

The Operational Impact of Architectural Alternatives for Radiological Imaging Workstations

Kevin M. McNeill

Characteristics of workstations for use in digital radiological imaging have been investigated for many years. However, much of this investigation has focused on the workstation in isolation, and has often been directed almost entirely at user interface issues. Certainly these issues are critical, but with the increasing use of commercial workstations it is important to look at the workstation in the context of the medical information environment and examine some important underlying characteristics required to meet the demands of digital radiology. This article examines the role of storage components in these workstations both architecturally and operationally. Both aspects are viewed with consideration of their impact on the internals of the workstation and its interaction with the external information system. By considering these aspects of the workstation it is apparent that local storage and image preloading are required to support diagnostic viewing. Additional operational and architectural strategies are required to efficiently manage information within the workstation.

Copyright © 1993 by W.B. Saunders Company

KEY WORDS: information system, medical image, workstation, image management, storage subsystems.

RADIOLOGICAL IMAGING takes place within an environment rich in interactions between information producing and consuming entities. This article examines the requirements for storage subsystems in workstations that support radiological imaging. However, to accurately define those requirements it is necessary to consider both the specific and general context in which such a workstation will operate. The specific context is the department of radiology. The general context is the entire hospital or medical organization.

THE MEDICAL INFORMATION ENVIRONMENT

Five categories of information required in the medical environment have been defined¹: (1) patient data; (2) medical knowledge; (3) communications between hospital entities; (4) information for patients; and (5) administrative/management information. These types of data and information exist in a variety of formats that may be standardized across hospitals, within a hospital, or not at all. Even today most of this information is likely to consist of hand written text on forms with perhaps five carbon copies.

Other formats include plots (eg, EEG or EKG), dictation tapes, and images on film.² It is important to note that although "data" and "information" are often synonymous, from an information systems point of view "data" refers to physically collected raw values and "information" refers to the meaning attributed to those values by a user.

The increasing application of computer technology in the medical environment has brought about the discipline of medical informatics. This discipline is concerned with applying computers as a source of clinical information by developing techniques that allow the clinician to quickly locate specific relevant information. This task is made especially difficult by the proliferation of computer-based information systems in medicine. These systems range from departmental information systems and research databases, to bibliographic retrieval systems and image management systems.

Organizational Considerations

Modern hospitals are an aggregation of many semiautonomous entities that provide very specific services. These services may involve direct patient care or they may be indirect support services. Each of these entities requires information as input (eg, an examination request form) and may generate additional information. This additional information may be directly used by another entity or may be supporting information that is archived by the producing entity but not distributed.

Each of these service entities may use computer-based information systems to support their operation, especially departments such as the clinical laboratory, pathology, and radiology. Such systems provide functions that may include patient tracking, examination scheduling,

From the Department of Radiology, University of Arizona, Tucson, AZ.

Address reprint requests to Kevin M. McNeill, MS, Department of Radiology, University of Arizona, Tucson, AZ 85711.

*Copyright © 1993 by W.B. Saunders Company
0897-1889/93/0602-0007\$3.00/0*

collection of demographic information, and report processing.

Hospital Information Systems

The concept of an "integrated" hospital information system (HIS) has existed for more than 20 years. In such a system, information would reside in a central database, stored in a coherent format that allows integration of data generated by different applications. Information would be easily accessed by all users with appropriate security clearance. The HIS would provide efficient means to support operation and management, hospital personnel needs, patient needs, and clinical research.

The task of developing a single system to support the information needs of an entire hospital has proven easier to imagine than to implement. As a result, most HIS implementations are actually a collection of independent information systems serving individual entities.³ These systems have a limited amount of coupling between systems to form an "integrated" information system. Typically, the capabilities to exchange data between systems is very limited.

THE INTEGRATED INFORMATION SYSTEM

Clearly the vision of medical informatics is to provide an integrated information system in which all of the types of data described above are available, easily and quickly, to appropriate health care workers where they work. This vision remains unrealized for reasons that fall into two basic categories: nontechnical (or sociopolitical) reasons and technical reasons.

