Labeled graphs as semantics of objects," *Proc. SIGDBS and SIGAI of Information Processing Society of Japan*, Oct., 1990.
- [18] K. Yokota, H. Tsuda, Y. Morita, S. Tojo, H. Yasukawa, "Specific features of a deductive object-oriented database language *QUILXOTE*," *Proc. Workshop on Combining Declarative and Object-Oriented Databases*, ACM SIGMOD, Washington, D.C., May 29, 1993.

## Computers in Radiology

# Asynchronous Transfer Mode Technology for Radiologic Image Communication

H. K. Huang<sup>1</sup>, Ronald L. Arenson, William P. Dillon, S. L. Lou, Todd Bazzill, Albert W. K. Wong

Image communication is an important component in picture archiving and communication systems (PACS) and teleradiology applications. Currently, local area networks (LANs) and wide area networks (WANs) use different technologies for image communication. Asynchronous transfer mode (ATM) is an emerging technology that can be used for both LANs and WANs. This article describes experimental results using an ATM network to transmit CT scans and digitized radiographs between the University of California, San Francisco (UCSF) and Mount Zion Hospital, an affiliated community hospital in the San Francisco Bay area. The WAN connection between the two hospitals is via an ATM main switch at Pacific Bell, a local communication carrier located in Oakland, CA, which uses single-mode optical fibers. Preliminary results show that, using the ATM Optical Carrier Level 3 (OC3) (155 Mbits/sec) specification, it takes 1.3 sec and 2.7 sec to transmit a 10-Mbyte digitized radiograph and a 20-Mbyte CT scan, respectively, between the two locations. Encouraged by these results, we have designed and implemented an ATM WAN and LAN between UCSF and Mount Zion Hospital. This is the first of a three-phase project of installing a WAN serving four hospitals and one clinic in the San Francisco Bay area.

### Background

A PACS or a teleradiology system consists of four components: image acquisition devices, several computers, an archival unit, and display workstations. These components are connected by an image communication network. The PACS or the teleradiology system also needs to connect to

other medically related databases, such as a hospital information system or a radiology information system.

If the communication network is within a local area connected by cables, the network is considered a LAN. If the network requires telecommunication carriers like a telephone company, microwave dishes, or a satellite, it is called either a metropolitan area network (MAN) or a WAN, depending on the distance between nodes. For this discussion, WAN is used to represent both MANs and WANs. Because the communication media used for LANs and WANs are different, the technology for their applications is also different. In general, technologies for LANs are more versatile, and their costs are lower than those for WANs. We have reported on some LANs that have signaling rates as high as 1 gigabit/sec with realization of 24 Mbits/sec in radiologic LAN applications (compared with the standard Ethernet signal rate of 10 Mbits/sec with realization of 800 Kbits/sec) [1-3]. On the other hand, radiologic applications using WANs have been limited to the dial-up digital services zero and one. The latter, with a maximum speed of 1.544 Mbits/sec, is sometimes referred to as the T1 [4]. Although the higher-speed digital service three is available, its application is limited by costs and serviceability. Therefore, a gap in technology development and required cost exists between LANs and WANs.

The current concept in radiologic image communication is that no physical or logical boundaries should exist between LANs and WANs (Holman BL et al., presented at the NCI, NIH, and Conjoint Committee Conference on Diagnostic Radiology,

Received November 23, 1994; accepted after revision January 20, 1995.

Presented in part at the CalREN (California Research and Education Network) inauguration meeting, San Francisco, October 1994.

Supported in part by NLM HPCC Program NO1-LM-4-3508 and by CalREN ATMN-007.

<sup>1</sup>All authors: Department of Radiology, University of California, San Francisco, 530 Parnassus Ave., Rm. CL-158, San Francisco, CA 94143-0628. Address correspondence to H. K. Huang.

AJR 1995;164:1533-1536 0361-803X/95/1646-1533 © American Roentgen Ray Society

November 1994). For this reason, ATM for both LANs and WANs has become the emerging technology [5]. ATM is a method for transporting information that splits data into fixed-length cells (53 bytes). Each cell consists of 5 bytes of header information pertinent to the ATM transmission protocol and 48 bytes of information. It is based on the virtual circuit-oriented packet-switching theory developed for telephone circuit switching applications [6].

Our department has received two grants to set up an ATM WAN and LAN for transmission of radiologic images and related medical data. This network connects five sites: UCSF; Mount Zion Hospital; San Francisco Veterans Administration Medical Center; San Francisco General Hospital; and San Francisco Magnetic Resonance Center. Each site will have an ATM switch that connects to the main ATM switch in Oakland, CA. Figure 1 shows the ATM connections in the testing area. UCSF serves as the expert center, and its ATM switch is connected to the departmental PACS infrastructure. Other sites are considered satellites. Images and related patient data are transmitted from the satellite sites to UCSF for interpretation and consultation.

The ATM network provides us with an opportunity to study both LANs and WANs using a single technology. This article describes the first phase of this study: the connection between UCSF and Mount Zion Hospital. We describe an experiment that demonstrates the successful connection of a WAN and an LAN in a clinical setting with a single high-speed communication technology. We provide the performance statistics based on this experiment and illustrate the implementation of the ATM network between the two sites.

#### The Satellite Site: Current Clinical Operational Environment

Mount Zion Hospital has a GE 9800 helical CT scanner (Milwaukee, WI) and a Siemens 1.0-T MR IMAPCT scanner (Iselin, NJ) that are connected to an acquisition computer (Sparc LX; Sun Microsystems, Mountain View, CA). Direct digital CT and MR examinations from Mount Zion Hospital are transmitted to UCSF as follows: The acquisition computer acquires both direct digital CT and MR images automatically during scanning opera-

tions using a method similar to that described previously [7]. After an examination is completed, the entire study is reformatted to the ACR/NEMA (American College of Radiology/National Electrical Manufacturers Association) standard with UCSF shadow group header information, which contains information about the patient's examination and history that is relevant to the UCSF clinical environment. The formatted image data is then immediately transmitted from Mount Zion Hospital to UCSF's PACS through a T1 line for display and for long-term archiving.

For this study, we developed an automatic time log in the UCSF PACS that tabulates the time required for each stage after the PACS detects the completion of an examination [7]. The stages include image transmission from the scanner to the acquisition computer at the satellite site, image reformatting, image transmission from the satellite site to UCSF, image archiving at the PACS controller, and image display in the viewing room.

We randomly selected 600 CT scans with an average size of 20 Mbytes per scan and recorded the average time required to complete each of the five stages. The result was about 800 sec, of which 200 sec was for the T1 communication [8]. We also recorded the time required to send a 10-Mbyte digitized radiograph through the T1 line (about 100 sec). These times were used as a baseline against which the ATM technology was compared.

#### The ATM Experiment: Experimental Set-Up and Simulation

We performed the following experiment to evaluate the performance of the ATM OC3 WAN and LAN between UCSF and Mount Zion Hospital. Table 1 lists the test equipment we used, and the experimental set-up is shown in Figure 2. At Mount Zion Hospital, an ASX 200 ATM switch (FORE, Warrendale, PA) is connected to a Sparc20 computer (Sun Microsystems, Mountain View, CA) with an S-bus ATM adapter board (FORE) using two multimode optical fibers. This ASX 200 ATM switch is connected to the ATM main switch via single-mode optical fibers.

At UCSF, another ASX 200 ATM switch is connected to Sparc10, 20, and 690 MP computers using multimode optical fibers. A Sparc20 computer was used to simulate the acquisition computer at Mount Zion Hospital; a Sparc20 and 690 MP computer were used to simulate the display workstation computer and the PACS controller computer at UCSF, respectively. During the experiment, we turned the connection(s) on and off as needed. The ATM WAN and LAN throughputs were measured under various conditions.

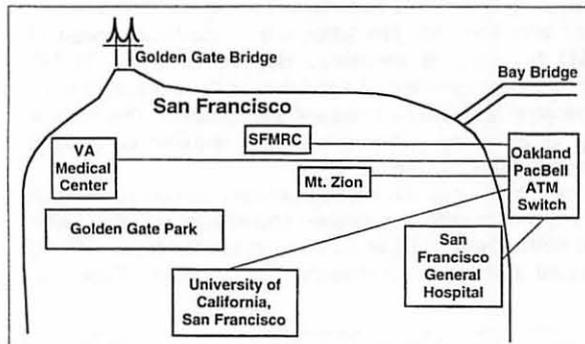


Fig. 1.—The San Francisco Bay area. Five clinical sites (University of California, San Francisco; Mount Zion Hospital; San Francisco Veterans Administration Medical Center [VA Medical Center]; San Francisco General Hospital; and San Francisco Magnetic Resonance Center [SFMRC]) are connected with asynchronous transfer mode (ATM) test network through main ATM switch at Pacific Bell, Oakland, CA. PacBell = Pacific Bell.

