Journal of Digital Imaging

VOL 9, NO 2 MAY 1996

Design and Implementation of a Picture Archiving and Communication System: The Second Time

H.K. Huang, Katherine Andriole, Todd Bazzill, S.L. Lou, Albert W.K. Wong, and Ronald L. Arenson

This report describes the authors' experience in the design and implementation of two large scale picture archiving and communication systems (PACS) during the past 10 years. The first system, which is in daily clinical operation was developed at University of California, Los Angeles from 1983 to 1992. The second system, which continues evolving, has been in development at University of California, San Francisco (UCSF) since 1992. The report highlights the differences between the two systems and points out the gradual change in the PACS design concept during the past 10 years from a closed architecture to an open hospital-integrated system. Both systems focus on system reliability and data integrity, with 24-hour on-line service and no loss of images. The major difference between the two systems is that the UCSF PACS infrastructure design is a completely open architecture and the system implementation uses more advanced technologies in computer software, digital communication, system interface, and stable industry standards. Such a PACS can withstand future technology changes without rendering the system obsolete, an essential criterion in any PACS design.

KEY WORDS: picture archiving and communication system (PACS), communication networks, infrastructure design, open architecture, display workstation.

THERE ARE generally three methods of approach to design and implementation of a picture archiving and communication system (PACS).[1] In the first approach, systems integration, a multidisciplinary team with technical know-how is assembled by the radiology department or the hospital. The team becomes a system integrator, selecting PACS components from various manufacturers. The team develops system interfaces and writes the PACS software according to the clinical requirements of the hospital. In the second approach, requirements specification and contracting, a team of experts, from both outside and inside the hospital, is assembled to write detailed specifications for the PACS for a certain clinical environment. A manufacturer is contracted to implement the system. In the third, or turnkey approach, the manufacturer develops a turnkey PACS and installs it in a department for clinical use. Each of these approaches has advantages and disadvantages. One advantage of the first, or systems integration approach, is that the research team can continuously upgrade the system with state-of-the-art components and therefore, the system will not become obsolete. The system so designed is tailored to the clinical environment and can be upgraded without depending on the schedule of the manufacturer. One disadvantage is that it requires a substantial commitment by the hospital to assemble a multidisciplinary team. In addition, the system developed will be one of a kind, and therefore, service and maintenance will be difficult because it consists of components from different manufacturers.

The primary advantage of the second approach (requirements specification and contracting) is that the PACS specifications are tailored to a certain clinical environment, yet the responsibility for implementing the PACS is delegated to the manufacturer. The department acts as a purchasing agent and does not have to be

From the Laboratory for Radiological Informatics, Department of Radiology, University of California, San Francisco 94143-0628.

Partially supported by PHS grant RO1 CA 51198, National Library of Medicine contract NO1-LM-4-3508, Cal REN grant ATMN-007, Fuji Medical Systems USA, Lumisys, and Abe Sekkei.

Address reprint requests to H.K. Huang, DSc, Laboratory for Radiological Informatics, Department of Radiology, University of California, San Francisco, CA 94143-0628.

R.L.A. contributes to the design of the UCSF PACS.

0897-1889/96/0902-0001$3.00/0

concerned with the installation. The disadvantages are that the specifications tend to be overly ambitious. The experts may not be familiar with the clinical environment, and they may underestimate the technical and operational difficulty. The designated manufacturer, who may lack clinical experience, tends to overestimate the performance of each PACS component. As a result, the completed PACS may not meet the overall specifications. The cost of contracting the manufacturer to develop a specified PACS is also high because of the manufacturer's narrow profit margin in building only one system.

The advantage of the third or turnkey approach is that it is a generalized production system, therefore, the cost tends to be lower. However, in this approach, the manufacturer needs a couple of years to complete the production cycle. By the time the system is commercially available, some components may have already become obsolete because of the fast moving computer and communication technologies. Also, it is doubtful whether a generalized PACS can be used by every specialty in a department and by every radiology department. In the past several years, these three approaches gradually merge as additional clinical data in PACS become available. As a result, the distinction among them has become blurred.

This report describes the authors' experience of design and implementation of two PACS, one at University of California at Los Angeles (UCLA) and the other at University of California at San Francisco (UCSF) during the past 10 years. The first system was based on the first approach and the second was based on combining the first and second approaches. Sections 2 and 3 describe the UCLA and the UCSF PACS, respectively. Section 4 compares the differences between these two systems.

THE FIRST PACS SYSTEM AT UCLA

We began the design of the UCLA PACS in 1983.[2] Its implementation went through three phases. Phase 1, from 1984 to 1990, encompassed the demonstration of the concept of PACS and the design of the PACS infrastructure. Phase 2, from 1990 to 1991, comprised clinical implementation of several PACS modules. Phase 3, from 1992 on, included the system's refinement, maintenance, and applications. This PACS was designed for the radiology department operation without consideration of a hospital integrated PACS.

Phase 1: Demonstration of Concept and Design of PACS Infrastructure

To show the concept of PACS to physicians, in 1987 we implemented two PACS modules; one in the pediatric radiology section within the department,[3] and the other in the coronary care unit.[4] The pediatric radiology section was selected because it operates independently from other radiology specialties and resembles a mini radiology department. It is an excellent model to study the implementation of a PACS for the entire radiology department. In this module, images were displayed on two 2,048-line monitors. The module was used for daily conferences and case reviews. The coronary care unit was chosen for the second PACS module because it explored the application of PACS outside of the radiology department. In this module, images were displayed on three 1,024-line monitors. Both modules were in clinical operation 24 hours a day, 7 days a week. The reactions from both radiologists and clinicians who used these two systems was very positive.

From 1988 to 1990,[5] we concentrated on the design of the PACS infrastructure. The critical components in the infrastructure were the communication system, PACS controllers, data base design, fault tolerance consideration, and system integration software. This infrastructure supported a digital-based radiology operation.

The infrastructure was implemented from 1990 to 1991.[6] There were 64 multimode and 48 single-mode fiber optic cables connecting the three buildings (Center for the Health Sciences [CHS], Medical Plaza, and Taper Building) housing the radiology department. There were two PACS controllers, one at the CHS and one at the Medical Plaza. The infrastructure was on-line in the beginning of 1991.