Nontechnical Issues

A common question arises in many large organizations with information intensive activities. The question is whether to implement centralized or distributed information management. There has been a long standing debate in the information management community about the relative merits of these two strategies and there is no clear indication that one is better than the other. Often, the relative merits are blurred by issues of control.^{4,5} Although centralized data management may be more efficiently implemented, some departments may resist relinquishing control over "their" information. In

the medical information environment, the application of information systems in many hospitals has historically come about as the implementation of stand-alone departmental systems. Attempts to transfer control of these systems to a centralized information department may encounter resistance.

In hospitals with a central information systems department having responsibility for operating information systems on the behalf of other departments, integration may present another problem. The field of medicine is very strongly divided into specializations, each of which maintain tight control over the information or knowledge that justifies their existence. Providing an integrated information system that alters the flow of information, or makes it possible to bypass existing specialized providers of information, are certain to meet with end user resistance.

Technical Issues

Technical problems in integrating information systems are usually easy to define and are commonly underestimated. These problems range from the mundane questions of connector compatibility to state-of-the-art questions of whether the technology exists to process, transfer, and store information as required by the application. Technology includes not just hardware but also software built on the capabilities of the hardware. For applications in the medical information environment another critical technological issue is that of standardization. To achieve the integrated HIS, a wide range of applications must communicate. The availability of widely supported communications and applications standards have a significant impact on the ability of software developers to realize the system.

RADIOLOGICAL IMAGING

A typical radiology department is a service entity within the hospital. The service provided by the department is the generation of a diagnostic report or, increasingly, an interactive consultation with other physicians. Services are performed as a result of consultation requests sent to the department by other physicians. In a hospital with an associated medical school, the

department will have the additional missions of providing teaching and conducting research.⁶

The Radiological Information Environment

In many respects, the information environment of the radiology department is simply a microcosm of the medical information environment. The radiology department consists of several divisions (eg, diagnostic radiology, ultrasound, nuclear medicine, film library, etc), each of which use different imaging techniques. There is a flow of information between these divisions as well as between the radiology department and external entities.

The radiology information system. The goals of applying computer-based information systems to radiology are similar to those of the HIS. Both systems aim to maintain a coherent store of information, make information available to qualified users at their location of work in a timely fashion, and to present the information to the users in the format required by the user. However, there are very challenging differences that have limited the effective application of information systems technology in radiology. Radiology deals with distinct categories of data: text, graphics, and images. The text and graphics information is similar in content to that found in other hospital entities. It consists of reports, demographic information, scheduling information, charts, curves, etc. Typically this information is organized in units of a few hundred or a few thousand bytes. In the past, the technical requirements for the broad application of computer-based systems to imaging were impossible to meet and only the textual information needs were addressed in the development of radiology information systems (RIS). Today, the technology is much closer to meeting the challenge of imaging; however, the distinction between RIS and image management systems is likely to remain long into the future.

Most large hospitals use an RIS in an essentially stand-alone operation. Integration with the HIS is either nonexistent or consists of a very limited exchange of data. The limited integration that can be found is generally a serial link, or personal computer-based gateway.

Radiological images and their management. When we include images in our consideration of information management in radiology, the magnitude of technical challenge increases dramatically. In the existing context, radiology must deal with the physical logistics of processing, handling, storing, and retrieving thousands of films every year. Current management protocols are oriented around the physical handling and storage of film, with films collected into physical jackets that are then collected into a library. Management of film requires a large amount of expensive floor space and personnel. Current procedures may also result in slow film retrieval or lost films.⁷

In the context of digital radiology, image information requires the handling of billions of bytes of information and presents enormous problems for storage, communication, and presentation. The situation for image management systems is even less advanced than that for the RIS because of the prototypical nature of most existing systems. Individually, most of these problems are technology-based and should be solved by near term advances. However, the integration of individual components into a full-scale image management and communication system that can replace the film-based system is still a very difficult problem.