TABLE 1: Test Equipment Used for the Asynchronous Transfer Mode Experiment Between the University of California, San Francisco and Mount Zion Hospital

Asynchronous transfer mode (ATM) optical carrier level 3 (OC3, 155 Mbits/sec) switch at Pacific Bell, Oakland, CA
ATM switch (ASX 200; FORE, Warrendale, PA)
ATM SONET (Synchronous Optical NETWORK) OC3 adaptor boards within the ASX 200 switch
Sun Sparc20, 10, 690 MP computers (Sun Microsystems, Mountain View, CA)
S-bus ATM adaptor boards (FORE)

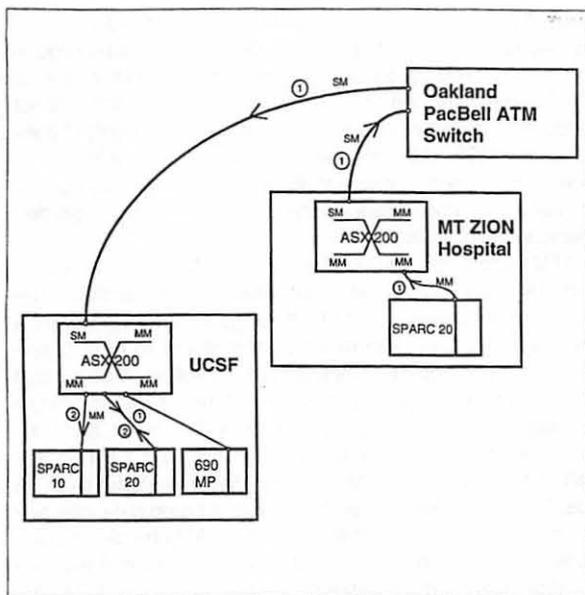


Fig. 2.—Experimental set-up for the asynchronous transfer mode (ATM) (OC3 specification, 155 Mbits/sec transfer rate) wide area network (WAN) and local area network (LAN) throughput test between University of California, San Francisco and Mount Zion Hospital using the ATM switch in Oakland, CA. Path marked with number 1 is for WAN performance measurement, path marked with number 2 for LAN performance measurement. Paths 1 and 2 combined is for both WAN and LAN measurement. PacBell = Pacific Bell, MM = multimode fibers, SM = single-mode fibers.

## Results

### Simulation

We first measured the ATM WAN performance by activating only the path marked with the number 1 in Figure 2. ATM LAN performance was measured by activating only the path marked with the number 2 in Figure 2. However, the PACS controller computer (Sparc690 MP) was running continuously, representing the background communication activities. The performance of both networks was measured by activating paths 1 and 2 simultaneously. The performance of the connection between the 690 MP and the ASX 200 was not measured because its performance represented image archiving, which always has a lower priority in the PACS design. The simulation used the following parameters: image buffer size of 128 K, measurement from computer memory to computer memory, a data set size of 256 Mbytes, and the Transmission Control Protocol/Internet Protocol (TCP/IP) communication protocol.

Table 2 shows that the ATM WAN performance is about 60 Mbits/sec and the ATM LAN is 66 Mbits/sec (or close to 40% of the 155 Mbits/sec signaling rate). When combining the WAN and the LAN concurrently, their performances decrease by 46% and 73%, respectively, but the aggregate performance reaches 77 Mbits/sec.

### Transmission of Digitized Radiographs and CT Scans

Using the same experimental set-up, we activated only path 1 to measure the WAN performance for transmitting dig-

TABLE 2: Asynchronous Transfer Mode Optical Carrier Level 3 Performance Statistics

	Wide area Network	Local area Network	Aggregate
Path 1: From Mount Zion Hospital Sparc20 (Sun Microsystems, Mountain View, CA) to UCSF Sparc20	60.64	—	60.64
Path 2: From UCSF Sparc20 to Mount Zion Hospital Sparc10 (Sun Microsystems)	—	66.64	66.64
Paths 1 and 2 concurrently	28.16	48.80	76.96

itized radiographs and CT body images. The digitized radiographs used were 10 Mbytes each, and the CT scans were about 20 Mbytes per examination. Six hundred sets of each type of examination were transmitted and measured. The statistical variance of transmission rate within each type of examination was less than 1%. The results were used as a comparison against those obtained by the T1 method, as shown in Table 3.

## Discussion

We have successfully completed an ATM performance test between UCSF and Mount Zion Hospital using an OC3 link with a transmission rate of 155 Mbits/sec. Experimental results demonstrate that we can obtain a transmission rate of about 60 Mbits/sec in both the LAN and the WAN. This performance can be translated as sending a 10-Mbyte digitized radiograph in 1.3 sec or a 20-Mbyte CT scan in 2.7 sec. We believe this rate would satisfy most radiologic image communication requirements. Our experimental results were based on the TCP/IP protocol. We have yet to optimize the ATM communication protocol, which could potentially improve its performance by 20–25%.

Because this experiment is a simulation, we have to extrapolate these results in the clinical environment. First, the images transmitted during the simulation were not obtained directly from the CT scanner host computer or from the laser film digitizer; they were from the acquisition computer. In an actual clinical network, a bottleneck occurs between the scanner/digitizer and the acquisition computer because performing a CT scan or

TABLE 3: Transmission Time Required Between Mount Zion Hospital and the University of California, San Francisco Using a T1 line and Asynchronous Transfer Mode Optical Carrier Level 3

	10-Mbyte Radiograph		20-Mbyte CT Study	
	One Image	Two Images	One Study	One Current and One Historical
T1 line <sup>a</sup>	100 sec (1.6 min)	200 sec (3.4 min)	200 sec (3.4 min)	400 sec (6.7 min)
Asynchronous transfer mode <sup>b</sup>	1.3 sec	2.7 sec	2.7 sec	5.4 sec

<sup>a</sup>Transmission rate = 1.5 Mbits/sec, realization rate = 100K/sec.

<sup>b</sup>Transmission rate = 155 Mbits/sec, realization rate = 7.5 Mbytes/sec.

digitizing a radiograph takes time. Second, aside from the acquisition node, we want to know if the simulation results can represent the remaining network performance in the clinical environment. This issue can be addressed by the characteristics of the ATM switches and adapter boards. First, ATM switch operation is, in principle, different from Ethernet connections. In the case of Ethernet, the transmission speed is reduced drastically when the number of connections increases. On the other hand, because the ATM switch operates in a star architecture, it has an aggregate data rate of 2.5 gigabits/sec at the switch. Until the number of connections to the switch reaches this limit, the performance between nodes would not deteriorate, regardless of the number of connections. When the switch reaches this limit, a second switch can be added. As for the performance at the computer connection, the ATM transmission rate depends only on the computer into which the ATM adapter board is inserted. Because we are using Sparc-series computers in the clinical environment, we do not anticipate a change in the ATM transmission performance. Thus, we are confident that the results from this simulation would be a good representation of an ATM network in a clinical environment.

Encouraged by this result, we have designed a total ATM WAN and LAN between UCSF and Mount Zion Hospital. Figure 3 shows the first phase of the system configuration. In this configuration, the existing T1 line is used as the back-up link for the WAN. For those PACS components with older computer buses that the ATM technology does not support, we use a LAN access

switch (LAX-20) that converts the ATM connection to an Ethernet connection. This technology is called the Star Ethernet configuration. In this configuration, each connection can achieve a rate of 10 Mbits/sec, and the performance of each connection is not reduced by other connections. Examples in this category of computers in UCSF PACS are the 2K display stations that use the older Sun computers with a VME-bus backplane. This network architecture has been implemented, and we are designing experiments to measure its performance.

ATM technology is supported by the ATM Forum [9], a consortium of over 500 private companies and universities. The development trends in ATM technology are moving toward a universal standard. The research trends are in ATM adaptable layer algorithm development, which will allow an optimal method for high-speed transmission of multimedia radiologic information including images, reports, video, voice, and text.