Communication system. We designed a three-tiered fiber optic communication system with Ethernet, FDDI (fiber distributed data interface), and Ultranet (a proprietary 1 gbit/sec network).[7] Ethernet was used to transmit images from acquisition devices to the acquisition computer. Because the acquisition device was slow in generating images, the transmission speed between these two nodes was not crucial. Images were reformatted at the acquisition computer and sent to the PACS controller by means of FDDI. Images were archived onto optical disks and distributed to the image display stations with the Ultranet. The three communication networks were coexistent in the infrastructure and served as backups for each other.

PACS controllers. There were two PACS controllers in the infrastructure.[8] Each controller was composed of an image server (4/490 SPARC; Sun Microsystems, Mountain View, CA) with 4-Gbyte magnetic disk storage, a 1-Tbyte optical disk library with write once read many (WORM) disks for archiving images, and a Sun 4/490 SPARC server running the data base (Sybase, Emeryville, CA) for patient directory and text information. The architecture of each controller was identical and could be used as the backup for the other. The PACS controllers were connected with the Ultranet. Images could be transmitted between the PACS controllers and display workstations at 4 to 8 Mbytes/sec.

Data base. Two identical Sybase data bases existed in each PACS controller and served as a mirrored system. Current patient image information was updated continuously on the data base of each controller.

Fault tolerance consideration. In the infrastructure, every critical component had a backup. There were two identical data bases one in each PACS controller. Each PACS controller was located in a separate building to avoid potential disaster. The three communication networks backed up each other, and all active fiber optic cables had spares. Each PACS controller was powered by an uninterruptable power supply with up to 20 minutes of uninterrupted power.

Systems integration software. The previously described components were integrated as the PACS infrastructure by means of an elaborate system software. The system software was written in C programming language and ran under the UNIX operating system.

Phase 2: Implementation of PACS Modules

To implement PACS modules in the clinical environment, two additional tasks were needed. The first task, completed in Phase 2, was to connect image acquisition devices to the PACS controller through the infrastructure. The second was to design and implement display workstations in the department and clinics in the third phase. In image acquisition, we connected three computed tomography (CT) and three magnetic resonance (MR) scanners with direct digital interfaces, as well as three computed radiography (CR) units and two film digitizers to the infrastructure.

Phase 3: Systems Refinement, Training, Maintenance, and Applications

Phase 3 was comprised in two stages.[9] Stage 1 was the development of display workstations and their clinical implementation. Stage 2 consisted of refining the PACS; upgrading the display workstation software; and establishing training, maintenance, and service.

Stage 1: Display stations and clinical implementation. In this stage, four stations, each with two 2,048-line monitors, were deployed in the pediatric radiology (two stations), neuroradiology, and genitourinary radiology section. Also, one laser imager printing station was installed as a hardcopy device. In addition, two three-monitor stations with 1K monitors were installed in the coronary care unit and pediatric intensive care unit (ICU). Figure 1 shows the UCLA PACS infrastructure and image acquisition and display stations as of October 1992.

Stage 2: Systems refinement and training, maintenance, and service. During clinical implementation, we set up procedures for training, system maintenance, and service. Three groups of personnel were trained. The first group was radiologists and clinicians to use the display stations. The second group included the PACS coordinator, technologists, and clerical personnel. This training was extensive and covered image quality assurance, updating the patient directory, and first-line troubleshooting. The third group was the PACS engineers. This training was most elaborate. It included all operational aspects of the PACS.

In September 1992, the authors transferred the responsibility of daily operation to a new PACS management team at UCLA, and we relocated to UCSF to develop a second generation PACS. This report only summarizes the

