



Exploiting Database Technology in the Medical Arena

A *critical* assessment of *integrated* systems
for picture *archiving* and *communications*

Picture Archiving and Communications Systems (PACS) are the backbone of the totally digital radiology department. They supply the communications, storage, and user-interface. PACS attempts to utilize computing power to increase the efficiency of the radiology department in a cost-effective manner. In this article, we describe the system, and the complications of database design and communications constraints. We conclude with a discussion of future research that should be addressed to further PACS. Included is a brief description of the functions and advantages of a commercial PACS system called CommView®.

The Manual Radiology Department In the manual radiology department, manual film libraries exist to store the film images. Images taken from various modalities (computed tomography, X-Ray, etc.) are stored in a patient folder. When the radiologist is ready to examine the image, he/she requests the image from the film library, schedules time to view the image, views the image, possibly views previous (older) images, discusses the case with the referring physician, and writes the examination report.

The examination report is placed with the film in the patient folder, and the folder is, once again, filed away. This process results in much wasted time and wasted resources. Images that are lost or misfiled must be rescheduled and retaken. Because the filing system is manual, the retrieval time of the image is, optimally, on the order of minutes. The lack of sufficient film reading stations makes it difficult to schedule time for the examination of the image, which in turn makes it harder for the referring physician to obtain the results in a timely manner. It is difficult to obtain a copy of the image without the need for a digitized hard-copy to be generated upon request. Additional information for a patient (billing, location, etc.) is difficult to obtain without requesting

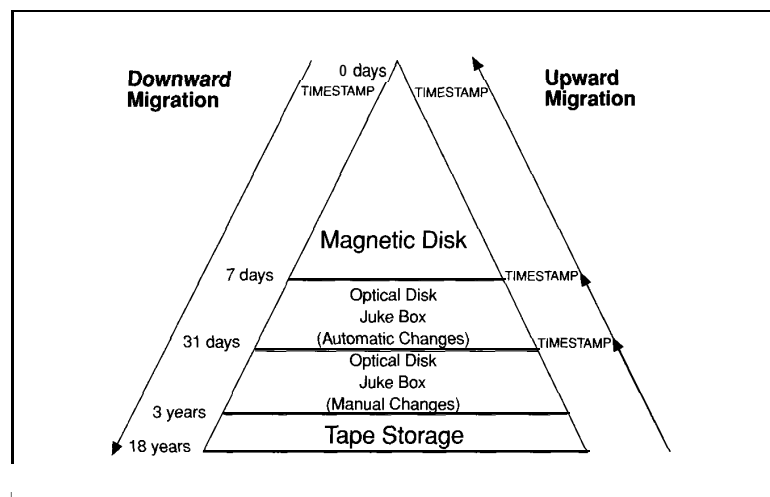
Lynn Allen
AT&T Bell Laboratories
Ophir Frieder
Dept. of Computer Science
George Mason University

that the information be retrieved from some of the other hospital systems. The automated PACS attempts to integrate all of the functions of the manual radiology department into a networked facility operating in a systematic way.

Database Issues PACS data consist of text, image, and voice [1] (multimedia data). The data consist of the patient information, including patient id, location, physician, etc. The image component consists of the multiple images that make up an examination. On average, the size of one radiological image is 10 MB [2] and can be as large as 50 MB [3] for a 2048-by-2048 [3]. Voice

is used in some PACS to store the oral voice version of the radiologist's examination report or to communicate remotely between the radiologist and the referring physician. To associate the proper voice data with the proper image, and the image with the proper text, database models, such as relational and object-oriented, are used.

Database Requirements PACS focuses on the radiological examination as the primary entity. One examination includes textual data, images, and possibly voice. A database design must be chosen such that a relationship exists that links the three data types to form one examination. Since response times maintain a two-second maximum [2], the database design must provide efficient retrieval. To maintain acceptable performance, images are compressed so that the network transfer time and utilized disk space are minimized. With the compression, though, the image must maintain an acceptable quality (lossless), and the



1. The storage level hierarchy for PACS data.

speed of the compression and decompression must be minimal [4].

When radiologists view an image during an examination, they often also access previous examinations. Therefore, the system must be able to retrieve previous data by both the primary key and by the content of the image [3].

Images/examinations are deleted if the patient dies or the image exceeds the legal 18 year retention period [2]. Since the most common operations are to add and retrieve data, emphasis should be placed on making these operations optimal. If a distributed database design is used, remote fragments must be efficiently identified to increase the performance of the PACS.

The transmission of image data must be reliable, and the system must be secure. All images and examinations in the database must be printable. Numerous people must be able to access an examination concurrently, especially when there is remote consultation between a radiologist and the referring physician.

PACS must remain operational even if one machine on the network goes down or a new one needs to be added. Through digitization, existing (pre-PACS) patient film reports must be entered into the PACS. Also, the system must be flexible enough to allow for an interface to the hospital information system (HIS) and the radiology information system (RIS).

Three-level Hierarchy

The size of the radiological images dictates that storage media are needed that are capable of handling the large volume of data. Since images must be saved for up to 18 years, three levels of the database are defined. The table shows the levels, the associated retention period, the types of storage media used, and the approximate storage capacity [5], with images older than 3 years residing in a tape storage. Figure 1 shows the leveled concept.

By using this leveled approach, the most requested images reside in the fastest memory, so that the retrieval time is minimized. Unfortunately, not all images requested by the radiologist are less than seven days old. Therefore, data migration is needed to move the older data into the faster memory for retrieval, and alternately, to move data from higher levels to lower levels once their age is greater than the maximum age for their resident level.