The list price of the ATM switch or the ATM LAN access switch is equivalent to that of an Ethernet router, but the ATM adapter board is lower than the fiber-distributed data interface board [1]. We anticipate that the price of ATM hardware components will drop drastically in the next 2 years and will then be affordable by most radiology departments. The ATM LAN has no cost besides the switches and adapter boards if the fiber-optic cables are already in place. On the other hand, the hidden cost of the ATM WAN is the long-distance connections that depend on the charges posted by the carriers who own the transcontinental and transoceanic fiber-optic cables. It is imperative that the radiologic community demonstrate the usefulness of this technology in health-care delivery so that the cost of using these cables will come down. The beauty of the ATM technology is that even if the ATM WAN is still too high to be cost-effective in the near future, we can still use the T1 line for the WAN, as shown in Figure 3, and the ATM LAN architecture remains functionally intact. We expect ATM to become the standard communication technology in the near future.

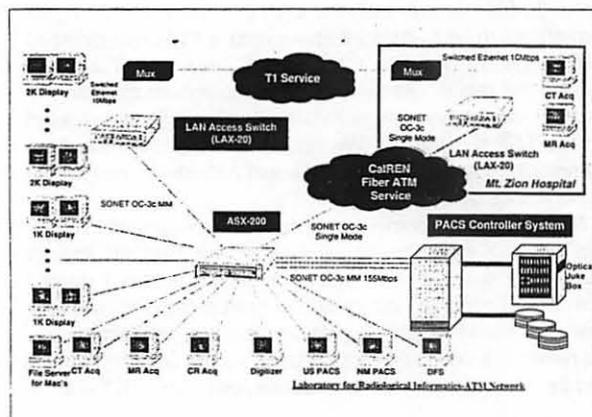


Fig. 3.—Implementation of first phase of asynchronous transfer mode (ATM) wide area network (WAN) and local area network (LAN) configuration at University of California, San Francisco (UCSF). WAN is connected by ATM switch ASX-200 at UCSF and LAN access switch LAX-20 at Mount Zion Hospital through main ATM switch in Oakland, CA. In this phase, LAX-20 is chosen at Mount Zion Hospital because most computer equipment there still uses older computer bus architecture. However, connection between two switches at Mount Zion Hospital and UCSF is true ATM speed. In this architecture, T1 line is used as WAN back-up. PACS = picture archiving and communication system, Acq = acquisition device, SONET = Synchronous Optical Network, CalREN = California Research and Education Network, Mux = multiplexer, OC-3c = optical carrier level 3, US = ultrasound, NM = nuclear medicine, DFS = digital fluoroscopy, Mbps = megabits/sec.

## REFERENCES

- Huang HK, Lou SL, Cho PS, et al. Radiologic image communication methods. *AJR* 1990;155:183-186
- Stewart BK, Lou SL, Wong WK, Huang HK. An ultrafast network for communication of radiologic images. *AJR* 1991;156:835-839
- Huang HK, Wong WK, Lou SL, Stewart BK. Architecture of a comprehensive radiologic imaging network. *IEEE J Sel Areas Commun* 1992;10:1188-1196
- Stewart BK, Dwyer SJ, Huang HK, Kangaroo H. Design of a high speed high resolution teleradiology system. *PACS VI, SPIE* 1992;1654:66-80
- McDysan DE, Spohn DL. *ATM theory and application*. New York: McGraw-Hill, 1995
- Cavanaugh JD, Salo TJ. Internetworking with ATM WANs. In: Stallings W, ed. *Advances in local and metropolitan area networks*. New York: IEEE Computer Society Press, 1994
- Lou SL, Huang HK. Assessment of a neuroradiology PACS in the clinical environment. *AJR* 1992;159:1321-1327
- Lou SL, Huang HK, Bazzilli T, Gould RG, Dillon WP, Schomer BG. Inter-hospital CT image communication: T-1 line versus carrier service. *Radiology* 1994;193(P):331
- The ATM Forum. *ATM user-network interface specification, version 3.0*. New York: Prentice Hall, 1993

u e

R

S

F

N

A

Laboratory

for

Radiological

Informatics

9

5

Chicago, IL  
November 26 - December 1

## REFRESHER COURSES

---

### Course No. 126

#### PACS Implementation of the Future

Meryll M. Frost, Jr, BSEE, Moderator, *Gainesville, FL*

H. K. Huang, DSc, *San Francisco, CA*

Steven C. Horii, MD, *Philadelphia, PA*

Meryll M. Frost, Jr, BSEE, *Gainesville, FL*

The radiology department of the future will use PACS exclusively for the acquisition, storage, distribution, and display of all radiologic images. The acquisition of all clinical images will be digital, eliminating the need for film while providing superior diagnostic performance. The seamless integration of diagnostic images with other conventional information systems will promote the efficient and effective practice of medicine in the managed-care environment.

This course will describe the three basic requirements of PACS for the radiology department of the future. First, the global need to merge all clinical information will be presented, with a description of the methods and benefits of integrating medical information systems with imaging technology. Next, the acquisition, storage, and distribution of radiologic and other clinical images, according to the DICOM standard, will be described. Communication issues for image distribution such as network types and speed will also be addressed. Finally, the diagnostic viewing room of the future will be introduced. Since the effective presentation of diagnostic information is not limited just to images, the most effective methods for providing access to all clinical data will be presented.

**Learning objectives:** To better understand the impact PACS will have on the future of radiology, along with how this technology will be used in the daily practice of clinical medicine.

### Course No. 353

#### How to Use a PACS Workstation ("Hands-on" Workshop)

Ronald L. Arenson, MD, Moderator, *San Francisco, CA*

Steven C. Horii, MD, *Philadelphia, PA*

Harold L. Kundel, MD, *Philadelphia, PA*

David E. Avrin, MD, PhD, *San Francisco, CA*

PACS continue to hold promise for reducing our dependence on film for acquisition, storage, and display of medical images. This course will provide radiologists and other health care professionals with hands-on instruction by using state-of-the-art workstations for PACS. **Purpose:** Participants will be led through important concepts, including (1) basic image manipulation (window/level, scroll/zoom, and patient selection), (2) image compression, (3) 3D reconstruction, (4) image processing (edge enhancement), and (5) folder manager processes (pre-fetch, auto-sequence, and auto-window/level). Where appropriate, comparison normal images will be shown along with the modified images to demonstrate the differences. No more than three participants will share each workstation, allowing sufficient individual time to master the techniques presented.

# SCIENTIFIC SESSIONS

## Sunday Morning • Room N227 (M-1)

204 • 11:25 AM

### Impact of US PACS

M.J. Moskowitz, PhD, San Francisco, CA • H.K. Huang, DSc • P.W. Callen, MD • R.A. Filly, MD • R.G. Gould, ScD

**PURPOSE:** To determine the cost-effectiveness of use of a mini-PACS system as the primary means of clinical reading and storage of US images. **MATERIALS AND METHODS:** A commercial mini-PACS module has been in operation in the US section of our academic institution for more than a year and connects 7 US units to 2 display stations. All studies are now read with soft copies, and images are archived with optical disks. Case-review time, retrieval time, and storage costs have been tabulated for a 2-month period before and after installation.

**RESULTS:** Cost savings were smaller than expected due to the large amount of data per US case and the expense of the optical storage medium. An increase in the average number of images acquired per case subsequent to the use of the PACS contributed to the storage costs. Initial results indicate that file-retrieval access times approach 30 seconds, on average.

**CONCLUSION:** US PACS is very efficient for clinical use. The departmental savings versus system costs are not as easily determined.

## Monday Afternoon • Room E353C

489 • 3:18 PM

### Coregistration of MR Images and High-Resolution PET Images for Quantitative Localization of Temporal Lobe Epilepsy

S.T. Wong, PhD, San Francisco, CA • R.C. Knowlton, MD • R.A. Hawkins, MD, PhD

**PURPOSE:** To optimize PET with fluorodeoxyglucose (FDG) and MR imaging localization of interictal seizure foci in partial complex epilepsy (PCE) patients by combining quantitative structural and functional information with qualitative visual image interpretation.

**MATERIALS AND METHODS:** High-resolution FDG PET images were acquired with a Siemens HR Exact system and were registered to MR images (Siemens 1.5 T) in 22 surgical candidates with PCE. PET and MR imaging data sets were coregistered with the Woods algorithm. Hippocampal and amygdala regions of interest (ROIs) were defined on MR images via interactive segmentation. Volumes were calculated from segmented MR imaging regions, and F-18 signal from corresponding ROIs on PET images was measured, normalized to hemispheric and whole-brain values. Visual inspection of coregistered images and quantitative MR image volume and metabolic FDG PET indices were related to other clinical data. **RESULTS:** Focal temporal zones of decreased FDG uptake were found in 13 PCE patients with atrophy and 9 patients without atrophy evident at MR imaging. Subtle abnormalities in 7 of 22 data sets were better detected with inspection of coregistered images.