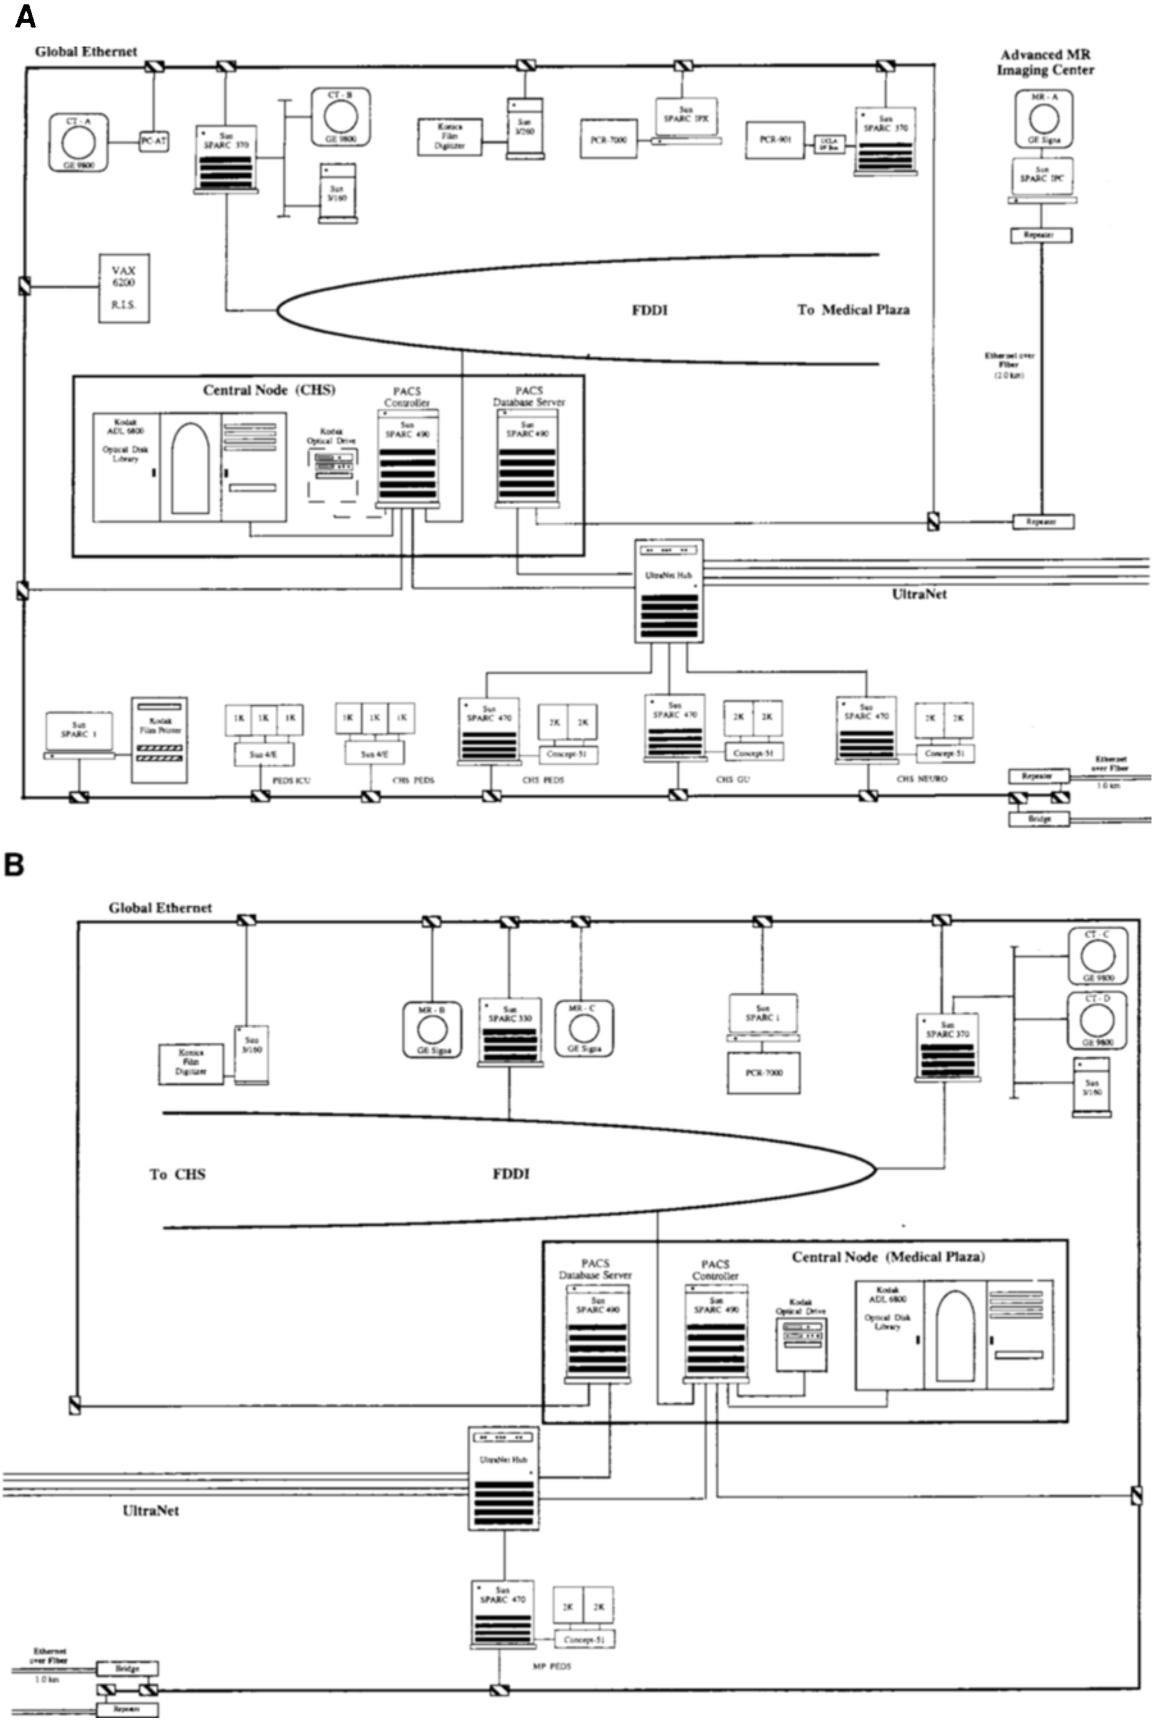
A
Global Ethernet
Advanced MR Imaging Center
CT - A
GE 9800
PC-AT
Sun SPARC 370
CT - B
GE 9800
Sun 3/160
Konica Film Digitizer
Sun 3/260
PCR-7000
Sun SPARC IPX
PCR-901
Sun SPARC 370
MR - A
GE Signa
Sun SPARC IPC
Repeater
VAX 6200
R.I.S.
FDDI
To Medical Plaza
Ethernet over Fiber (2.0 km)
Central Node (CHS)
Kodak ADL 6800
Optical Disk Library
Kodak Optical Drive
PACS Controller
Sun SPARC 490
PACS Database Server
Sun SPARC 490
Repeater
UltraNet Hub
UltraNet
Sun SPARC 1
Kodak Film Printer
1K
Sun 4/E
PEDS ICU
CHS PEDS
Sun SPARC 470
2K
Concept-51
CHS PEDS
CHS GU
CHS NEURO
Repeater
Bridge
Ethernet over Fiber
1.0 km
B
Global Ethernet
Konica Film Digitizer
Sun 3/160
MR - B
GE Signa
Sun SPARC 330
MR - C
GE Signa
Sun SPARC 1
PCR-7000
Sun SPARC 370
CT - C
GE 9800
CT - D
GE 9800
Sun 3/160
To CHS
FDDI
Central Node (Medical Plaza)
PACS Database Server
Sun SPARC 490
PACS Controller
Sun SPARC 490
Kodak Optical Drive
Kodak ADL 6800
Optical Disk Library
UltraNet Hub
UltraNet
Sun SPARC 470
2K
Concept-51
MP PEDS
Ethernet over Fiber
1.0 km
Bridge
Repeater

UCLA system up to September 1992. Further development of the UCLA system has since been done by the new management team.

THE SECOND TIME, UCSF

We started to plan the second generation PACS at UCSF in October 1992. In addition to following our previous PACS design philosophy at UCLA, we have redesigned the PACS as a hospital-integrated system,[10] and built the framework in the infrastructure for future PACS-based radiology research.[11]

There are several major differences between the UCSF and the previous UCLA PACS design, among them intelligent image archiving and distribution; integration of hospital information system (HIS), radiology information system (RIS),[12] and other manufacturer's PACS components; new network architecture and technology; and collaboration with manufacturers to develop new display workstations. This section describes these major features.

Intelligent Image Archive and Distribution

The UCLA PACS was designed with focus on system reliability and data integrity, promising 24-hour on-line service and no loss of images. Images were managed in the individual PACS component on a first come-first serve basis, which resulted in inefficient image distribution and retrieval.

The second generation UCSF PACS design includes more intelligent and thereby minimizes access time for both current and historical images. The system is hospital-integrated and based on a composite staging mechanism using multiple storage media, HIS and RIS, and the client server concept.

Two major aspects are 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 system and system efficiency, which minimizes access time for images at the display stations. The following describes some major components.