The Image Management and Communication System

The image management and communication system (IMACS) is envisioned as a completely computer-based system for managing image information from acquisition to diagnostic display, fully integrated with the RIS and HIS. This environment will provide the radiologist with a wealth of information to improve the accuracy of diagnosis, thus improving the quality of patient care and increasing the efficiency with which the department provides service. The problems and expense of the cumbersome film-based system will disappear.⁸ This is certainly an optimistic goal for IMACS. However, given the lack of integrated HIS, the poor coupling between RIS and HIS, and the technical challenges of image information management, a more realistic short-term goal would be to match the capabilities of the current film-based system at a reasonably similar cost. The

remainder of this article will focus on successively smaller pieces of the IMACS puzzle to look at the requirements for components that are critical to even this simplified goal.

Components. An IMACS is built from the following high level component systems: (1) image acquisition systems, (2) image archiving systems, (3) image display systems, and (4) the communications network. Each of these components may have subtypes. Because these components themselves are systems, they are each comprised of individual components. Each of these high level systems provide their own set of requirements and challenges (Fig 1).

Architectures. In addition to the components, a critical questions is presented: "How do we connect these components to make an IMACS?" The architecture of the IMACS is very closely related to the communications network component; however, it is a broader issue that includes not just consideration of physical connections but also protocol issues, global software architecture, and integration with external systems. The system architecture feeds back requirements and design constraints to the component systems. For example, the system architecture will influence the choice of centralized or distributed archiving systems. The architecture of the IMACS can reflect the structure of the department, or it can force a restructuring of the department. Additionally, the choices for architectures will be highly constrained by the capabilities of the system components. Some representative architectures are shown in Figs 1 and 2.^{9,10}

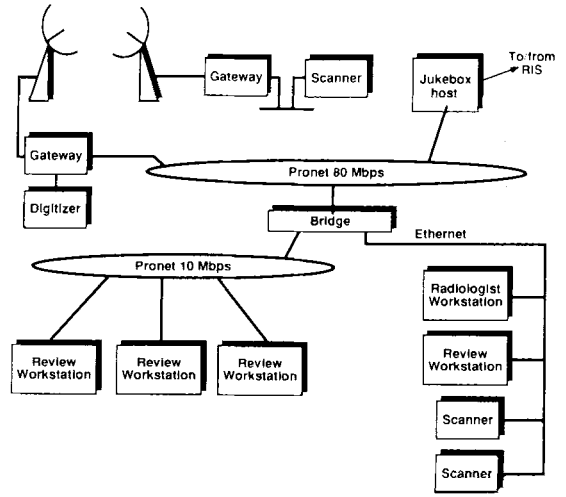


Fig 2. University of Pennsylvania IMACS architecture.

THE RADIOLOGICAL IMAGING WORKSTATION

If we were to select the single most critical component of IMACS, we have no alternative but to choose the image display workstation that is used for diagnosis. The service provided by a radiology department is centered around the diagnostic interpretation of images. In IMACS, the workstation becomes the tool by which this interpretation is accomplished. If this component is not accepted by the radiologist, the entire system will not accomplish its goals, regardless of the capabilities of the other components. In this section, we will reduce the workstation to its components and look at a few common architectures for arranging these components.

Components

The workstation must satisfy the primary mission of displaying high resolution images for diagnostic interpretation. Support for that mission requires components to communicate with external systems, store information, perform required processing, and interface with the human user. Note that the human computer interface must provide for display of information to the user and acceptance of commands from the user. There is a great variety of technology available for each of these components. However, the criteria for selecting technology for the radiological imaging workstation is simple: we must choose the leading-edge technology. We require the highest perfor-

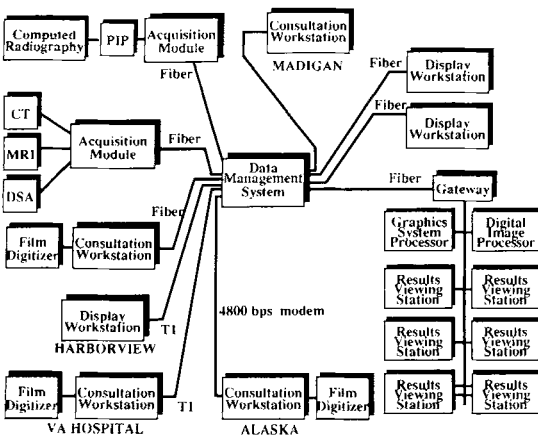


Fig 1. University of Washington IMACS architecture.

mance in terms of capacity, resolution, and transfer rates.