Downward migration is triggered by the time stamp associated with an image. Periodically, the software in charge of each level scans a table containing the time stamp for the resident images. If the date is greater than the maximum for that level, the image is moved to the next lower level, thereby freeing the faster memory for the

| Level | Retention Period | Typical Storage Media | Approximate Storage Capacity |
|-------|---------------------|---------------------------------|------------------------------|
| 1 | 1 to 7 days | Magnetic Disk | 21 GB |
| 2 | 8 to 31 days | Automatic Optical Disk Juke Box | 31 GB |
| 3 | 3.2 days to 3 years | Manual Optical Disk Juke Box | 3.2 TB |

more recent images. The image date is then recorded in the lower level time stamp table. If an image is accessed from the highest level, or the image from a lower level is updated, the image is assigned a new time stamp and resides in the highest level for a minimum of seven more days.

Upward migration occurs upon receipt of a request by a radiologist or physician for an image that is more than seven days old. When this image is accessed, the date count is reset, and the image will reside in Level 1 for a minimum of seven days. This request could either be on-line or off-line. If the request is on-line, the image retrieval could take up to a few minutes because of the access time of the slower storage media. If the request is made off-line, the images are retrieved prior to the radiologist accessing them. (The images are pre-loaded prior to his/her arrival at the viewing workstation.)

Database Design

Database design is not a trivial task. Hedgecock, et al. [6] have identified some of the required information that comprises an examination. The data include the patient's name, identification, date of birth, contact name, hospital identification, date and time of the examination, type of modality, examination type and location, images, and attending and referring physicians. The final report contains the interpretation of the examination, as well as the original ordering information. Since the images must be included with the textual information and possibly voice, standard commercial databases alone do not suffice.

Typically, only data from a single patient at a time need be viewed or accessed together. Therefore, a design for PACS that mimics the manual approach has been suggested. This approach is based on the patient folder concept [7]. When the radiologist receives the folder for the patient, he/she has access to almost all of the information needed to make the examination report. The folder contains all of the textual information stated above, plus all of the images (old and new) for the patient. Therefore, the radiologist has at hand all supporting information from previous examinations.

Many PACS databases try to mimic the folder approach in an automated fashion. It is difficult, however, to provide all of the images for a given patient whenever only one is requested. Therefore, pre-fetch algorithms and expert systems are used to predict the next requested image and to load the local device before it is requested.

Several methods have been defined to utilize alternate database approaches. An extension of the relational database model and an object-oriented approach have been formulated, as described below.

Relational Considerations

The relational database approach is popular in commercial applications. To use the simplicity of the relational database approach for text, and to enhance the model to allow for images, Martinez and Nemat [3] have defined an extension to the relational database model—the database archive system.

There are three key components in their extended relational approach: the front-end interface, the system manager, and the storage system. The front-end handles the communications using the ACR-NEMA Protocol [8]. (The ACR-NEMA Protocol was developed to facilitate point-to-point communication between imaging devices.) The system manager is the heart of the system, where database control takes place. The storage system is the storage area for text and images, and uses the three-level hierarchy discussed previously. To increase the retrieval rate of images from the Level 1 hierarchy, the image should be split across disks such that the reading and writing can be done in parallel. This approach is a modification to the horizontal and vertical partitioning schemes commonly employed in multiprocessor databases [9].

The system manager is composed of multiple functional elements and is responsible for processing queries and storing new data. The functional elements include a request processor for parsing and scheduling requests, as well as executing them. The request processor controls the allocation and de-allocation of the system resources. A database manager is responsible for accessing (storing and retrieving) the image and patient files. To

monitor the space on the storage media, the file manager oversees the physical storage by controlling all of the physical allocations and de-allocations. It also returns the physical address associated with an image. A migration manager provides the functionalities previously discussed. Lastly, an expert system processes all requests for image data by content. The system manager is that part of the system which performs the relational aspects.

The textual information relating to a patient, e.g., name, billing information, etc., is stored in the patient file. The image data file contains information about every image. This information includes specifics on the characteristics of the image data, how it is to be accessed, and a logical disk address of the location of the image. To enhance the retrieval of additional images by the radiologist or subsequent images from a multi-image examination, relationships are stored in a knowledge descriptor, which maps images to other related images based on relationships. For example, a patient has an examination that required five images. These images are related to each other by the origination of the examination, e.g., MRI, the date of the examination, the referring physician and the radiologist. In the system manager, a storage table is kept to map the logical to physical address for a given image, and to keep track of the amount of used/unused storage space [3].

Another relational database design has been prototyped by Karmouch, et al. [1]. This design is similar to the previous one, in that a relational database is used for the textual data. In this scheme, the database system is divided into two portions. The first, the radiology information database, manages the information regarding the examination such as status, and descriptors that link logically to the multimedia database system. The multimedia database system contains the physical unformatted text/image/voice. For efficiency and standard retrieval, each data type is located on a separate device such that information specific routines are implemented to handle the retrieval of data. For example, only routines that are needed to retrieve image data are installed on the image device.

With this design, the functions provided by the relational database model (select, etc.) are not available to the user. Instead, custom made functions are the interface to the user. For example: Get archived reports (Patient#).

One key problem with the relational approach is the complexity required to query the database by the content of the image. Radiologists, however, typically retrieve all images that contain a specific feature for

a particular patient. For example, "retrieve all images for John Smith that include the left hand." Designs utilizing an object-oriented approach have been developed that support such retrieval [10, 11].

Object-Oriented Design

In the PACS domain, the object-oriented database approach selected enables images to be retrieved by content. Whereas in the relational database approach, image, text, and voice are handled separately, the object-oriented approach attempts to link these types and the operations on them. In other domains, an object-oriented, network-wide, multimedia prototype system [12] collects information used by the biological community for the study of nematodes.

To store images by the object-oriented approach, descriptive information is included with the images. By developing some distinct operations on the different types of images, the user can query the database and identify those images that contain a particular surface.