**CONCLUSION:** Interpretation of FDG PET and MR images in nonlesional PCE patients is improved with image coregistration analysis. Subtle abnormalities may be best recognized with visual inspection, but quantitative analysis may produce generalizable thresholds for lateralization.

## Monday Afternoon • Room N139 (L-2)

581 • 2:30 PM

### Multimodal Image Fusion Techniques for Minimum Invasive Surgical Planning in Complex Seizures

S.T. Wong, PhD, San Francisco, CA

**PURPOSE:** To develop improved noninvasive imaging methods for the accurate localization of seizure foci, replacing invasive intracranial surgical planning methods currently in use and improving the outcome of the surgery for treatment of seizures.

**MATERIALS AND METHODS:** Selected patients referred to our Epilepsy Center entered an evaluation protocol consisting of 4 noninvasive imaging tests: MR imaging, F-18 fluorodeoxyglucose PET, H-1 MR spectroscopy, and magnetoencephalography, as well as intracranial electroencephalography (EEG). Multimodal image fusion analysis was applied to coregister, extract, and quantitate the relevant structural and functional

information of these 4 neuroimaging modalities. Results were compared with those from the established surgical planning method (ie, intracranial EEG studies) and with surgical outcomes.

**RESULTS:** Initial results of 22 selected epilepsy patients who underwent surgery treatment demonstrated concordance between computed image fusion analysis and the intracranial EEG findings. All patients were seizure free or improved significantly after surgery. In addition, 2 surgical candidates missed with intracranial EEG were identified by the new image fusion method.

**CONCLUSION:** The combination of anatomic, biochemical, and physiologic imaging maps of ictal foci should lead to more effective, economic, and safe preoperative evaluation for seizure surgery. The image fusion methodology should be extendible for studying other neurologic disorders.

585 • 3:02 PM

### Distributed Image Processing with Asynchronous Transfer Mode Technology

X. Zhu, PhD, San Francisco, CA • H.K. Huang, DSc

**PURPOSE:** To use asynchronous transfer mode (ATM) networked computers to perform computationally intensive image processing tasks in time-critical radiology applications. With ATM technology (OC-3, 155 Mbit/sec), the network communication rate is greatly enhanced over conventional Ethernet technology, thus removing 1 of the major bottlenecks limiting the performance of the distributed computing.

**MATERIALS AND METHODS:** Time-consuming computed radiographic (CR) image preprocessing and enhancement requires convolving of a CR image with a mask of various dimensions. Instead of relying on proprietary hardware, we distributed the task to a computer cluster connected by a local-area ATM network via software developed on top of Parallel Virtual Machine, a distributed computing software package. Processing times were determined as a function of the number of clustered computers.

**RESULTS:** Image processing time was reduced inversely with the number of computers used. With improved bandwidth of the ATM network, more computers can be included to participate in distributed image processing.

**CONCLUSION:** With readily available computers in the radiology department, distributed computing provides a cost-effective way to shorten the image processing time. As the emerging ATM technology becomes more widespread, distributed computing can find wider applications in radiologic image processing.

## Tuesday Morning • Room N133 (L-8)

787 • 10:30 AM

### DICOM-compliant Teleradiology Server for a PACS Environment

D.E. Avrin, MD, PhD, San Francisco, CA • K.P. Andriole, PhD • A.C. Zintgraff

**PURPOSE:** To present a DICOM-compliant teleradiology client-server technology as an alternative to the traditional video frame-grab method.

**MATERIALS AND METHODS:** The DICOM client-server technology was connected to our large-scale PACS. Technique-modifier messages passed through our radiology information system (IDXRad) were read through an HL-7 interface to the PACS controller. The PACS controller routes specific examinations to the on-call radiologist through the teleradiology server. The DICOM server dials up and transmits the full 12-bit image data and relevant information to the remote client station, which is capable of receiving 12-bit data. The client station has window and level tools to optimize image viewing. The DICOM server also supports 8-bit viewers, allowing for a transitional, mixed environment.

**RESULTS:** A significant improvement in teleradiology image quality can be realized by transmission of the original, full-resolution (both spatial and density) 12-bit data to the remote viewing station. Window and level capabilities provide improved image visualization.

**CONCLUSION:** A DICOM compliant server-client system may be preferable to the video frame-grab method for on-call teleradiology applications, particularly in an institutional environment that already has a clinical PACS for digital modalities. (This work was supported in part by EMED, Inc.)

**789 • 10:30 AM****Feature-based Neuroimaging Databases**

S.T. Wong, PhD, San Francisco, CA • K. Soo Hoo, BSc • H.K. Huang, DSc

**PURPOSE:** A framework is lacking for integrating medical images, functional data, and reports into a general database environment. This work presents methodology that enables access to digitally stored multimodal medical images by features rather than by artificial keys, such as patient ID.

**MATERIALS AND METHODS:** A Unix workstation was used to retrieve neuroimaging image and text data directly from our departmental PACS. Feature extraction was carried out both automatically and via interactive segmentation techniques with use of an image processing package, VIDA. The quantitated feature data and original and annotated images were then organized in a remote multimedia database server. A multimedia graphic user-interface software package, Gain Momentum, was then used by the physician to access the remote database by image features.

**RESULTS:** An experimental neuroimaging database system, including MR imaging, MR spectroscopy, PET, and magnetic source imaging data in 25 patients, has been developed and used for aiding in seizure diagnosis. Image features used for indexing include lobar and fine brain anatomic volumes and associated functional data, such as glucose uptake, bioenergetic metabolite intensity, and dipole localization.

**791 • 10:30 AM****DICOM Image and Message Transmission over an Asynchronous Transfer Mode**

M.J. Moskowitz, PhD • H.K. Huang, DSc • T.M. Bazzill, BS • X. Zhu, PhD, San Francisco, CA

**PURPOSE:** DICOM image transmission is typically implemented by using TCP/IP over Ethernet. The purpose of our study was to quantify network performance and to implement strategies intended to increase overall performance.

**MATERIALS AND METHODS:** Central Test Node software was modified to incorporate TCP/IP calls on top of the higher-bandwidth ATM network. Our ATM network is both intrahospital and interhospital and connects acquisition computers, linked to digital imaging equipment, to a central archiving node and to display stations. Interhospital image transmission and message parsing times were evaluated by using Ethernet, T-1, and ATM technologies.

**RESULTS:** Initial results indicate that DICOM-ATM performance, while limited due to the TCP/IP overhead, is significantly faster than T-1 and Ethernet, roughly corresponding to their respective inherent bandwidth differences. Overall network performance is improved by use of a modified DICOM protocol for increased throughput over ATM, compared with the unmodified DICOM over ATM.

**CONCLUSION:** DICOM transmission over high-bandwidth mediums (T1, ATM, etc) is effective to employ, in terms of the network performance gained.

**794 • 10:30 AM****Improved Reading Efficiency by Using Montage in a Clinical PACS Display**

M.A. Linnolt, BS, San Francisco, CA • S. Lou, PhD

**PURPOSE:** With the expanding use of CT and MR and increased number of images per study, a more efficient soft-copy display format needs to be implemented. This paper presents our experience with a new approach of selecting only the significant images from a study into a montage file.

**MATERIALS AND METHODS:** Neuroradiology was chosen as the clinical setting in which we evaluated this new approach. The neuroradiologic 2K workstation, a Sun 4/470 with two 2K monitors and a parallel transfer array disk, was outfitted with the new montage software module. After appropriate clinical training sessions, the neuroradiologists were allowed to use the new function. This function creates user-selected montages containing only the relevant images needed for reading a case. Usage of this function was monitored over a 6-month period with a separate statistical software module.

**RESULTS:** The neuroradiologic service uses the PACS 2K workstation to read over 20 cases per day. The usage statistics collected show significant improvement in reading efficiency. The radiologists' response has been overwhelmingly laudatory.

**CONCLUSION:** Selecting a small subset of clinically relevant images from a study into a montage dramatically improves case display on a PACS workstation, cuts radiologist's reading time, and improves service efficiency. This approach should be adopted for all clinical PACS operations for CT and MR modalities.

**795 • 10:30 AM****Workstation in the Intensive Care Unit for the Assessment of Respiratory Distress**

E. Pietka, PhD, San Francisco, CA • O. Ratib, MD, PhD • J.C. Chevrolet, MD

**PURPOSE:** In acute respiratory distress (RD) an abnormal accumulation of water and solute in the extravascular space is a common finding. A good correlation has been reported between radiographic findings and severity of RD as measured with invasive monitoring techniques. Computer-assisted extraction of chest radiographic findings should provide a more objective and quantitative evaluation of anatomic indexes.