Architectures

As with the IMACS itself, given the components we must choose an architecture for connecting them into a system. In the case of the workstation, there are fewer options for interconnection. Most typically, the system architecture follows that of general purpose computers with components interconnected by a shared bus such as VMEbus or Multibus. To achieve the higher performance required when dealing with large images, it is also common to find secondary buses dedicated to image transfer.¹¹⁻¹³ In these systems, the control operations are performed using the standard bus, but image transfers take place over special data paths that have higher bandwidth. This additional bandwidth is provided by using higher speeds, wider paths, or both. For example, add on imaging subsystems such as the PIXAR image computer (Pixar, San Rafael, CA) have 256-bit wide data paths. This requires less than 10 million transfers per second to achieve 240 MBytes/s. Generally, the path of critical importance is the path from storage to processor to display (s.p.d.). A general representation of this path is shown in Fig 3. Most differences in architecture are focused on the performance of this data path. The impact of these differences are discussed in greater detail later. There may be a several variations on the s.p.d. in which any of the components other than the display may be absent.

STORAGE SUBSYSTEM

From the standpoint of image management, we are concerned with the storage of images in the workstation and the movement of images between storage components and other system components, especially those in the critical path to the display and the external communications




Fig 3. Diagram of the path from storage to processor to display.

component. The storage component must consist of one or more selections from the memory hierarchy. The memory hierarchy spans the continuum from magnetic tape to high-speed registers. We are generally concerned with random access memory (RAM) and magnetic disks. How we use these two forms of storage depends on the architecture of the workstation and the design philosophy of the IMACS.

One strategy for workstation design would use no disk storage at the workstation. The only storage at the workstation would be the RAM used for the frame buffer(s). Each time an image is to be displayed, the information is fetched from the archive across the communications network. This philosophy has very important implications. First, the communications interface is inserted into the s.p.d. pathway because the storage components (the disk and the RAM) are now connected by the external network (Fig 4). Second, the requirements for the communications network escalate dramatically, both in terms of capacity and transmission bandwidth. Third, the performance requirements for the archive components are much greater because the response time of the archive is in the critical path for every image access at every workstation served by that archive. This strategy simplifies workstation design but greatly complicates the overall IMACS design and does not really reflect film-based operation.

An alternate strategy would be to more closely emulate the current system. In the current film-based system, local image storage is associated at each viewing location. Images are typically fetched before the scheduled reading time. The number of requests for images in the library (archive) is small. This strategy decouples the communications network from the s.p.d. pathway, results in lower instantaneous demand on the network, and eases the response time requirements of the archive. Determination of the set of images to be prefetched can be accomplished by an image retrieval expert system.¹⁴ The resulting increase in complexity and cost of the workstation is more than offset by the benefits to the overall IMACS design. The following sections will focus on workstation storage requirements assuming selection of this strategy.