To identify the content of an image, the image must first be broken down into smaller parts. At the highest level, an image is separated into segments. A segment defines information on objects contained in that segment. For example, if segment A contains an object, such as a tumor, the segment defines the coordinates of the tumor within the segment. The segments are broken down further into smaller parts, which provide more de-

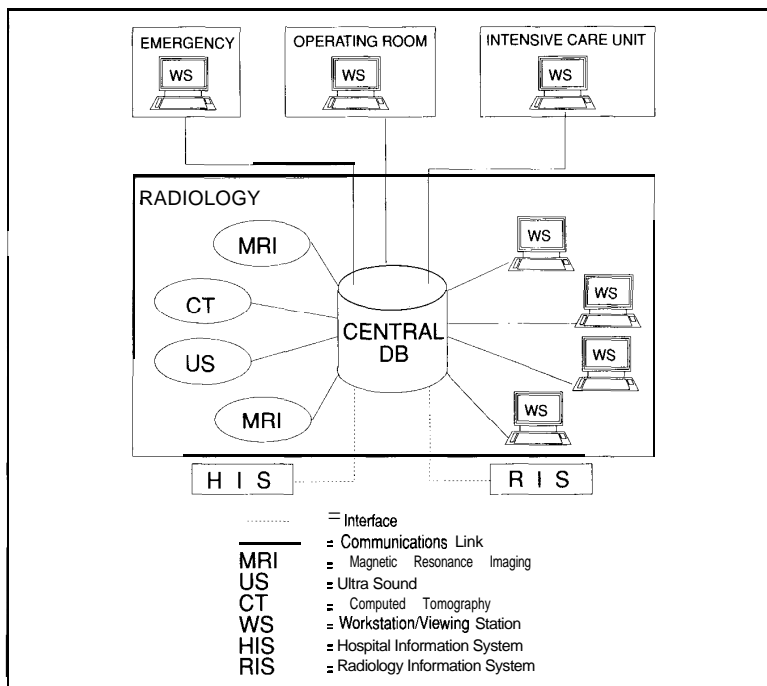
tailed information such as linear segment midpoints. Finally, the segment parts are related to the object.

Centralized Approach

A PACS database can either be located central to all modalities and display stations or it can be distributed throughout the hospital, the campus, the city, or even nationwide. For tutorials on medical image modalities and display systems, see [13] and [14], respectively. This section covers the centralized approach.

Centralized Databases were considered the "norm" in some of the first PACS systems. Once images are acquired, they are stored in the central database, and any workstation that needs an image submits a request over the network to the database. The image is then transferred over the network to the requesting workstation, where it is stored in the workstation's local memory. With this technique, the image and text data are controlled by one central node, with a relatively simple interface. Figure 2 shows a typical centralized PACS approach.

Many studies use the centralized approach [15, 16]. This approach uses one database to store and process all of the data. Requests for data come to the central point and are processed there. Although the centralized approach offers some good advantages, such as consistency and integrity, several problems have been detected. First, with a large PACS, a request to the centralized database is expen-



2. Example of a centralized PACS.

sive. The time for the central PACS to parse and execute the request, and then return the requested information, can range anywhere from several seconds to several minutes, depending on the system load. Therefore, in a very busy system, the performance is degraded and the "acceptability" of the system by users may not be forthcoming.

Second, in a centralized system, there is only one database, and any hardware

or system problems may halt all of the radiological operations. If the hardware or software needs to be upgraded, the system may need to be shut down as well, and this scenario is completely unacceptable for areas such as emergency and operating rooms. The centralized database is connected to the requesting devices via a network link through which all database requests are routed. Thus, the link is likely to become a bottleneck if the number of requesting devices connected is high. Because the data returned are images, large information packets are transmitted, which may cause the link to degrade system performance.

When dealing with a PACS, the needs for radiological images extend beyond just the radiology department. There are needs in the emergency room, the operating room, the intensive care unit, etc. In a large hospital, where there are often multiple buildings, remote workstations and acquisition modalities can be widely separated. Take also the case where a hospital supports a mobile x-ray unit that acquires images and possibly views them from locations that are far from the central database. Poor image resolution may result from noisy transmissions, more likely at greater distances.

Distributed Approach

To solve the deficiencies of the centralized database, designs using a distributed approach have been investigated [17, 18]. The distributed database is broken up into fragments of data (sections of the database). These fragments are located at several sites throughout the system. The location of each fragment must be carefully planned such that the operations from one section of the hospital are most likely to use the data in the closest local segment. If the data are stored remotely, they can be routed to the workstation's local memory to allow for decreased response time in subsequent retrievals. Also, if most of the images that need to be retrieved are stored locally, the network will not be occupied with as many requests. To determine the storage site of the data requested, the local stations must contain a directory of locations. This directory must be maintained by the system to reflect any changes.

If a user requests an image, the system will look for the station on which the image resides. If the image resides on the current station, the system only needs to acquire the image from the local disk. If the image is located on a remote station or archive, the station requests a copy of the image from the remote site. The remote station then transfers the image to the local station and loads the image into local memory. Optimally, most of the image requests from a distributed database location should be found on the local system if the PACS distributed database fragments has been effectively planned. Figure 3 shows an example of a distributed PACS where the database is broken into fragments spread across the network. Fragments exist in the emergency room, the intensive care unit (ICU), and pediatrics, as well as in the radiology department. There is a repetition of fragment C in both the ICU and the radiology department, since this are critical data and must be referenced quickly from both locations.