**MATERIALS AND METHODS:** Digital images have been measured on a specially designed workstation. Clinically approved measures are extracted, including cardiothoracic index, vascular pedicle width, pulmonary blood flow, and edema distribution. The presence of pleural effusions, peribronchial cuffs, septal lines, and air bronchograms is assessed.

**RESULTS:** The analysis is performed in 15 intensive care unit patients with RD and 15 patients with normal cardiorespiratory function. For radiographic examination, computed radiography is used. The RD patients are being longitudinally followed up. Visual comparison of images and assigned scores are replaced by real values reflecting the size, length, or width of predefined findings.

**CONCLUSION:** The semiautomatic measurements are more reproducible than conventional manual measures. A computer-aided analysis permits a better assessment of the severity of RD. Comparing a series of images enables a more accurate evaluation of the results of treatment.

**Tuesday Morning • Room N137 (L-6)****797 • 10:30 AM****Limitations of Transmission Control Protocol on High-Speed Radiologic Asynchronous Transfer Mode Networks**

S.T. Wong, PhD, San Francisco, CA • H.K. Huang, DSc

**PURPOSE:** To point out certain limitations of running popular transmission control protocols (TCPs) over asynchronous transfer mode (ATM) networks and to highlight the important issues for effective use of ATM technology in radiology.

**MATERIALS AND METHODS:** A wide-area OC-3 ATM network test bed (155 Mbit/sec) is being implemented in our department, connecting the 4 University of California, San Francisco (UCSF) hospitals in the San Francisco Bay area via a Pacific Bell ATM public switch and Fore System ASX200 local switches. The ATM traffic supported is best-effort service class. TCP is the communication method for image data transfer over these ATM links. Data throughput tests obtained the median bandwidth available to the TCP connection. Subsequent analysis on the cause of low transmission performance is thus conducted.

**RESULT:** Preliminary testing of TCP connections for memory-to-memory image data transfer between 2 UCSF hospital sites achieved approximately 65 Mbit/sec, less than 50% of the specification, after fine tuning TCP window and packet sizes. Further performance improvement is not expected in the best-effort service class of the ATM network.

**CONCLUSION:** To maximize the use of large ATM bandwidth, the slow start transmission algorithm used in TCP must be changed, and the buffer capacity of available ATM switches must be increased.

**798 • 10:38 AM****Authentication and Security Mechanisms for Digital Images**

S.T. Wong, PhD, San Francisco, CA • H.K. Huang, DSc • R.G. Gould, ScD

**PURPOSE:** An unresolved issue of medical imaging infrastructure is the authenticity and security of diagnostic images accessible over networks. This study investigates a practical solution based on rigorous cryptographic methods.

**MATERIALS AND METHODS:** An experimental cryptographic test bed was developed on top of our departmental PACS, with digital encryption and time stamping methods. Digital encryption provides data security and prevents unauthorized alternation and reading of image documents. Digital time stamping assures image authenticity, generates a characteristic "digital fingerprint" for a document when it is originally generated, and checks that it has not been modified. The appropriateness of these 2 cryptographic techniques for diagnostic images is considered, and encryption/decryption/transmission times are determined.

**RESULTS:** A model was derived for a digital image trust center to assure security and authenticity in diagnostic information. The cryptographic protocol used in the test bed demonstrated a trade-off between timing performance and security measures.

**CONCLUSION:** A practical solution to address data authenticity and security of medical image data bases is proposed. The use of cryptographic techniques involves a trade-off between security and speed.

801 • 11:02 AM

**Adaptive Process Control: A Case Study on Computed Radiographic Preprocessing for PACS Archive and Display**

J.G. Zhang, PhD, San Francisco, CA • S.T. Wong, PhD • K.P. Andriole, PhD • H.K. Huang, DSc

**PURPOSE:** Computational defects and software and operator errors can dominate the cause of failures in processing images acquired from various imagers into standardized formats for storage and display. Adaptive processing control allows computed image processing systems to adapt to changes in their environment by constantly measuring their own performance and adjusting their own parameters. The implementation of this technique for computed radiography (CR) image preprocessing is presented.

**MATERIALS AND METHODS:** Our radiology department operates 2 CR systems (Fuji 9000 and AC-2), generating about 1.6 gigabits of image data daily. These data must be preprocessed before archiving in our PACS, including reformatting, background removal, rotation alignment, lookup table correction, and image encoding. An adaptive process control program is developed to monitor the status of these tasks in real time and to correct system errors and computational defects automatically based on event-driving and task-backtracking mechanisms.

**RESULTS:** Real-time monitoring and error recovery of CR image preprocessing is implemented. Operator errors or system faults are detected and corrected immediately.

**CONCLUSION:** Adaptive process control techniques can improve the reliability of CR image acquisition, preprocessing, communication, and administration. They can be extended to other computed image processing.

802 • 11:10 AM

**Performance Characteristics of a Digital Imaging Network with Asynchronous Transfer Mode Technology**

A.W. Wong, BS, San Francisco, CA • H.K. Huang, DSc • T.M. Bazzill, BS • J.K. Lee, MS • I. Bobkov, BS • M.J. Maskowitz, PhD

**PURPOSE:** Asynchronous transfer mode (ATM) technology (OC-3, 155 Mbit/sec) has been implemented within our radiology department's hospital-wide PACS as well as in a wide-area network (WAN) connecting affiliated hospitals. This paper describes the network performance observed in a clinical setting.

**MATERIALS AND METHODS:** The ATM network connects 4 major campus buildings and a remote hospital, providing intra- and interbuilding communication for radiologic images, including CT, MR, computed radiographic, US, and digitized screen-film images. The network links these modalities to a central archiving node and to display stations. Network performance was measured in terms of image transfer rates a 24 hours a day, 7 days a week.

**RESULTS:** Network performance varied with time (peak hour or off hour) and location (local area network [LAN] or WAN). The maximum memory-to-memory transfer rate achieved between nodes was 66.4 Mbit/sec, allowing a 20-Mbyte CT examination to be transmitted in less than 3 seconds.

**CONCLUSION:** Even with the constraint of slow disk input/output speed, the ATM network still outperformed the Ethernet by 180%. For this reason, such networks are a suitable choice for large-scale PACS involving both LAN and WAN.

**Wednesday Morning • Room N230 (M-6)**

1310 • 10:46 AM

**Impact of a PACS Display Station in a Medical-Surgical Intensive Care Unit**

K.P. Andriole, PhD, San Francisco, CA • M.L. Storta, MD • M.A. Linnoit, BS • G. Gamsu, MD • H.K. Huang, DSc

**PURPOSE:** To assess the changes in the behavior of health-care professionals caused by introduction of a PACS display station in the intensive care unit (ICU), by means of pre- and post-PACS evaluations.

**MATERIALS AND METHODS:** Pre-PACS measurements obtained for 290 ICU digital chest and abdominal images included the number of films per patient day; the number of referring clinician visits per patient day; and measurements of elapsed time from exposure to the time of radiologist-unit interaction, review by referring physician, and clinical action based on image information. Post-PACS measurements of the same parameters were obtained by means of a workstation function utilization tracking program.

**RESULTS:** Before introduction of the PACS station, the average number of films per patient day was  $1.3 \pm 0.6$ ; the average number physician visits to the radiology department,  $2.18 \pm 1.35$  per day, with the physician and a radiologist interacting only 27.4% of the time; the average elapsed time to radiologist-unit interaction, 1 hour 42 minutes  $\pm$  2 hours 3 minutes; the average time to physician review, 2 hours 32 minutes  $\pm$  1 hour 50 minutes; and the average time to action, 3 hours 21 minutes  $\pm$  3 hours 24 minutes. Preliminary post-PACS results suggest that the time to clinical action may be reduced with an ICU display station.

**CONCLUSION:** Introduction of a PACS display station in the ICU may reduce the number of examinations per patient day and modify the behavior of health-care professionals. (This work was supported in part by a donation from Fuji Medical Systems U.S.A., Inc.)

**Thursday Morning • Room N137 (L-6)**

1622 • 11:18 AM

**Real-Time Structure-Lossless Mammographic Image Compression**

J. Wang, MSc, San Francisco, CA • H.K. Huang, DSc

**PURPOSE:** High-resolution digitization of an 8 x 10-inch mammogram creates 10-40 Mbytes of data. A structure-lossless mammographic image compression technique, implemented in real time during digitization, has been developed to reduce the image size.