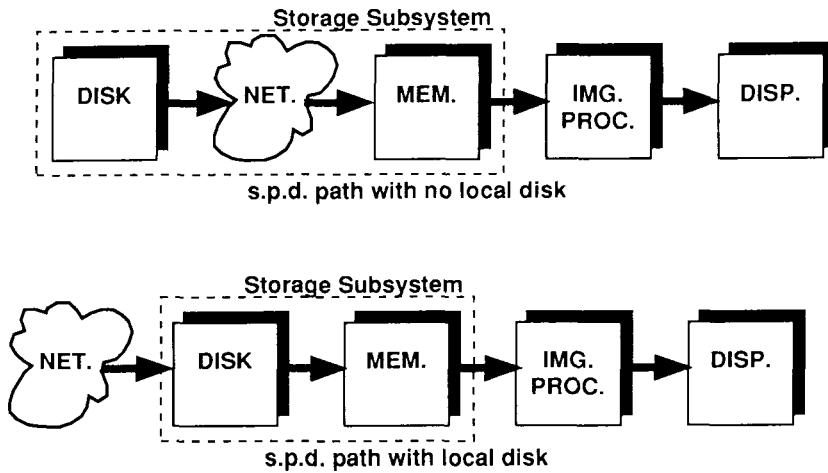


Fig 4. Storage location strategies.

General Requirements

To determine the general requirements for workstation storage, we must have an idea about the operational characteristics of the workstation. These characteristics are determined primarily by the task requirements of a radiology department. However, there are many modalities and types of reading tasks, such as alternator-based reading, outpatient reading, and wet reading.¹⁵⁻¹⁷ To determine general requirements for a workstation in a complete IMACS environment, it will be sufficient to use the most stringent requirements. If these can be met, those of less demanding tasks should follow. In terms of storage requirements, the most demanding type of reading is alternator based.

Operational issues. Alternators are used for reading inpatient films. The alternator is equipped with a working surface, a hot light, a dictation machine, and a telephone. A status board associates patient names with a panel index. Films are placed on the alternator by film library clerks before the arrival of the radiologist.¹⁷ The clerks select images from the patient folder based on their knowledge what studies are used for each type of case. The patient folder, with unselected images, is available in a rack below the reading panels. The radiologist sequentially selects a patient, indexes the associated panel, interprets the films, and dictates a diagnostic report. Infrequently, a film may be retrieved from the folder and viewed. Also, a film is occasionally taken off the panel and viewed using the hot light. This is approximately

equivalent to manipulation of contrast/brightness. Very infrequently a radiologist may need a film that was not selected by the film clerks to be placed at the alternator. In the vast majority of cases, the films viewed at the alternator are all of the same type. The frequency of reference or comparison with films of other modalities is very low.¹⁸ The average time required for the reading process typically ranges from 30 to 100 seconds.^{16,17}

Capacities and performance. Alternators vary in capacity with a range of 20 to 56 panels. Each panel consists of four boxes. Each box can hold one 14 × 17-in film, and two panels are visible at any time. If we assume that the alternator is generally fully loaded, a workstation must provide equivalent image storage capacity. The alternator holds a total number of images ranging from 80 to 224. If we consider 2,000 × 2,000 pixel 12-bit chest images, uncompressed each image is approximately 8 MBytes. Therefore, we require between 671 MBytes and 1.8 GBytes of storage. Additionally, the workstation should have approximately 20% excess capacity for wet readings or other unplanned image storage. Other researchers have derived similar estimates.¹⁹

In general, a single case will consist of four to six films and will not span more than two alternator panels. The time required to access a case will depend on the distance from the current panel to the panel containing the new case and will range from 2 to 60 seconds with an average of approximately 30 seconds.^{17,20} For a six image case, we will have 50 MBytes of image

data to move between local disk and RAM. Assuming we have sufficient RAM to buffer all the images, we need a sustained throughput of between 50 MBytes/s and 1 MBytes/s to emulate the alternator response time. How closely we can approach the upper bound is determined by the architecture of the workstation.

Architectural Considerations

The storage subsystem is a component of a larger system. It's performance characteristics must, therefore, be considered in their relationship to other components of the system. This is especially true of the other components in the s.p.d. pathway. In addition, the storage subsystem is composed of two major components: RAM and magnetic disk. Both disks and RAM vary widely in performance. For disks, parallel transfer drives offer among the highest performance coupled with great capacity. For example, a Storage Concepts S2 Disk Array controller (Storage Concepts, Irvine, CA) can provide approximately 32 GBytes of capacity and a sustained transfer rate of up to 20 MBytes/s over special high speed interface channels. However, we can also find very capable disk subsystems for standard buses such as VMEbus. The Storage Concepts SM Disk Controller connects to the VMEbus with a capacity of up to 1.6 GBytes and sustained transfer rate of up to 15 MBytes/s.