Local storage management via PACS intercomponent communication. To ensure data integrity, the UCSF PACS always retains two copies of an individual image on separate storage devices until a successful archive of the image to the long-term optical disk library has been made. This backup scheme is achieved via the PACS intercomponent communication:

(1) At the radiologic imaging system: Images are not deleted from the imaging devices unless technologists have verified the successful archiving of individual images via the PACS terminals. Should any failure of the acquisition process or the archival process occur, images can be resent from these imaging systems to the PACS; (2) At the acquisition subsystem: Images acquired in the acquisition subsystem remain on its local magnetic disks until the archive subsystem acknowledges back to the acquisition subsystem a successful archive. These images are then deleted from the magnetic disks so that storage space from these disks can be reclaimed; (3) At the archive subsystem: Images received in the archive server from various acquisition nodes are not deleted before their successful archiving to the optical storage. On the other hand, all archived images are stacked in the archive server's cache magnetic disks and will be deleted based on their aging criteria (eg, number of days an examination is performed, discharge or transfer of a patient, etc).

Multiple storage media. The storage management system features three levels of user-accessible storage media: (1) redundant array of inexpensive disks (RAID) in the display station for immediate access for current images; (2) magnetic disks in the archive server for fast retrieval of cached images; and (3) erasable magneto-optical disks and WORM disks in the optical disk library for retrieval of any historical images. On the other hand, all local magnetic disks in the radiologic imaging systems and the acquisition subsystem are used for storing newly acquired images. These images are deleted once they have been successfully archived to the optical disks. Table 1 illustrates the configuration of these multiple level storage media.

Fig 1. (A) UCLA PACS network at the Center for Health Sciences, and remote MR site. (B) UCLA PACS network at Medical Plaza. They are connected together. CHS, Center for the Health Sciences; GenUn, genitourinary radiology; PCR, Philips computed radiography; Peds, pediatric radiology; RIS, radiology information system (Reprinted from Computerized Medical Imaging & Graphics, Vol 17, Huang HK, Taira RK, Lou SL, et al, Implementation of a large scale picture archiving and communication system, pp 1-11, 1993, with permission from Elsevier Science Ltd, The Boulevard, Langford Lane, Kidlington OX5 1GB, UK. 9)

Table 1. Multiple Storage Devices for Images in the Storage Management System

	Storage Media	Location	Purpose
Level 1	Redundant array of inexpensive disks (temporary storage)	Display subsystem (display host)	Provides immediate access to both current and selected historical images
Level 2	Magnetic disks (temporary storage)	Archive subsystem (archive server)	Provides fast retrieval of current images
Level 3	Magneto-optical disks (longer-term storage)	Archive subsystem (optical disk library)	Provides retrieval of historical images
	WORM disks (permanent storage)	Archive subsystem (optical disk library)	Provides retrieval of historical images

RAID technology applied to PACS storage. All high-resolution (2,048 × 2,048 pixels) display stations in the UCSF PACS are configured with 5-Gbyte high-performance RAID. With this configuration, a 2,048 × 2,048 × 10-bit (8-Mbyte) CR image can be displayed in less than 2 seconds.

Folder manager. The storage management system is characterized by its on-line patient folder management.[13] 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, this folder remains in the display station(s) for immediate access until the patient is discharged, transferred, or other aging criterion (eg, two days after an out-patient visit) is met. The patient's admission, discharge, and transfer (ADT) 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 prefetching procedures will be performed. By applying the folder manager concept, the prefetch mechanism is only performed once per hospital stay of an individual patient.

Implementation of the Intelligent Archive Server

HIS/RIS/PACS interfacing. Interfacing the HIS allows the storage management system to receive patient ADT messages. Interfacing the RIS, on the other hand, allows the storage management system to receive information such as patient arrival, examination scheduling, examination cancellation, examination completion, etc. These events trigger the storage management system to perform the prefetch, studies grouping, and platter management mechanisms. Exchange of messages among these heterogeneous computer systems is conducted in the Health Level Seven (HL7) standard data format[14] with the use of Transmission Control Protocol/Internet Protocol (TCP/IP) protocols on a client/server basis.

Integration with other manufacturer's PACS components. The UCSF second generation PACS can integrate with other manufacturer's PACS components using the Digital Imaging and Communications in Medicine (DICOM) 3.0 standard. An example is the Aegis ultrasound PACS (Acuson, Mountain View, CA).[15] In this case, PACS treats the Aegis as a PACS acquisition device and coordinates the US images the same way as CT and MR images in the patient's image folder.

Image routing. Before successful archiving to long-term optical storage, all current images arrived at the archive server from various acquisition nodes are immediately routed to their destination display station(s). This routing mechanism minimizes access time for current images at the display stations. The routing process is driven by a predefined routing table composed of parameters including examination type, display station site, radiologist, and referring physician. The routing algorithm performs table look-up based on these parameters and determines where an image should go.

Image stacking. Stacking current images in the archive server's cache magnetic disks allows these images to be retrieved from the high-speed magnetic disks instead of the low-speed optical disks. The archive server holds as many images in its magnetic disks as possible and manages these images on the basis of their aging criteria. During a hospital stay, for example, images belonging to a given patient remain on the magnetic disks of the archive server until the patient is discharged or transferred.

Image aging. Aging criteria such as number of days since an examination was performed, discharge or transfer of a patient, or class of the patient (in-patient or out-patient) are used by the storage management system to control the migration of images from one storage device to another (eg, from a magneto-optical disk to a WORM disk) or the deletion of images from their resident storage devices.

Image prefetching. The prefetching mechanism[16] is triggered by means of a patient arrival message from the RIS. Selected historical images are retrieved from the long-term optical storage. These images are then distributed to the destination display station(s) before completion of the patient's current examination. Prefetching historical images to the display stations minimizes on-line image retrieval, hence relieving peak-hour workload of the archive subsystem and the networks. The prefetch algorithm is based on predefined parameters such as examination type, disease category, radiologist, referring physician, location of display station, and the number and age of the patient's archived images. These parameters determine which historical images should be retrieved from the long-term archive.

Studies grouping. During a hospital stay, a patient may undergo different examinations on different days. Images from these examinations are archived to the erasable magneto-optical disks, where they are scattered across different platters. When a patient is discharged or transferred, these images are then grouped from the magneto-optical disks and copied contiguously to a single WORM disk or to consecutive WORM disks for permanent storage. Once these images have been archived permanently, they are removed from the magneto-optical disks so that storage space in the magneto-optical disks can be reclaimed. Studies grouping allows all images from a patient during a hospital stay to be archived contiguously to optical disk(s), hence optimizing future retrieval of a patient's images from multiple examinations.