**MATERIALS AND METHODS:** A laser scanner was used for film digitization. During the digitization, data are initially stored in a 32K buffer that holds 4-8 lines of data, depending on the digitization matrix, then are dumped to a host computer. Software, running on the host computer, was developed to segment the data, discard the background, and implement compression by using lossless predictive coding. This process is completed between buffer dumps. The laser scanner's speed is normally 50-100 lines per second. Twenty mammograms were digitized by using a 50-mm sampling, and the compression ratio achieved was determined.

**RESULTS:** Real-time compression reduces mammogram image size by a factor of 3-4 with no loss of information.

**CONCLUSION:** The used real-time structure-lossless image compression does not increase digitization time, yet it reduces storage requirements and image transmission time.

**Friday Morning • Room N231 (M-4)**

1868 • 11:50 AM

**Clinical Evaluation of Computed versus Plain Radiography for Application to Neonatal Intensive Care Units**

K.P. Andriole, PhD, San Francisco, CA • R.C. Brasch, MD • C.A. Gooding, MD • R.G. Gould, ScD • H.K. Huang, DSc

**PURPOSE:** To investigate the clinical utility of computed radiography (CR) versus plain radiography for neonatal intensive care unit applications.

**MATERIALS AND METHODS:** The latest-version CR imaging plate (ST-V), with a scanning density of 10 pixels/mm, was used in a modified film cassette containing the CR imaging plate, a conventional screen and film, allowing simultaneous acquisition of CR and plain film images. For 100 portable neonatal chest and abdominal examinations, plain film was subjectively compared with CR hard- and soft-copy. Hard-copy images were masked to disguise the origin of the film and were presented in random order for scoring. Soft-copy images were displayed on a high-resolution 2K x 2K display station. Two experienced pediatric and 3 junior radiologists graded overall image quality on a scale of 1 (poor) to 5 (excellent), visualization of various structures in the chest (ie, lung parenchyma, pulmonary vasculature, spine detail, tubes/lines) and visualization of pathologic findings (ie, pulmonary interstitial emphysema, interstitial thickening, pneumothorax).

**RESULTS:** Preliminary results indicate that the image quality of both soft- and hard-copy CR images is comparable to that of plain film, with CR yielding excellent pulmonary vasculature visualization.

**CONCLUSION:** High-resolution CR may be a suitable alternative to screen-film radiography for bedside neonatal imaging.

## SCIENTIFIC EXHIBITS

---

### SPACE 074PH

#### Framework of Medical Imaging Database Building on a Hospital-integrated PACS

*S.T. Wong, PhD, San Francisco, CA • K. Soo Hoo, BSc • H.K. Huang, DSc*

The framework for integration of medical images and reports into a general database environment is lacking and is urgently needed. The purpose of this work is to present the new image data analysis tools and the top-down multimedia engineering approach in the development of medical imaging databases that facilitate information retrieval by means of image features. An experimental medical imaging database test bed is built on top of a hospital-integrated PACS and provides tools for image registration, image feature extraction and quantitation, multimedia data organization, content-based image retrieval, and volumetric image visualization. Current medical images include MR imaging, PET, MR spectroscopy, and magnetoencephalography. The method developed in this experimental test bed should be extendable to other medical imaging database applications. [See also scientific paper 789.]

### SPACE 345ED

#### Computer-aided Mammography Instruction

*M.J. Moskowitz, PhD, San Francisco, CA • J. Wang, MSc • S.T. Wong, PhD • H.K. Huang, ScD • J.D. Allen, MD • E.A. Sickles, MD*

This exhibit is intended to display the state of the art in high-resolution computer-aided instruction for mammography. It is expected that the hardware used in this project will become less expensive in 2-5 years, which will make this technology available for widespread use. A commercial system (CREDO, Vicom) for viewing digitized mammograms was used as the platform for creation of an interactive teaching file for mammography. Radiology residents were given an examination, then trained on a digital teaching file made up of 20 cases, and retested.

## InfoRAD EXHIBITS

---

### SPACE 9824

#### MoSearch: A Radiologist-friendly Tool for Finding-based Diagnostic Report and Image Retrieval

*M.R. Ramaswamy, MD, Galveston, TX • D.S. Patterson, MD • B.W. Goodacre, MD • E. VanSonnenberg, MD • L. Yin, BA • J.K. Lee, MS • et al*

The diagnostic report is an end product of any imaging study or therapeutic intervention performed in our radiology department. Currently, over 1 million such reports are stored within the department's RIS, detailing 5 years of clinical experience. Although the RIS facilitates patient-based access to these data, finding-based access is limited and time consuming. This exhibit demonstrates a user-friendly PC-based software package that allows the radiologist to conduct sophisticated real-time searches of diagnostic reports based on patient characteristics, examination modality and anatomy, and imaging findings and to easily review, refine, and output the results. A notable feature of this system is its use of synonym-matching and syntactic cues, allowing it to identify findings within the text of a diagnostic report much more accurately than can simpler "keyword" searches.

#### Learning Objectives:

- Understand the wealth of information contained within diagnostic reports.
- Realize the value of computerized tools in recovering this information, much of which may have been otherwise "lost."
- Identify ways in which this information may be used to facilitate the day-to-day practice of radiology.

### SPACE 9848

#### Asynchronous Transfer Mode-distributed PACS Server for Intensive Care Unit Application

*H.K. Huang, DSc, San Francisco, CA • A.W. Wong, BS • T.M. Bazzill, BS • K.P. Andriole, PhD • J.K. Lee, MS • J.G. Zhang, PhD*

This exhibit simulates the intensive care unit (ICU) operation at the University of California, San Francisco, where images and reports are routed automatically from the PACS controller to a distributed ICU server. The server consists of a Sun SPARC 20 with a 30-GB storage array and an asynchronous transfer mode (ATM) (155 Mbits/sec) switch. In this exhibit, the server is preloaded with ICU patients' report and computed radiography images in their original size. Two ICU stations will be at the exhibit, each with a SPARC 20, a DOME Md2/SUN board driving two 1,600 × 1,200 monitors, and an ATM board connected to the switch. The user initiates a designated ICU at 1 of the 2 stations and requests patient images and reports. The transmission rate from the server to the station is about 6 Mbytes/sec.

#### Learning Objectives:

- The user-friendly ICU workstation software.
- The transmission speeds of ATM.
- The I/O transfer rate of the RIAD disk.
- The architecture of the distributed server.

# Special Invited 1996 Issue of Computerized Medical Imaging & Graphics on Medical Image Databases

## Table of Contents

Editorial by STC Wong and H. K. Huang

Forward to the special issue by S Shiffman and EH Shortliffe

### Content-Based Image Retrieval

"I<sup>2</sup>Cnet: Content-Based Similarity Search in Geographically Distributed Repositories of Medical Images" by S Orphanoudakis, C Chronaki, and D Vamvaka

"Medical Image Collection Indexing: Shape-Based Retrieval Using KD-Trees" by GP Robinson, HD Tagare, JS Duncan, CC Jaffe

"Active Index For Content-Based Medical Image Retrieval" by SK Chang

"Integrating content-based retrieval in a medical image reference database" by G Bucci, S Cagnoni, and R De Dominicis

"Holographic Image Archive" by JI Khan and DYY Yun

### Database Systems

"A Client/Server System for Internet Access to Biomedical Text/Image Databanks" by GR Thoma, LR Long, and LE Berman

"A 3D Medical Image Database Management System" by M Arya, W Cody, C Faloutsos, J Richardson, and A Toga

"Design Methods and Architectural Issues of Integrated Medical Image Database Systems" by STC Wong and HK Huang

### Image Data Modeling

"Object Manifestations For Medical Imaging Applications" by NR Adam, RD Holowczak, and WS Li

"A Medical Image Object-Oriented Database With Image Processing And Automatic Reorganization Capabilities" by F Aubry, V Chameroy, and R Di Paola

### Temporal Image Databases

"A United Timeline Model and User Interface For Multimedia Medical Databases" by JDN Dionisio, AF Cardenas, RK Taira, DR Aberle, WW Chu, MF McNitt-Gray, J Goldin, and RB Lufkin

"Temporal Image Database Design for Outcome Analysis" by X Zhu, KN Lee, DL Levin, STC Wong, HK Huang, K Soo Hoo

## Editorial

The special issue on medical image databases witnesses a new, evolving development of biomedical imaging. Medical image database is inherently multidisciplinary. As Shiffman and Shortliffe have stated in their foreword of this special issue, the convergence of technologies of image management and textual data management is leading to the development of a new set of relationships between biomedical imaging and informatics communities. Shiffman and Shortliffe described the importance of image information management for efficient delivery of health care and the necessity and opportunities in sharing data and applications between imaging and clinical information systems. Before all these, however, efforts are required to stimulate collaborations, joint training, and cross publications between medical imaging and medical informatics communities. This special issue vividly exemplifies their observations. The invited articles represent a synergy among leading workers of the biomedical imaging, computer science, and clinical communities in developing this new kind of information systems.