Relationship to other system components. The storage subsystem must act as a sink for data arriving through the external communications component. We would prefer the disk not to be a bottleneck, but there are several available communications technologies that are easily capable of overwhelming even the fastest available disk subsystems. It is apparent that the available throughput of both disk and communications channel will continue to increase for some time. However, the gap can be expected to widen because communications technology can be photonic-based and provide much greater bandwidth than disk technology which is limited to electronics for the foreseeable future.

The storage subsystem must act as a source for the image processor and display components. Generally, we can expect that the storage subsystem will consist of RAM and magnetic disk, and that the image data which is currently

available for display will reside in the RAM. This division offers performance advantages that will be discussed later. Higher speed disk technology that would negate the need for a RAM buffer may only be a few years away.

In addition to the components of the s.p.d. pathway, the storage subsystem will typically have to interact with the central processing unit (CPU). This interaction will primarily consist of exchanges of command and status information. If we use the same disks for storage of all data, then there will be data accesses by the CPU, and possibly other components, to store/retrieve text, programs, graphics or other data. Whether such accesses occur or not depends on the logical storage architecture.

In bus architectures. Such as those shown in Fig 5, a shared bus provides the s.p.d. pathway. These architectures may use a single bus (top) or two buses (bottom). This architecture provides an advantage because most commercial components (memory boards, cpus, etc) are designed to interface to industry standard busses. Many standard busses can support the highest performance disk controllers. For example, VMEbus will support transfer rates of approximately 20 MBytes/s and the highest performance disk controllers currently available operate at up to 15 or 16 MBytes/s. There are also industry standard secondary busses, such as VSB, which are designed to provide increased performance by either providing a wider data

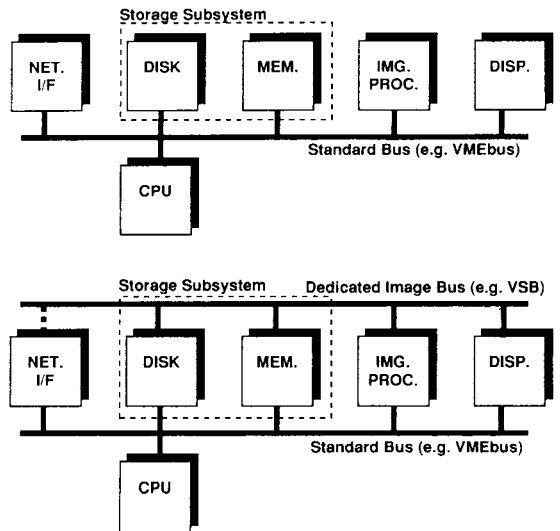


Fig 5. Bus oriented architectures.

path or higher transfer rate. These busses may also be oriented toward block mode transfers that move large amounts of data efficiently.

An important disadvantage of bus oriented architectures is contention. Because the bus is shared, any concurrent activity between system components must contend for the shared bus. This contention is worse in the single-bus architecture, but it is still a factor in dual-bus architectures. As an example, we may soon be able to initiate the display of a case from disk when a message arrives at the network interface. Interactions between the CPU and the network interface (eg, interrupt processing) contend with the disk controller for the use of the bus. In the dual-bus architecture contention with components not on the s.p.d. pathway is not an issue. However, contention will still prevent overlapping of input/output (I/O) and processing because we cannot move an image from disk to memory while the processor is accessing memory to move an image to the display (Fig 6).

The pipelined architecture eliminates contention between components of the s.p.d. pathway. However, these paths impose more restrictive requirements on the components that may make it difficult to use standard commercial equipment and limit the choice of components that can be interconnected. It also requires the components, such as the memory, to have multiple ports. This architecture will support concurrent operations that can decrease the time required to display a case. However, we must still deal with contention (now in the memory) which will prevent us from realizing the ideal overlap of I/O and processing operations.