Platter management. Platter management allocates the storage space reserved in the WORM disks for future images in case a patient revisits or is readmitted to the hospital. In this way, images of a patient from multiple hospital visits can be accumulated in a single WORM disk or in consecutive disks, reducing excess disk swapping and consequently minimizing retrieval time for these images. However, preallocating storage space in an optical disk for a particular patient is expensive. Logically grouping consecutive optical disks into one volume, on the other hand, can reduce disk swapping time and hence minimizes the retrieval time for images stored in different disks within the same volume.

Networking

One distinct difference in networking between the UCSF and the UCLA PACS is the availability of asynchronous transfer mode (ATM) technology in the UCSF system. The UCLA networks were mainly a local area network (LAN) with Ethernet, FDDI, and Ultranet. In the UCSF PACS networks, ATM is used both in wide area network (WAN) and LAN with the conventional T-1 and Ethernet as back-up, respectively.[17] Figure 2 shows the logical network connection and Fig 3 shows the physical ATM connection.

Image Display

The UCSF PACS image display system 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 (Apple Computers, Cupertino, CA). In the first type, we modified the two-monitor 2K display workstations developed at UCLA by adding the HIS/RIS interface and some extra display functions. An example is the Montage function that allows the assembly of images from different examinations into one file. These workstations are used in the neuroradiology and pediatric radiology sections. Second, we worked

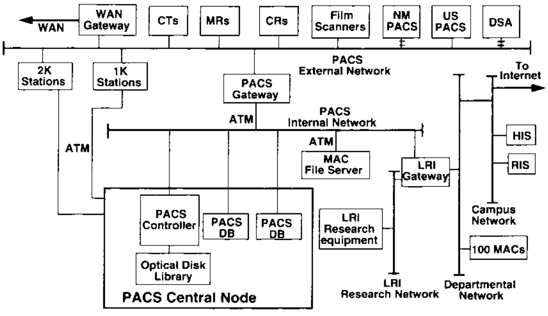

Fig 2. Department of Radiology, UCSF Logical Network Architecture, which includes WAN and LAN. The external network is open whereas the internal network has a firewall protection.

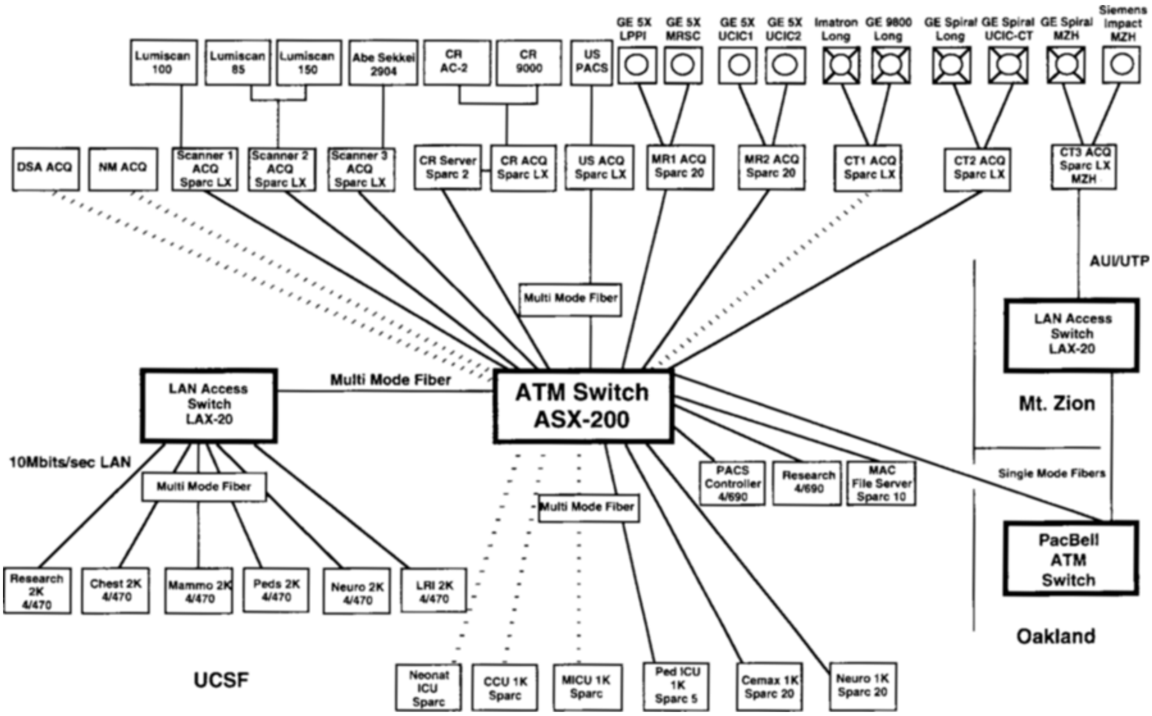

Fig 3. First phase ATM connection between UCSF and Mt Zion Hospital; (—) completed, (------) second phase. WAN and LAN ATM (155 mbits/sec) Network Physical Connection at UCSF. For those older components that do not support ATM, an ATM to LAN switch is used which supports 10 mbits/sec for every node connected to the switch.

with ISG Technologies Inc (Toronto, Ontario, Canada) to develop the two-monitor 1,600-line display stations for ICU applications. Figure 4 shows such a workstation.

The third method is to develop a file server to distribute integrated PACS images and textual data to the Macintosh desk top computers for individual review, teaching, and research.[18] Figure 5 shows the distributed network and Fig 6 depicts a page on the Macintosh screen.

MAJOR DIFFERENCES BETWEEN THE UCSF AND THE UCLA PACS

This section summarizes some major differences between the UCSF and the UCLA PACS.

PACS Controller

The PACS controller is an intelligent machine that controls the flow of data within the entire PACS from acquisition (data input) to archiving (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 data base; (4) archiving images to optical disks; (5) routing images and HIS/RIS data to display workstations; (6) handling retrieval requests from display stations.