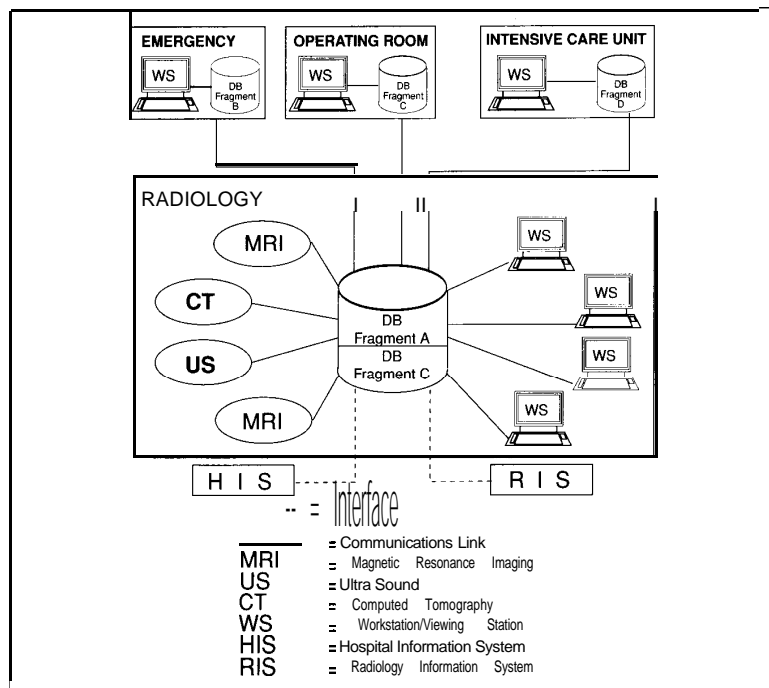
Strategically placing the data where they are most likely to be accessed increases the system performance by reducing network access time. However, with the distributed approach, network transmission issues such as data security are further complicated. Also, new concerns such as data replication are introduced.

The distributed approach to PACS allows for the planning and modular

implementation. Since the system is spread out, if one machine goes down, the other locations connected over the network remain operational. More workstations and fragments can be added to the system, thus allowing theoretically unlimited growth without requiring shut down of the system. An example of a distributed software architecture that supports the updating of programs and data without interruption is described in [19].

A distributed approach is more cost-effective since it supports incremental growth. Since PACS is quite expensive and cost-benefit analysis is inconclusive, modularization satisfies many needs without a large price tag. However, system implementation costs are higher for the distributed system (software, etc.) because additional functionality must be added to handle the database fragments at remote stations. Lastly, since the images are either transferred remotely to the local memory, or the images are immediately available locally, the quality of the image will not be degraded.

Currently, data placement strategies that attempt to maintain balanced processing times across all nodes are being developed by Siegelmann and Frieder [20]. An evenly distributed processing load results in average query response times. The communication issues associated with the design of a modern data network environment, e.g., a distributed PACS environment, have been described [21].



3. Example of a distributed PACS.

The cost of PACS is a driving force in the push for distributed processing. With a modular design, it is feasible to build the PACS piece by piece. With the centralized design, the cost of buying an entire PACS can be quite expensive [22]. All in all, distributed PACS offer more advantages over the centralized approach when the system is large and/or physically separated. Centralized PACS is good for smaller radiological systems.

Data Retrieval Techniques

Relational databases retrieve images based solely on the accompanying textual data (e.g., Patient Name, date). With an object-oriented approach, image data are retrieved based on rules that scan the characteristics of an image for matches (e.g., body parts, curves). These rules basically define an expert system. An expert system can also be incorporated into the relational database approach to allow query by image content. A function known as pre-fetching [23, 24] will assist the radiologist by fetching any related images and loading the images locally before requested. The ability of a radiologist to have comparison images pre-loaded, significantly reduces image retrieval time and hence the total diagnostic time.

Two methods utilized in pre-fetching are rule-based systems and machine learning [23]. Pre-fetching can use a standard set of rules installed on the workstation, with a probability of occurrence for that rule. The rules trigger questions about the image requested. Based on the "weights" associated with rules, the one with the highest net weight is executed. The weights are determined by consulting experienced radiologists and using their expert knowledge to determine the most common actions used in studying an image. If the radiologist frequently requests the most recent study of a patient, but rarely requests the next most recent study, the most recent study will have a higher weight than will the next most recent study [23]. For example,

RULE:

IF study = any study

AND a previous comparison study of the same modality is available

THEN prefetch the most recent study (Probability 1.0)

AND prefetch the next most recent study (Probability 0.15)

The radiologist can customize the rules to conform to a preferred retrieval scheme.

In identifying the rules in a system using pre-fetching, studies are performed to determine the usage patterns. A "knowledge acquisition" approach can assist in identifying these rules. This process involves observing the operations of the radiology department and the radiologist. Lui Sheng, et al., identified criteria for retrieval to include patient name or id, age, sex, section, current disease, medical history, modality, anatomical portion, procedure number and age of films. They showed that the retrieval patterns also seemed to be specific for certain functions of the hospital. For example, the retrieval patterns for the pediatric section were uniform throughout, but differed from the other sections of the hospital (cardiac care, etc). Therefore, different sets of rules need to be defined specific to each section of the hospital [22].

Through utilizing an expert system, the rules can grow as familiarity by the radiologist increases, and more specific rules are identified. Since learning techniques have only recently been employed in the PACS domain, much work is needed to determine their effectiveness and efficiency. Levin, et al. [23] suggest supplying the workstation with a "machine learning" mechanism to monitor the success of the radiologist in selecting one of the pre-loaded images. This learning mechanism would deduce the radiologist's pattern in selecting a next image, then alter the individual machine rules. Intense calculations are not possible, otherwise performance of the individual workstation would diminish. This method could prove productive, provided that the overall performance is not degraded significantly.