Operational issues. The digital workstation must support the same throughput as the alternator to be accepted by the radiologist. Figure 6 shows the serial processes involved in displaying an image. Several of the major timing factors are indicated in the diagram (Fig 7).

In a simple scenario of serial operation, the typical case described above (six images totaling 50 MBytes) the best time to display the image will be given by:

$$T_{disp} = 6*(T_s + T_{x_{dm}} + T_{x_{fpd}}) + \epsilon$$

$$T_{x_{dm}} = I_{size}/R_{sustained}$$

Where ϵ represents time imposed by other factors such as bus conflict (if applicable). This time will vary for each access and should be negligible so it will be ignored in this discussion. The value I_{size} indicates the size of a single image and the value $R_{sustained}$ is the sustained transfer rate from disk to memory. For some representative calculations we will consider a VMEbus based architecture, using the following data size and rate numbers:

$$I_{size} = 8.39 \text{ MBytes}$$

$$R_{sustained} = 15 \text{ MBytes/s}$$

$$T_{x_{dm}} = 8.39 \text{ MBytes}/15 \text{ MBytes/s} = 0.56 \text{ s}$$

$$T_{x_{fpd}} = 8.39 \text{ MBytes}/40 \text{ MBytes/s} = 0.21 \text{ s}$$

$$\text{Assume average seek time } T_s = .035 \text{ s}$$

$$T_{disp} = 6*.035 + 6*0.56 + 6*0.21 = 0.21 + 3.36 + 1.26 = 4.83 \text{ s}$$

These calculations assume we have sufficient memory to buffer all six images, and that the entire process is serial. They also assume that

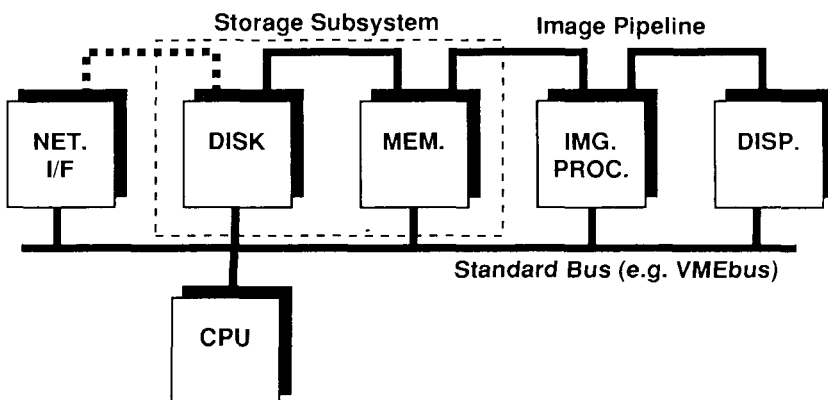


Fig 6. Pipeline oriented architecture.






Fig 7. Timing factors in image transfer. T_s , seek time; $T_{x_{dm}}$, transfer time: disk to memory; $T_{x_{ipd}}$, transfer time: process and display.

even though we have only 12 significant bits of data, we have to move around a 16-bit word. We could compact the data, but unless the compacting/uncompacting function is provided by hardware, we will pay a very heavy penalty in data transfer performance. It is also assumed that image data is stored contiguously on disk to optimize performance so we only require a single seek per image. Because these results are based on maximum available performance, the only way to decrease the time to display a case is to overlap operations. If we can move an image from memory through the processor to the display as soon as the image is transferred from disk, we can bury nearly all of the processing time (Fig 8).


It is important to realize that the values for the various timing factors may be modified by interaction of the components. For example, we may be tempted to simply plug in the manufacturer's value for sustained disk transfer rate and the size of the image to compute $T_{x_{dm}}$. However, we have to make sure that the path connecting the disk to the memory will sustain



Sequential I/O and Processing : $T_{disp} = 6*(T_1 + T_2)$



Overlapped I/O & Processing When $T_1 > T_2$: $T_{disp} = 6*T_1 + T_2$



Overlapped I/O & Processing When $T_1 < T_2$: $T_{disp} = 6*T_2$

Fig 8. Overlapped versus sequential I/O. $T_1 = T_s + T_{x_{dm}}$; $T_2 = T_{x_{ipd}}$.

that transfer rate and that the memory will sustain that transfer rate. Otherwise, as some have found, the realized performance of the system will not match the calculated performance.¹¹

Relationship between storage subsystem components. Now that we have looked at the storage subsystem relative to other components of the workstation and considered some operational issues, we can look in closer detail at the relationship between storage subsystem components. This section will discuss both architectural and operational issues.