The PACS controller implemented in the UCLA PACS was a single-processor Sun SPARCserver 490 computer, from which a one terabyte optical disk library based on 14 inch platters was attached. A three-tiered communication network comprising Ethernet, FDDI, and the proprietary fiber-optic UltraNet network was used to provide independent paths for data transmission between the PACS controller and other PACS computers. Images acquired from radiologic imaging devices were transferred from various acquisition nodes to the PACS controller, where they were then routed to destination display stations and were archived chronologically on WORM optical disks. With its fault-tolerant design, the PACS controller focused on system reliability and data integrity, promising 24-hour on-line service and no loss of images.

The PACS controller developed at UCSF includes more intelligence and thereby minimizes access time for both current and previous imaging studies. The computer system is based

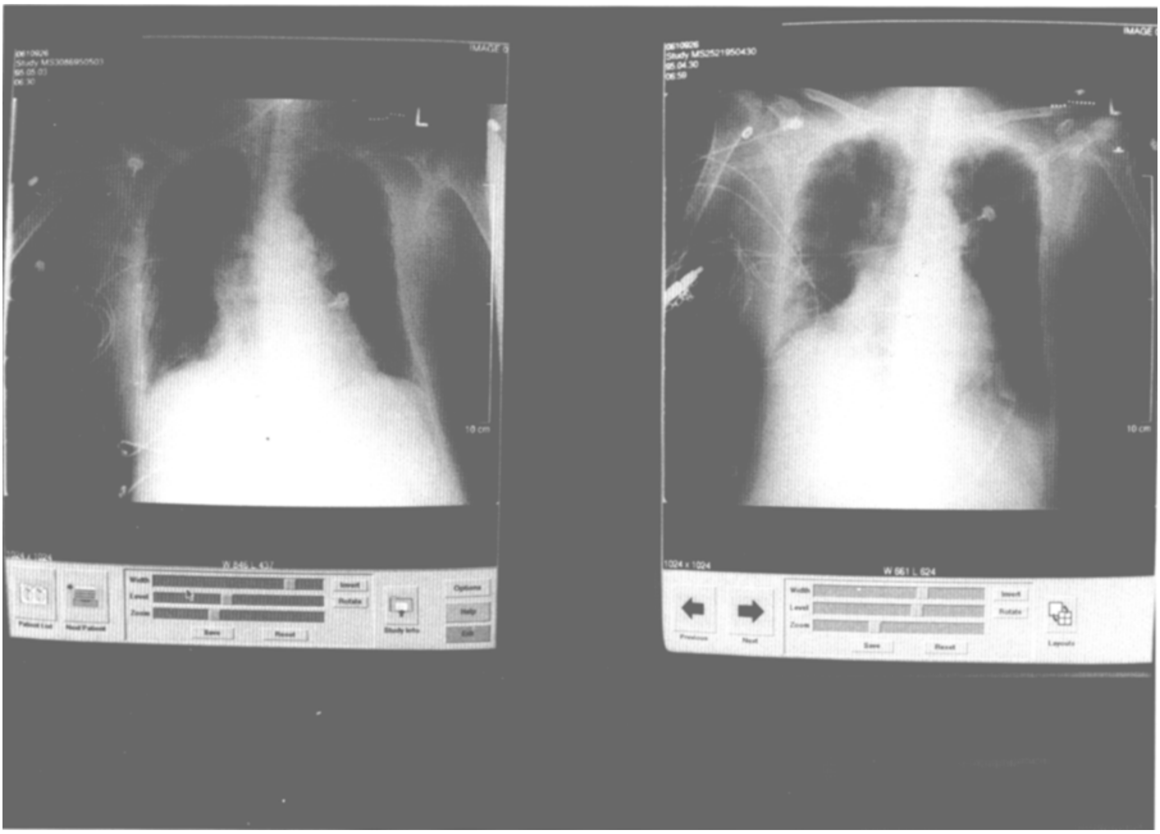

Fig 4. A 1,600-line two monitor workstation for ICUs. This station consists of a SPARC 20 with 128 Mbytes memory and two TurboGXplus cards, all off-the-shelf components. The display software was developed by ISG Technologies Inc based on UCSF specifications. The communication interface between workstations and the PACS controller was codeveloped by both parties based on the DICOM standard.

on the SUN 690 (Sun Microsystems) with four central processing units (CPUs), which allows multiple processes to run simultaneously with minimal shared CPU time. Two standard network interfaces, the Ethernet and the 155-bit/

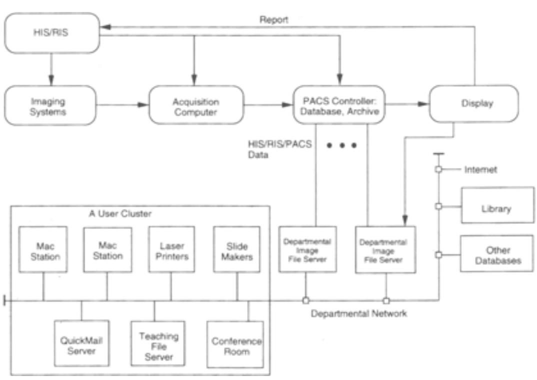

Fig 5. Distributed image file servers connected to the PACS controller. Each server provides specific applications for a given cluster of users. One of this clusters is to serve the Macintosh users in the department to access PACS data.

sec bandwidth ATM, OC 3 networks, are used for receiving and distributing images. The optical disk library 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). This second-generation PACS controller differs from the first-generation system in several of its new features: multiple storage media, image stacking, automated image prefetching, studies grouping, platter management, and HIS/RIS/PACS interfacing. Table 2 shows the major differences in the infrastructure design between the two systems.

In Table 2 there are several items that were not considered in the original UCLA design, for example, integration to the HIS/RIS and other