Image Compression

With the disk storage size of a medical I image on average 10 MB, it can take from seconds to minutes to retrieve an image in its original unaltered state. For this reason, techniques have been developed to compress images so that transmission time and storage space are minimized, without appreciably degrading image quality. In the medical domain, the eye must not be able to detect any change in the original image as a consequence of compression and then decompression. The quality of the image should not deter the timely reporting of an examination, nor should it cause false interpretation. To this end, much research has been done to identify techniques that will fill this need. Several accepted techniques include discrete cosine transform coding (DCT), and block truncation coding (BTC). These techniques are beyond the scope of this article. The reader is referred to [25] and [26] for further details. Rather, we deal in this section with the perfor-

formance aspect of compression, how it affects the PACS, and for which features of the PACS it is appropriate.

An image in its original unaltered state is known as an original-raster scanned image [27]. When such an image is transferred over a communication link to a remote workstation, it must be segmented into smaller pieces, and header information must be added to identify the image and enable the receiving workstation to reassemble the image. Therefore, this image cannot be displayed until the entire image is received and the individual segments reassembled, a time consuming process. With compression, however, the image can be reduced to a size that can be sent in one segment. The image is sent, and then decompressed and displayed, since only one image transmission and one image segment is required. Compressed images may also fill more than one segment to allow for higher resolution. This method has become very effective when employed in image "browsing" [27].

Through image browsing, a radiologist can effectively scan each image in a given patient file to view the actual image being sought. Upon startup of the image display, the first compressed image fitting into a single segment is displayed. If the user is not interested in viewing this image, the next single-segment compressed image is displayed, and so on. When wishing to view a particular image, the user simply stops browsing. While the image is being viewed, higher resolution images are sent over the link, and it overwrites the displayed image. Therefore, the image will continue to increase in resolution while being viewed. Since the radiologist is generally searching for specific characteristics of an image, and most of the images viewed do not meet these specifications, a majority of the images are never viewed long enough to require increased resolution. Thus, the communication link is not burdened with the heavy traffic load associated with the transmission of original raster scanned images, and the user may view the image immediately. The compression method used for this type of operation is known as pyramid data structures [28]. In particular, the reduced-difference pyramid data structure takes the difference between two-by-two nodes, with three of the four differences retained. This method reduces the size of the resultant image by 25 percent.

Each compression method has its advantages and disadvantages. With the adaptive discrete cosine transform coding (ADCT), the image quality is good, but the quality of the sharp edges is not maintained. The block truncation coding (BTC) reconstructs sharp edges well, but the general quality of the image is degraded. To reconstruct an image of high

quality, Nakagawa, et al. [29] have defined a "hybrid" technique, which utilizes both ADCT and BTC compression methods. In the hybrid technique, an original image is divided into 16-by-16 pixel blocks, and the edges are extracted by a 3-by-3 edge operator. If the maximum edge value in the block exceeds a certain threshold, the block is judged to have edges. For blocks which do not have edges, ADCT is applied to ensure good quality reconstruction. For blocks having edges, a BTC-based method is first applied to the image, then ADCT is applied to the difference between the BTC-reconstructed image and the original image to improve image quality. This hybrid technique has been shown to produce images of a quality comparable to the original, while delivering storage and transmission savings as well [29]. More detailed information on image compression is available [4, 25, 28, 29, 30].

Storage Media

Storage devices that can hold gigabytes of data with relatively efficient access time are in great demand in the PACS environment. Research continues to consistently improve PACS by providing a storage media that can hold many images and is accessed quickly.

For level 1 in the migration hierarchy, magnetic disks are preferred due to the required access rate. Currently, the most used storage devices for Level 1 are the fixed Winchester disk and the removable disk. These drives offer quick access and are inexpensive, but they can only store on the order of hundreds of megabytes of data.

Generally, for levels 2 and 3, optical disks are used. Because optical disk technology compromises rigid response time requirements, they can not be used for level 1. The more popular optical devices include CD-ROM (compact disk, read only memory), WORM (write once, read many) optical disks, and optical erasable disks. Since data can only be stored on the CD-ROM at the time of manufacture, with the exception of educational purposes [31], this device is not useful to a PACS. The WORM optical disks are the most common for levels 2 and 3. They offer large storage up to and exceeding 2 GB per disk. Typically, at these levels, the WORM optical disks will be used in a multiple-disk (jukebox) package.

An optical disk jukebox is based on the record changer concept. It allows access to more than 89 optical disks in one jukebox, and selects which disk to read from or write to. The jukebox changes a disk in approximately eight seconds [32]. The optical erasable disk can be reused by erasing existing data and writing new data. Because erasing an area on a disk and then re-writing over it is a time-con-

suming task, the time to access or write to the disk is greater for optical erasable disks than for WORMS. Since PACS use mainly add and retrieve operations, even a small amount of processing time lost due to using an erasable disk is not warranted. Furthermore, the cost of erasable optical disk technology is currently significantly greater than WORMS. Therefore, data storage WORM disks in a jukebox configuration are the best solutions, to date, for levels 2 and 3.

Additional archival storage devices include optical tape and cards. Optical tape stores approximately 1 TB of data per 12 inch reel, and the transfer rate is 3 MB per second [32]. Optical tape is useful for level 4 data (3+ years). Because optical tape stores significantly more data than a magnetic medium, fewer tapes are needed to store long term data. Another option is the optical card, which is the size of a credit card. Its storage capability currently exceeds 200 MB. Its primary application is for the storage of image data that need to be transported. Thus, the card, in effect, would be a miniature, digitally capable, integrated "filmless" patient folder [33, 34].

Communication

PACS communications involve local-area, and wide-areanetworking. Before the issue of communication technologies can be addressed, one must know how to communicate between multi-vendor modalities. In the early days of PACS, no standard existed on which equipment manufacturers could base their communication design; hence, manufacturers developed their own methods. This technique would have been effective if every modality that would ever communicate in a particular hospital was manufactured by the same company. This, however, was not the case, and a standard was thus needed.