Content-based indexing constitutes one of the most important challenge of medical image databases. Several research groups are addressing this issue under a variety of approaches. Orphanoudakis, Chronaki, and Vamvaka discussed a network of servers that can provide content based query services through a World Wide Web server. They proposed description types for the representation of image content and presented performance results of a description type implemented in their network. Robinson, Tagare, Duncan, and Jaffe of Yale proposed a similarity measure and an indexing mechanism for non-rigid comparison of image shape using KD-Tree algorithms. They have tested their shaped-based retrieval techniques on a database containing 85 cardiac MR images. SK Chang introduced the new concept of active index for content-based retrieval of medical images. The applications of the active index are image prefetching and similarity matching. Bucci and colleagues proposed a content-based search engine for tomographic image databases based on the Kahrnunen-Loève transform. They aim to integrate the search engine into a radiologic image databases developed in the earlier European Union's Telemed effort for references and education purposes. Khan and Yun from Hawaii presented a drastically different mechanism. Rather than relying on symbolic representation of image content to mediate the search, they incorporate the

concept of holographic associative memory into computation form to facilitate image retrieval by content.

At the systems level, Thoma, Long, and Berman described a client/server system prototype developed in the National Library of Medicine for accessing multimedia biomedical databanks through the ubiquitous Internet. The challenge is to provide access to both the images, i.e., x-rays, CT, and MR, and associated data by a global community for broad range of biomedical applications and research. Arya, Cody, Faloutsos, Richardson, and Toga described a prototype database system developed at IBM to query and visualize 3D spatial data, focusing on the mapping of structural and functional relationships of the brain. They list the requirement of the application, discuss the database design issues, and present timing results of their experiments. Wong and Huang describe their effort in designing and implementing vertically integrated medical image database systems. The approach they used is to integrate content-based indexing and knowledge base techniques within the picture archiving and communication systems to make the whole image data useful in medicine. They discussed research issues and provides examples from applications implemented in the hospital integrated PACS environment at the University of California, San Francisco. Adam, Holowczak, and Li proposed the notion of object manifestations of medical image objects in a heterogeneous client computing environments. They provided examples of how to apply constraints and rules to represent data within objects and to cater for variations in computer platforms, networks, and user preferences. Aubry, Chameroy, and Di Paola of INSERM presented a more formal specification for the logical design of medical image databases, drawing on their experience in the European MIMOSA project. They tested their model in two projects relating to medical image communication and image processing monitoring.

The ability to define and track temporal relation in images and related data is becoming an essential components of medical image databases. Cardenas and his UCLA colleagues in the departments of radiology and computer science presented a multimedia medical database model to support timeline-based presentation of information. A proof-of-concept prototype is developed for thoracic oncology and thermal tumor ablation therapy of the brain. Zhu, Kim, Wong, Soo Hoo, and Huang, on the other hand, described a different model of temporal image database system with the ability to quantitate and monitor the

progress of lung cancer of a patient under therapy treatment. Even for similar medical problems, it is interesting to observe that the former system took the perspective of data modeling while the latter stressed on the importance of medical image processing.

Medical image databases is a rapid growing field that has the potential to change the practice and research of biomedical imaging and to improve the efficiency of health care delivery in the coming decade. This special issue is an initial effort to bridge the gap between biomedical imaging and informatics communities in addressing this new development of medical information management. Finally, we would like to thank Professor Robert S Ledley, Editor-in-chief and Mrs. Blair V Mossman, Managing Editor, for providing us the opportunity to serve as Guest Editors of this issue.

Guest Editors  
Stephen Wong  
H.K. Huang

NEW FROM VCH PUBLISHERS

# PACS

## Picture Archiving and Communication Systems in Biomedical Imaging

**H.K. HUANG, D.Sc., FRCR (Hon.)**

*Director,  
Laboratory for Radiological Informatics,  
Professor and Vice-Chairman,  
Department of Radiology,  
University of California,  
San Francisco, CA*

**May 1996. Approximately 450 pages.  
Tables, figures, references, and index.  
ISBN 1-56081-685-6. Hardcover \$150.00**

**To Place Your Order,  
Please Use the Convenient Form  
on the Reserve Side.**

For special discounts on bulk orders  
for corporate or premium use, contact the  
VCH Special Sales Department at:

VCH Publishers, Inc.

220 East 23rd Street • New York, NY 10010-4606  
PHONE: 212-683-8333 • FAX: 212-481-0897



**VCH Publishers, Inc.**  
New York • Weinheim • Cambridge • Basel • Tokyo

Picture archiving and communication systems (PACS) is a concept perceived by the radiology and health-care communities as "the wave of the future." PACS consists of image acquisition devices, storage archiving units, display stations, computer processors, and data base management systems. These systems are currently being installed worldwide for special clinical applications and large-scale hospital use. Written by the recognized pioneer in PACS and medical imaging research, PACS provides a systematic approach to this technology which will serve as a foundation for future PACS installation, applications, and training.

PACS is the first book on this topic to be written by a single author with 15 years' worth of research experience in developing PACS. The author has implemented two large scale PACS now in clinical operation and much of the material in this title has been developed through his 15 years of research and field experiences, especially in regards to the "do's and don't" of PACS technology and installation. The title presents a comprehensive treatment of all radiologic acquisition devices including conventional x-ray, computed tomography (CT), ultrasound (US), magnetic resonance imaging (MRI), computed radiology (CR), and laser digitization. PACS contains the first chapter ever written on methods of interfacing hospital information systems, radiology information systems, and PACS. Additional chapters describe the planning and implementation of a digital radiology department, clinical experience with PACS in leading university hospitals, and an extensive chapter on image compression. This title contains material at the cutting edge of current archiving, communications, and biomedical imaging technology, much of which has yet to appear in journal publications.

PACS is organized such that a reader unfamiliar with these systems can skip the more advanced material and still grasp the overall concepts involving PACS technology; while the reader who has had more experience in these areas may delve into the clinical details underlying the applications. Although PACS was introduced by the radiology community, it can be applied to any scientific field that requires the management of information in the form of pictures.

**AUDIENCE:** Bioengineers; Radiologists; Optical Scientists; Biomedical Researchers; Lecturers in Informatics; Research, University, and Industrial Librarians; Computer Scientists; Electrical Engineers; Physicists; Health-Care Administrators, and Advanced Undergraduate and Graduate students in the preceding areas.