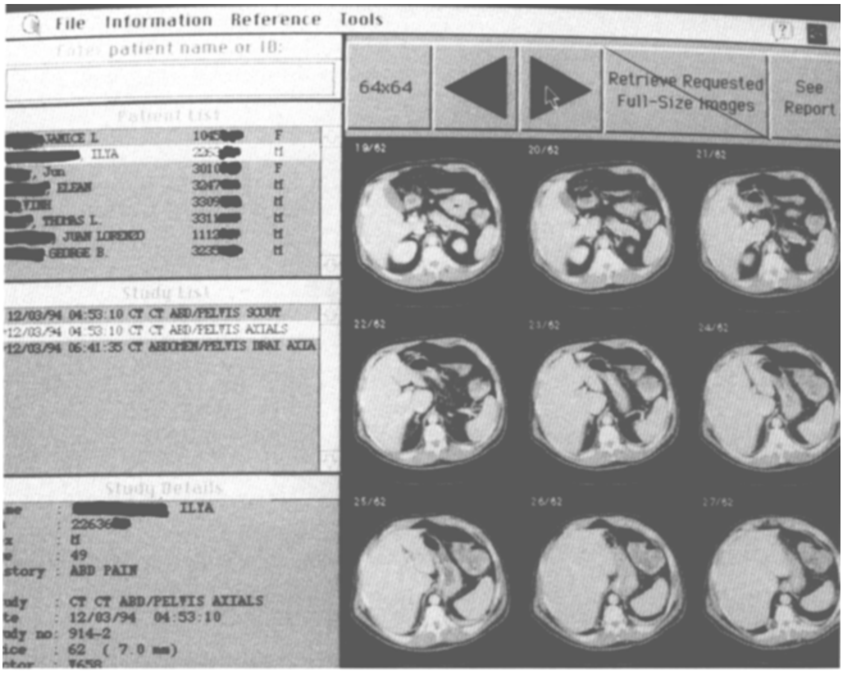

Fig 6. A page on the Macintosh screen. The user can directly access both images and textual information including diagnostic reports from PACS. The "Tools" kit allows user to manipulate individual images which includes programs like National Institutes of Health Image, etc.

vendor's PACS. The reason was because the hospital-integrated concept was not introduced until 1992. Also, the American College of Radiology-National Electrical Manufacturers Association (ACR-NEMA) and DICOM were not ready for implementation in 1991. On the other hand, the UCSF PACS does not have a duplicate PACS controller and all archiving is done within one site are because of limited space and financial resource.

Table 2. Differences in the Infrastructure Design

	UCLA (as of September 1992)	UCSF
Open architecture	limited	Y
Connectivity	limited	Y
Standardization		
ACR-NEMA, DICOM	N	Y
HL 7	N	Y
TCP/IP	Y	Y
Interface to HIS/RIS	N	Y
Interface to other vendor PACS modules	N	Y
Mirrored data base	Y	Y
Duplicate PACS control and archiving	Y	N
Different sites for archiving	Y	N
Auto routing	Y	Y
Image prefetching	N	Y
Image sequencing	N	in progress

Abbreviations: Y, Yes; N, No.

Networking

Table 3 shows the major differences in network design between these two systems. The UCLA network was mainly a LAN with very long fiber optic cable connections between buildings, whereas the UCSF system includes both WAN and LAN.

It is seen in Table 3 that the UCSF network design is more sophisticated with up-to-date network technology. On the other hand, the UCLA system had the higher speed proprietary Ultranet for high speed image communications. The UCLA system had no firewall protection for data integrity and security, which was a main draw-back in network design.

Image Acquisition Component

CT and MR. The UCLA CT/MR image acquisition systems were all made by General Electric Medical Systems (Milwaukee, WI). In this regard, the GE 9,800 CT scanners used GE IDNET-1 configuration for interfacing, and the image acquisition processes of GE high-speed CT scanners and Signa 5× MR scanners were based on GE proprietary communication protocols. The acquisition processes in the PACS did not have patient ID verification algorithm, no automatic process recovery mechanism, and no

Table 3. Differences Between the Two Networks

UCLA*	UCSF
• Fibers were used to connect remote buildings, but the backbone primarily consisted of Thicknet Ethernet (copper) that stretched through multiple buildings and levels. • Clinical acquisition, Research, and Distribution networks were all physically connected to the same subnet. Ultranet used for 2K distribution only. • AUI cables were connected directly to the backbone as opposed to CAT 5 UTP cables. Radiology Dept did not have a direct connection to the Internet. • Ultranet was used for transfer of images from the PACS Controller to the 2K display stations at a rate of 125 Mb/s. The transfer of information between the two Ultranet hubs in CHS and Medical Plaza was striped between four 250 Mb/s fibers giving a total transfer rate of 1,000 Mb/s or 1 Gb/s. • No routers or firewalls were used. • No WAN connections. • Only Med Plaza had a UPS.	• The backbone consists of primarily fiber optics, but CAT 5 UTP is the media used for station connections. • ATM via fiber used for acquisition and distribution. • The design is a distributed star configuration with multiple subnets. For examples, Genesis network is used to connect the digital modalities ie, CT, MR to the PACS controller. The departmental network is used to distribute images to Macintosh's users, etc. • PACS External network is used for transferring image data from the imaging modalities to the acquisition computers. The internal network, with a firewall protection is used to transfer data from the acquisition computers to the PACS Controller. • The External network is also used as a backup network for the display systems. • Distribution to the 1K stations is done via ATM, which uses 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 station. • Radiology has direct connections to the Internet via the departmental network which has over 100 connections to faculty and staff offices. • Routers and gateways are used to divide subnets and insure security among the networks. • The network has T1 and ATM WAN connections to two affiliated hospitals. Network hubs are backed up by UPS's.

*As of September 1992.
Abbreviations: AUI, attachment unit interface; CAT UTP, category 5 unshielded twist pair; UPS, uninterrupted power supply.

automatic paging system for engineering service. The image header format was based on a UCLA internal design and was not standard.

The UCSF CT/MR image acquisition systems consist of multivendor equipment including GE, Siemens (Erlangen, Germany), and Imatron (South San Francisco, CA). The GE 9800 CT scanner uses GE IDNET-II configuration for interfacing, and image acquisition processes of GE spiral CT scanners and Signa 5× 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. The PACS uses both ACR/NEMA 2.0 as well as DICOM 3.0 header information format.

CR. The UCLA PACS used earlier versions of CR technology: Philips 901 and two Philips 7000 laser plate readers (Philips Medical Systems, Shelton, CT), and ST-III type photostimulable phosphor plates; while the UCSF PACS incorporates the latest in CR technology: Fuji FCR AC2 and FCR 9,000 plate readers (Stamford, CT), and ST-V (standard) and HR-V (high-resolution) photostimulable phosphor plates. The interfacing of the CR devices to PACS required three different methods: two different methods for the two different models of Philips CRs and one method compatible with both Fuji CRs. The Philips 901 digital interface consisted of a RS-422 connection to a Data Interception Circuitry (buffer) board and a DR-11W link over which the image data was transmitted to the SUN acquisition computer; textual information (demographic and image header data) came over an RS-232 cable connection. The Philips 7,000 interface device included a PCR interface processor (PIP) which contained an auto data transfer board, memory buffer board and Ethernet board and transmitted both images and text to the SUN acquisition computer over Ethernet. The digital interface to PACS for both the Fuji AC2 and the 9,000 consists of a DMS Bus with a RS-485 cable (which is a combination 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 small computer system interface (SCSI) disk that transmits both textual and image data to the SUN acquisition computer over SCSI cable. Table 4 summarizes the differences in CR between these two PACS systems.