CR-NEMA Protocol

The American College of Radiology A and the National Equipment Manufacturers Association (ACR-NEMA) joined together to develop a standard for the point-to-point interface between two imaging devices. The protocol provided a "hardware interface, a minimum set of software commands, and a consistent set of data formats for communication across an interface between an imaging device and a network interface unit or another imaging device" [8]. A standard was needed to enable imaging equipment manufactured by different vendors to communicate. The then-available data communications standards did not support high speed transmission (140 Mbps or greater), nor did they support rapid access time for large volume image data. The ACR-

NEMA Committee attempted to standardize the method of communication between two imaging devices, but fell short of providing an optimal solution to the communication problem.

The ACR-NEMA Protocol is based on the OSI (open systems interconnect) model. There are six layers defined (from highest to lowest): application, presentation, session, transport/network, data link, and physical [8]. In brief, the physical layer supports the transfer of data over the hardware interface between imaging devices or image device and a network interface unit. The data link layer supports data flow control by framing the data and adding control and error checking headers to the frames. The transport/network layer uses virtual channels to send segments of a message over the interface. The session layer ensures that an end-to-end communication link is maintained during transmission and receipt of data. The presentation layer organizes a message into groups and elements. The application layer instructs the receiving device on what operations to apply to the data.

Although the ACR-NEMA Protocol did provide a much needed standard, it has not been fully applied by the equipment manufacturers or product developers. Several problems exist with the present standard. First, it is vague; the protocol specification is incomplete. For example, Toshiba and Matrix Corporations each developed ACR-NEMA interfaces for a pilot PACS [35]. Because of a misinterpretation of the protocol, both companies implemented the frame check values differently. Without constant supervision and coordination, multiple vendors can still implement the "standard" in different ways.

Second, the packet size, now set at 4 KBytes (2K words at 2 bytes per word [8]), is insufficient to support the vast volume of data transferred and should be set higher to fully use the available bandwidth. Also, the number of data lines should be doubled to use the physical interface to its fullest. Since a PACS is highly networked, the additional specifications for a separate network layer need to be provided. McNeill, et al., describe the ACR-NEMA physical layer as a bottleneck of the PACS environment. The need to connect devices via network interface units makes the standard cumbersome to use because of difficulties with the flow control [36].

One of the greatest contributions of the ACR-NEMA protocol is the standard message format. It provides descriptive information about the image including patient name, physician name, etc. Although the performance issues are a limitation of the ACR-NEMA protocol, and are likely to be changed, the message

format is likely to survive. Because the message format is so descriptive, messages can be reassembled when the image fragments are sent separately. Also, the message format contains information that is key to the image database in relating the image to the patient.

Communication Media Use

The needs of the Radiology department and the hospital will dictate the types of communications available. For an inter-hospital PACS, a LAN fiber optic communication link can be used [35, 37, 38]. Other links can be used (Ethernet, etc.), but fiber provides the best overall transmission.

The IEEE 802.6 DQDB MAN standard allows the simultaneous transfer of radiological images and voice communications. Thanks to its high transmission rate (exceeding 150 Mbps per bus), the DQDB MAN is very effective with systems that allow browsing. It can also be used for remote consultations between the radiologist and the referring physician, where voice is actually carried over the network [27].

ISDN is a worldwide public network standard governed by the International Consultative Committee for Telegraphy and Telephony (CCITT), and the American National Standards Institute (ANSI). ISDN's goal is to provide efficient, end-to-end digital communications of voice and other types of packet-switched and circuit-switched data, simultaneously over the public switched network [39]. ISDN technology will allow a radiologist at one hospital to communicate with other offices via a phone line [40, 41] at an acceptable rate.

Satellites have also been proposed to communicate over an inter-hospital PACS. The advantage to this approach is to have the database located in one hospital, while the viewing stations may be located in different hospitals. Also, satellite transmission can be readily used to transmit images to different local databases at other locations [42].

There are many communication options. The one that is optimal for a particular hospital depends on the size of the hospital, the location of the workstations, and the stated requirements. There is no one formula for PACS. It is a custom designed system.

The CommView® System

CommView is a PACS system developed jointly by AT&T and Philips for radiology applications. The Radiology Operations System (ROS) integrates the PACS with the radiology information system (RIS) and the hospital information system (HIS). The ROS, based on a modular approach, allows for

the incremental growth of the PACS within a hospital.

Within CommView, the PACS portion implements the image management by providing the functionality for image acquisition, transmission, display, and archival. The RIS handles the administrative portions of radiology such as patient scheduling, film tracking, report preparation, billing, and operations analysis (measures of productivity). Lastly, the HIS assists with patient registration, lab system functions such as ordering, order entry, radiology procedures, and financial systems and statements. Although each system may provide similar functions separately, CommView allows for their integration while providing consistent data and reducing redundancies. Overlaying these three operational systems is a communication architecture that makes the integration complete [43].

CommView acquires images via an acquisition module (AM), which transmits the image from the particular imaging modality to the data management system. Each acquisition module can connect up to five modalities. Short distance transmissions take place over LANs such as Ethernet, StarLan, and AT&T ISN/Datakit. Remote transmissions are via contract microwave services, dial up analog services, ACCUNET, and ISDN. CommView supports the bus, ring and star LAN topologies. Viewing capabilities include a standard viewing screen that allows for 1024-by-1024-by-8 bit image display, an advanced viewing station that is capable of 1024-by-1280 resolution with 64 concurrent images displayed, and a high resolution 2K-by-2K display station. CommView has good capabilities and may be configured in a variety of ways.