See Reverse Side for Table of Contents and Order Form ➔

# PACS: PICTURE ARCHIVING AND COMMUNICATION SYSTEMS IN BIOMEDICAL IMAGING

## Table of Contents

1. INTRODUCTION • Introduction • Organization of this Book • History of Picture Archiving and Communication Systems (PACS) • PACS Activities and Related Research in Various Countries • 2. DIGITAL RADIOLOGIC IMAGE FUNDAMENTALS • Terminology • Density Resolution, Spatial Resolution, and Signal to Noise Ratio • Test Objects and Patterns • Spatial Domain and Frequency Domain • The Fourier Transform (FT) • The Discrete Fourier Transform • The Fast Fourier Transform (FFT) • Performing a 2-D FFT on Radiologic Image • Cosine Transformation Using the FFT • Measurement of Image Quality • Measurement of Sharpness • Measurement of Noise • 3. IMAGE ACQUISITION I - CONVENTIONAL PROJECTION RADIOGRAPHY • Principles of Conventional Projection Radiology • Some Standard Procedures Used in Conventional Projection Radiology • Image Receptor • X-Ray Film Scanner • Video Scanning System • Drum Scanner • Solid State Camera • Laser Scanner • Digitization • Parameters Affecting the Quality of the Digitization • Aliasing Artifacts • 4. IMAGE ACQUISITION II - IMAGING PLATE AND DIGITAL RADIOLOGY • Laser-Stimulated Luminescence Phosphor Plate Technology • Principle of the Laser-Stimulated Phosphor Plate • A Computed Radiography System Block Diagram and its Principle of Operation • Description of System Components • Specifications of the Image Reader and Recorder • Operating Characteristics of the Scanning Laser-Stimulated Luminescence System • Some Clinical Applications • Digital Fluorography • Description of System Components • Operational Procedure • Imaging Mode • Processing Mode • Low-Scatter Digital Radiographic System • Line Scan Technique • Flying Spot Scan Technique • Dual-Energy Imaging • The Concept of Aluminum (AL) and Plastic (PL) Component Images • Theory of Dual-Energy Imaging and its Computational Procedure • The Linear Theory • 5. IMAGE ACQUISITION III - SECTIONAL IMAGING • Image Reconstruction from Projections • The Fourier Projection Theorem • The Algebraic Reconstruction Method • The Filtered (Convolution) Back-Projection Method • Transmission X-Ray Computerized Tomography (XCT) • Scanning Mode • Operation Principle of an XCT Scanner • Emission Computerized Tomography • Single Photon Emission CT (SPECT) • Positron Emission CT (PET) • Nuclear Medicine and Ultrasound Imaging • Nuclear Medicine • Ultrasound Scanning • Magnetic Resonance Imaging • Basic Physics • Pulsing Sequences and Relaxation • Magnetic Resonance Image Production • 6. IMAGE COMPRESSION AND RECONSTRUCTION • Terminology • Background • Error-Free Compression • Clipping and Bit Truncation • Run-Length Coding • Huffman Coding • Two-Dimensional Irreversible Image Compression • Introduction • Block Compression Techniques • Full-Frame Compression • Measurement of Differences Between the Original and the Reconstructed Image • Quantitative Parameters • Qualitative Measurement • Acceptable Compression Ratio • Quality of the Reconstruction Image • Receiver Operating Characteristics Analysis • Three-Dimensional Image Compression • Background • Wavelet Theory and Multiresolution Analysis • Three-Dimensional Image Compression with Wavelet Transform • Color Image Compression • Examples of Color Image in Radiology • The Color Spaces • Compression of Color US Images • Legal and Regulatory Issues • Use of Irreversible Compression • Measure of Image Compression • Image Post-Processing • Legal Standards of Compression • Product Liability • Summary and Research Directions • 7. PICTURE ARCHIVING AND COMMUNICATION SYSTEM (PACS) I:

INFRASTRUCTURE DESIGN AND IMAGE ACQUISITION • Introduction • PACS Components • Design Concept • Two Industry Data and Image Interface Standards • Health Level 7 (HL 7) Interface Standard • ACR/NEMA DICOM Interface Standard • Interface with Imaging Modalities and PACS Modules • Background • Automate Image Acquisition Interface Methods with Imaging Modalities • Fault Tolerance Methods for Automate Image Acquisition • Interface with a PACS Module • Imaging Preprocessing • 8. PICTURE ARCHIVING AND COMMUNICATION SYSTEM (PACS) II: COMMUNICATIONS AND DATABASES • Background in Communications • Terminology • Network Standards • Network Technology • Connecting Networks Together • Cable Plan • Types of Network Cables • The Hub Room • Cables for Input Sources • Cables for Image Distribution • Video Broadband Communication System • Broadband Technology • The Fiber-Optic Broadband Technology Video Communication System • Digital Communication Networks • Background • Design Criteria • Network Models • PACS Network Design • Teleradiology Network Design • PACS Central Node • Design Concept • System Configuration • System Software • System Operations • Interface with HIS and RIS • Background • HIS and RIS • Reasons for Interfacing PACS with HIS and RIS • Some Rules • Common Data in HIS, RIS, and PACS • Implementation of RIS-PACS Interfaces • Interface with Other Medical Databases • Multimedia Medical Data • Multimedia in the Radiology Environment • Integration of Heterogeneous Databases • 9. PICTURE ARCHIVING AND COMMUNICATION SYSTEMS (PACS) III: DISPLAY AND INFORMATION RETRIEVAL • Softcopy Display Workstation • Image Buffer and Processing Hardware • Display Monitor • Image Storage Device • Types of Display Workstations • Ergonomics of Image Workstations • Glare • Ambient Illumination • Acoustic Noise Due to Hardware • Image Processing and Display Functions • Image Enhancement Functions • Image Display and Measurement Functions • Montage • Basic Functions for a Display Workstation • The Laser Film Imager • The Block Diagram • Evaluation of the Performance Characteristics of a Laser Film Imager • Information Retrieval and Display • Patient Folder Manager Concept - Preliminary • Patient Folder Manager - Modules • Physician's Desk Top Retrieval • Distributed Image File Server • Some Issues Related to Image and Information Retrieval • 10. PICTURE ARCHIVING AND COMMUNICATION SYSTEM (PACS) IV: SYSTEM INTEGRATION AND IMPLEMENTATION STRATEGIES • System Implementation Strategies • Background • Three Methods of PACS Implementation • Method 1: System Integrations • Method 2: Requirements, Specifications and Contracting • Method 3: Turnkey Approach • Planning to Install a PACS • Film-Based Operation • Digital-Based Operation • PACS System Evaluation • Subsystem Throughput Analysis • System Efficiency Analysis • Image Quality Evaluation • 11. PICTURE ARCHIVE AND COMMUNICATION SYSTEM (PACS) V: CURRENT DEVELOPMENT TRENDS AND FUTURE RESEARCH DIRECTIONS • Recent Trends in PACS Development • Large Scale PACS • Teleradiology • ICU Modules • Emergency Room Module • PACS Future Research Directions • Medical Image Database Management • Architecture of the Medical Image Database Server (MDS) • Computation and 3-D Rendering Node • Image Content Indexing • Concept of Distributed Computing • Distributed Computing in a PACS Environment • Authenticity for PACS Images and Records • REFERENCES • INDEX

## ORDER FORM

 **VCH Publishers, Inc.**  
Customer Service Center  
303 Northwest 12th Avenue  
Deerfield Beach, FL 33442-1788  
New York • Weinheim • Cambridge • Basel • Tokyo

To order, simply fill out this convenient order form, enclose your payment, and mail to: VCH Publishers, Inc. at the above address.

FOR FASTER SERVICE IN THE U.S. AND CANADA, CALL TOLL-FREE 1-800-367-8249. OR FAX: 1-800-367-8247. OUTSIDE THE U.S. AND CANADA, CALL 1-305-428-5566. OR FAX: 1-305-428-8201.

VCH Now Takes Orders by E-mail! E-mail address: order@vch.com

To ensure prompt service, please include author/editor, title, ISBN, quantity, PO # or credit card #, billing address, shipping address, shipping address, and telephone # with your order.

YES! Please send me \_\_\_\_\_ copy(ies) of PACS: Picture Archiving and Communication Systems in Biomedical Imaging, by H.K. Huang.  
SBN 1-56081-685-6. Hardcover \$150.00

Subtotal: \$ \_\_\_\_\_  
Postage & Handling: \$4.00 for first book, \$1.00 for each additional: \$ \_\_\_\_\_  
Sales tax: Florida residents please add 6% sales tax. \_\_\_\_\_  
New York State residents please add appropriate local sales tax. \_\_\_\_\_  
Canada residents please add 7% GST. \$ \_\_\_\_\_  
Total: \$ \_\_\_\_\_

For orders outside of North America, please contact VCH Verlagsgesellschaft mbH, P.O. 10-11-61, 69451, Weinheim, FRG.

*100% Satisfaction Guarantee! If you are not completely satisfied, return your purchase in salable condition within 30 days for a full refund.*

### METHOD OF PAYMENT:

Check/Money Order enclosed (Payable in U.S. dollars to VCH Publishers, Inc., drawn on a bank in the United States).  
 Charge My Credit Card:  AMEX  Mastercard  VISA  
Account No. \_\_\_\_\_ Expiration Date \_\_\_\_\_  
Signature \_\_\_\_\_

(Credit card orders not valid without signature)

Bill Me  
Purchase Order No. \_\_\_\_\_

### SHIP TO: (please print)

Name: \_\_\_\_\_  
Title: \_\_\_\_\_  
Address:  Home  Business \_\_\_\_\_  
City \_\_\_\_\_ State \_\_\_\_\_ Zip \_\_\_\_\_  
Telephone (Daytime): \_\_\_\_\_

### LIBRARY RECOMMENDATION FORM

To the Acquisitions Librarian, Books: I recommend the library purchase a copy of the title(s) I have indicated. Order through your book distributor or by contacting VCH Publishers, Inc., Customer Service Center, 303 Northwest 12th Avenue, Deerfield Beach, FL 33442-1788.

From \_\_\_\_\_ Department \_\_\_\_\_  
Signature \_\_\_\_\_ Date \_\_\_\_\_

11.95 DH 10

*Price is subject to change without notice.*