The Display Component

The UCLA PACS display component consisted of the 2K display stations, 1K display stations, and a film printing station. The UCSF display component is comprised of 2K display stations, 1K display stations, and Macintosh stations. Both 2K stations in the UCLA and UCSF PACS are of the same hardware platform. We developed the software in-house at UCLA, and carried it over to UCSF. Certain enhancements have been added to the software at UCSF based on users' feed-back. In particular, there is a local montage feature for users to select images from different files for display. A feature with 1 on 1 for current images and 4 on 1 for historical images was specifically 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 UCLA 1K system was developed based on an in-house designed display board which supported three 1,024 monitors. The UCSF 1K system is developed in collaboration with ISG Technologies Inc based on off-the-shelf components. This 1K system can support either 2, 3, or 4 1,600 line monitors, and has very easy to use user interface.

The UCSF system currently does not support a film printing station, instead, it supports over 100 Macintosh users to retrieve images and related PACS data for research, teaching, and case review. From these Macintosh computers, the users can select other printing resources in the department for hardcopy output. Table 5 summarizes the major differences in the display components between these two PAC systems.

Table 4. Differences Between the CR Acquisition Components

	UCLA*	UCSF
Devices	Philips 901 Philips 7,000 (two)	Fuji AC2 Fuji 9,000
Phosphor plate type	ST-III	ST-V, HR-V
Interface	DR-11W board to 901 PIP to 7,000	DASM to AC2 and 9,000

*As of September 1992.

Table 5. Major Differences in the Display Component Between the Two PAC Systems

Types of Workstation	UCLA*	UCSF
2K	SUN 4/470, Megascan display boards with two monitors (4)	Same with new enhancement software (4)
1K	SUN 4, in-house display board with three 1,024 monitors (2)	SPARC 20, Tubor GX+ boards supporting 2, 3, or 4 monitors (5)
Film printing	Y	N
PC station	N	100

*As of September 1996.

DISCUSSION

Both the UCLA and the UCSF PACS are designed in-house based on the SUN workstation platform. Because of the evolution of the PACS concept, the UCSF system is designed as a hospital-integrated PACS with the ACR-NEMA and DICOM and HL7 standards. Relevant data from HIS and RIS is automatically incorporated in the PACS. Other manufacturer's PACS components that 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, prefetching, and auto-sequencing have influenced our design in the UCSF PACS controller, which is quite a drastic difference from the UCLA 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 was not available in the UCLA system because of a lack of such knowledge on our part as well as not understanding the importance of data distribution to every radiologist and clinician at that time.

To our disappointment we have not seen any improvement in the 2K display technology in 5 years. Both the UCLA and UCSF 2K display systems are practically the same in terms of hardware platform and display software. On the other hand, 1K display has improved in both performance and cost.

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 additional acquisition modalities and display workstations. PACS system refinement will continue evolving for higher reliability, better performance, and new application functions as its development is a continuous dynamic process.

REFERENCES

1. Huang HK: Three methods of PACS research, development, and implementation. Radiographics 12:131-139, 1992

2. Huang HK, Mankovich NJ, Barbaric Z, et al: Design and implementation of multiple digital viewing stations. *Proceedings of the Society of Photo-Optical Instrumentation Engineers Second International Conference on Picture Archiving and Communication Systems for Medical Applications (PACS II)*, vol 418. 1983, pp 189-198

3. Taira RK, Mankovich NJ, Boechat MI, et al: Design and implementation of a picture archiving and communication system (PACS) for pediatric radiology. Am J Roentgenol 150:1117-1121, 1988

4. Cho PS, Huang HK, Tillisch J, et al: Clinical evaluation of a radiologic picture archiving and communication system for coronary care unit. Am J Roentgenol 151:823-827, 1988

5. Huang HK, Kangarloo H, Cho PS, et al: Planning a totally digital radiology department. Am J Roentgenol 154:635-639, 1990

6. Huang HK, Taira R: Infrastructure design of a PACS. Am J Roentgenol 158:743-749, 1992

7. Huang HK, Wong WK, Lou SL, et al: Architecture of a comprehensive radiologic imaging network. IEEE JSA Comm 10:1118-1196, 1992

8. Wong WK, Taira RK, Huang HK: Digital archive center: Implementation for a radiology department. Am J Roentgenol 159:1101-1105, 1992

9. Huang HK, Taira RK, Lou SL, et al: Implementation of a large scale picture archiving and communication system. Comp Med Imaging & Graphics. 17:1-11, 1993

10. Osteaux M: A Second Generation PACS Concept. Berlin, Germany, Springer-Verlag, 1992, pp 1-322

11. Huang HK, Arenson RL, et al: Multimedia in the radiology environment: Current concept. Comp Med Imaging & Graphics 18:1-10, 1994

12. Breant CM, Taira RK, Huang HK: Interfacing Aspects Between the PACS, RIS, and HIS. J Digit Imaging 6:88-94, 1993

13. Arenson RL, Avrin DE, Wong AWK, et al: Second Generation Folder Manager for PACS, in Boehme JM, Rowberg AH, Wolfman NT (eds), Proc 12th Comp Appl to Assist Radiol, SCAR 94:601-605, 1994

14. Health Level Seven: An Application Protocol for Electronic Data Exchange in Healthcare Environments. Ann Arbor, MI, Health Level Seven Inc, 1991

15. Moskowitz M, Gould RG, Huang HK, et al: Initial Clinical Experience with Ultrasound PACS. Proc SPIE Med Imaging. 2435:246-251, 1995

16. Wong AWK, Huang HK, Arenson RL, et al: Multimedia Archive System for Radiologic Images. Radiographics 14:1119-1126, 1994

17. Huang HK, Arenson RL, Dillon WP, et al: Asynchronous transfer mode (ATM) technology for radiologic communication. Am J Roentgenol 164:1533-1536, 1995

18. Ramaswamy MR, Wong AWK, Lee JK, et al: Accessing a PACS's text and image information through personal computers. Am J Roentgenol 163:1239-1243, 1994