Future Considerations

Now that the first set of Picture Archiving and Communication Systems are being implemented, there are still some major issues that will affect the second wave of PACS. The ACR-NEMA protocol should be enhanced. There is also a need for fast networking capabilities. Storage devices that provide high density storage and minimal access time are required. Until the day when we see a read/write optical jukebox capable of handling hundreds of terabytes of data with a response time of under two seconds, the PACS community will not be satisfied.

Database technology will continue to pursue a flexible method of image retrieval by content. Knowledge acquisition techniques and expert systems will continue to be used and analyzed for their effectiveness. Methods of pre-fetching to load the next requested image need to be studied in greater detail.

Motivated by the rigid response time demands, an integrated, parallel, archival and display system is being studied at George Mason University. Multiprocessor based image generation algorithms, such as the isocontour extraction algorithm [44], are used to reduce the image generation time. By storing the data in a hypercube multicomputer database system and executing the image generation algorithms on the same database engine, an integrated parallel archival and display system is achieved [9]. Other areas of research include user interface and workstation design [14, 45, 46, 47], a wide variety of networking [35, 37, 38, 48], interfaces to HIS and RIS [49], and networking capabilities to transmit image and voice at very high speeds [50].

Conclusion

A Picture Archiving and Communication System is the solution to the progressing need to computerize the radiology department. PACS differ in their demands from many computer systems in their composite requirements of fast response time, high quality display, remote and local access, large data requirements, and long data retention (up to 18 years). PACS help to provide efficient support to radiologists and other physicians, providing better accessibility to examination results. By allowing remote consultations, time and money are conserved. By permitting hospitals over a wide area to share images, the patient will not be subjected to, and will not need to pay for, duplicate images. To gain fully from the PACS, it should be integrated with the HIS. Thus, whenever an order is placed for an examination, the HIS will cross-check with the PACS system and ensure the patient's previous images are retrieved [51].

PACS research has not ended. It will continue to grow with each new technology. The quest will continue to improve the performance, the display, and the accuracy of the PACS. Although different techniques are applied to utilize a PACS, a new PACS is modeled for each radiology department. Unfortunately, there is no one formula or design for a general PACS. Until such a design is available, PACS research will need to continue and will be based on the needs of the individual hospital.

Lynn Allen is a Senior Technical Associate at AT&T Bell Laboratories. Previously, she was an Information Management Member of the Consumer Communications Services Business Unit of AT&T. From 1988 to 1989, she was a Member of the Technical Staff in the Defense Systems Division of Computer Sciences Corporation.

Allen received her B.Sc. (1988) in computer science from Trenton State College. She received her M.Sc. (1991) in computer science from the New Jersey Institute of Technology.

Ophir Frieder received his B.Sc. (1984) in Computer and Communications Science, and his M.Sc. (1985) and Ph.D. (1987) in Computer Science and Engineering, all from the University of Michigan. From 1987-1990, he was a Member of the Technical Staff in the Applied Research Area of Bell Communications Research.

Currently, he is on the faculty of Computer Science at George Mason University. His research interests include parallel and distributed architectures, database systems, operating systems, and medical imaging architectures. He is a member of Phi Beta Kappa and the IEEE Computer Society. Address for correspondence: Department of Computer Science, George Mason University, Fairfax, VA 22030.

References

1. **Karmouch A, et al:** Multimedia Database Design and Architecture in Radiology Communications. *SPIE Medical Imaging IV: PACS System Design and Evaluation*, 1990.
2. **Liu Sheng 0, et al:** Requirement Analysis for PACS Database System. *SPIE Medical Imaging IV: PACS System Design and Evaluation*, 1990.
3. **Martinez R, et al:** Image Data Base Archive Design Using Parallel Architectures and Expert Systems. *SPIE Vol. 914 Medical Imaging II*, 1988.
4. **Ahranson J, et al:** Data Compression in the PACS Environment. *SPIE Vol. 1091 Medical Imaging III: Image Capture and Display*, 1989.
5. **Martinez R, et al:** Image Migration in a Three Level Data Base Archive System. *SPIE Vol. 914 Medical Imaging II*, 1988.
6. **Hedgecock MW, Jr, et al:** Database requirements for PACS. *SPIE Medical Imaging IV: PACS System Design and Evaluation*, 1990.
7. **Seshadri SB, et al:** Software Considerations in the Design of an Image Archive. *SPIE Medical Imaging IV: PACS System Design and Evaluation*, 1990.
8. American College of Radiology and the National Electrical Manufacturing Association: Digital Imaging and Communications - ACR-NEMA Standards Publication No. 300-1988, *National Electrical Manufacturers Association*, Washington, DC, 1988.
9. **Frieder O:** Multiprocessor Algorithms for Relational Database Operators on Hypercube Systems. *IEEE Computer*, 23(11), November 1990.
10. **Auhry F, et al:** An Image Handling System for Medical Image Processing. *SPIE Vol I 137 Science and Engineering of Medical Imaging*, 1989.
11. **Kofakis P, et al:** Image Archiving by Content: An Object-Oriented Approach. *SPIE Vol. 1234 Medical Imaging IV: PACS System Design and Evaluation*, 1990.
12. **Schatz B:** Building an Electronic Scientific Community. *IEEE 24th Hawaii International Conference on System Sciences*, vol 3, 1991.
13. **Stytz MR, Frieder O:** Three-Dimensional Medical Imaging Modalities: An Overview. *CRC Critical Reviews in Biomedical Engineering*, 18(1), July 1990.
14. **Stytz MR, Frieder O:** Computer Systems for Three-Dimensional Diagnostic Imaging: An Examination of the State of the Art. *CRC Critical Reviews in Biomedical Engineering*, 19(1), July 1991.
15. **Cho PS, et al:** Centralized vs. Distributed PACS for Intensive Care Units. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
16. **Morioka C, 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.
17. **Lui Sheng 0, et al:** Distributed Database Design and Modeling for PACS. *SPIE Vol. 1234 Medical Imaging IV: PACS System Design and Evaluation*, 1990.
18. **Lui Sheng 0, Chen Garcia 0:** The Design of Medical Image Databases: A Distributed Approach. *IEEE 1990 Phoenix Conference on Computers and Communications*, 1990.
19. **Frieder 0, Segal ME:** On Dynamically Updating a Computer Program: From Concept to Prototype. *Journal of Systems and Software*, North-Holland, 14(2), February 1991.
20. **Siegelmann HT, Frieder 0:** A Genetic Algorithm for Document Allocation in Multiprocessor Information and Retrieval Systems. *IEEE International Conference on Systems, Man, and Cybernetics*, 1991.
21. **Frieder 0, Shuey RL:** Communication Needs in a Data Engineering World. To appear in *Computer Networks and ISDN Systems*, Elsevier Science Publishers B.V. (North Holland).
22. **Freiherr G:** PACS Industry Dismembers Itself in Hopes of Building for Future. *Diagnostic Imaging*, March 1989.
23. **Levin K, et al:** Methods to Prefetch Comparison Images in Image Management and Communication Systems (IMACS). *SPIE Vol. 1234 Medical Imaging IV: PACS System Design and Evaluation*, 1990.
24. **Liu Sheng 0, et al:** IRES Image Retrieval Expert System. *SPIE Medical Imaging IV: PACS System Design and Evaluation*, 1990.
25. **Nakagawa Y, et al:** A High Speed Image Compression Method for Medical Images. *SPIE Vol. 914 Medical Imaging II*, 1988.
26. **Delp EJ, Mitchell OR:** Image Compression Using Block Truncation Coding. *IEEE Transaction on Communications*, Vol. Comm-27, 1979.
27. **Vallee R, et al:** Modeling and Simulation of Multimedia Communications Networks. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
28. **Wang L, Goldberg M:** The Reduced-Difference Pyramid: A Data Structure for Progressive Image Transmission. *Optical Engineering Journal*.
29. **Nakagawa Y, et al:** Prototype of Image Compression System for Medical Images. *SPIE Medical Imaging IV: PACS System Design and Evaluation*, 1990.
30. **Chameroy V, et al:** Applications of High Compression Coding to Dynamic Medical Image Series. *SPIE Vol 1137 Science and Engineering of Medical Imaging*, 1989.
31. **Mankovich NJ, et al:** Image Archiving: Hardware and Database Technology. *SPIE Vol. 1091 Medical Imaging III: PACS System Design and Evaluation*, 1989.
32. **Britt MO, et al:** Optical Archive Organization and Strategies for the 1990s. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
33. **Cho PS, et al:** Personal Digital Image Filing System. *SPIE Vol. 1234 Medical Imaging IV: PACS System Design and Evaluation*, 1990.
34. **Kerlin BD:** The Optical Memory Card as a Transportable Image Archiving Medium in a Digital Imaging Network. *SPIE Vol. 914 Medical Imaging II*, 1988.
35. **Dallas WJ, et al:** A Fiber-Optic Network System for PACS. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
36. **McNeill KM, et al:** Network Software for Picture Archiving and Communications Systems. *IEEE 1990 Phoenix Conference on Computers and Communications*, 1990.
37. **Martinez R, et al:** Design and Performance Evaluation of a High Speed Fiber Optic Integrated Computer Network for Picture Archiving and Communications System. *SPIE Vol. 1091 Medical Imaging III: PACS System Design and Evaluation*, 1989.
38. **Ricca SP, et al:** Impact of advanced fiber optics and ISDN technologies on PACS networking. *SPIE Medical Imaging III: PACS System Design and Evaluation*, 1989.
39. **Roca RT:** ISDN Architecture. *AT&T Technical Journal*, vol 65, Issue 1, Jan/Feb 1986.
40. **Blaine GJ:** ISDN: Early Experiments as a Wide Area extension to LAN-based PACS. *SPIE Vol. 1234 PACS System Design and Evaluation*, 1990.
41. **Jost RG, et al:** PACS- Is there light at the end of the tunnel? *SPIE Vol. 1091 Medical Imaging III: PACS System Design and Evaluation*, 1989.
42. **Seshadri SB, et al:** Satellite Transmission of Medical Images. *SPIE Vol. 914 Medical Imaging II*, 1988.
43. **AT&T and Philips:** Discover Your Entryway to PACS. AT&T Document Number 4535 983 00805, 1989.
44. **Frieder 0, Stytz MR:** Dynamic Detection of Hidden Surface Within a MIMD Architecture. *Proc of the Third IEEE Symposium on Computer-Based Medical Systems*, Chapel Hill, North Carolina, June 1990.
45. **Braudes RE, et al:** Workstation Modeling and Development: Clinical Definition of a Picture Archiving and Communications System (PACS) User Interface. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
46. **Dullien RC:** Progress in PACS -Definition of the Principal Workstation Types. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
47. **Gee JC, et al:** User Interface Design for a Radiological Imaging Workstation. *SPIE Vol. 1093 Medical Imaging III: PACS System Design and Evaluation*, 1989.
48. **Kohli J:** Medical Imaging Applications of Emerging Broadband Networks. *IEEE Communications Magazine*, Dec. 1989.
49. **Lodder H, et al:** HIS-PACS Coupling: First Experiences. *SPIE Vol. 914 Medical Imaging II*, 1988.
50. **Lindsey M, et al:** Design of a Network Interface Unit for a Picture Archiving and Communication System Network Operating at 200-500 Mbps. *IEEE 1990 Phoenix Conference on Computers and Communications*, 1990.
51. **Cannavo MJ:** Low-risk Strategy for PACS Calls for Modular Phase-in. *Diagnostic Imaging*, July 1988.