Storage subsystem components that support a pipelined architecture clearly offer the greatest potential for providing high throughput. Given such an architecture, what are the implications for hardware and how can we operate that hardware to maximize its performance?

The most obvious implications are the requirements for dual porting. We must provide dual porting of the memory to allow concurrent access of the the memory by the disk controller and the image processor. However, in predicting performance we must be aware that the transfer rates specified in vendor literature may not be accurate for simultaneous accesses at both ports. We may also require dual porting of the disk controller with ports to the memory and to the network interface. This is not as critical of a requirement given our strategy of prefetching image from the archive before scheduled reading. However, the same caveats apply with regard to advertised transfer rates of dual ported disks as apply to dual ported memory.

In addition to how we connect the physical components it is important to carefully consider how we operate them. An important issue is the format of image storage on the disks. To effectively use high-speed paths we must optimize data storage toward fast movement of image data. This means we should separate image data from other information such as text and graphics, and especially operating systems storage such as programs, paging files, etc. The maximum separation can be made by putting nonimage data on a completely separate disk. This prevents contention between disk access for image data and other types of disk accesses that are less critical. Figure 9 shows this separation and indicates the different data paths taken by

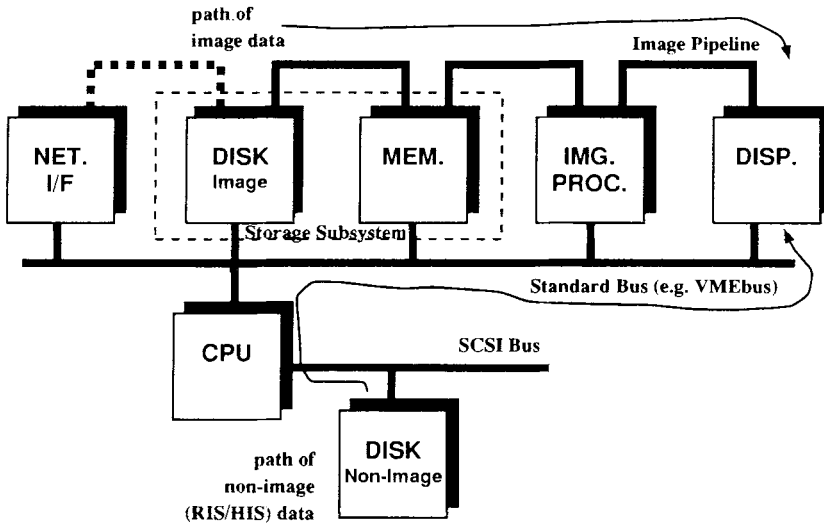


Fig 9. The maximum separation of image data from other information can be made by putting nonimage data on a separate disk. This diagram indicates the different data paths taken by image and RIS/HIS data as they are moved from storage to display.

image and RIS/HIS data as they are moved from storage to display, assuming display of RIS/HIS data on the same display as image data (eg, on graphic overlays). An equally valid architecture might have display of RIS/HIS information on a completely separate, nonimage display. Such a display would not be on the s.p.d. pathway, but would connect to the common system bus.

An additional technique for optimizing image transfer concerns the organization of the data on the disk. Image data should be stored contiguously. This ensures that we can use the full bandwidth of the disk to memory path and will not waste bandwidth waiting for the disk to seek each segment of a noncontiguous file. These techniques are critical to effectively using the bandwidth of the s.p.d. data path. However, they imply that we may not be able to use the general purpose file system for image storage to achieve efficient image transfer.

Relationship to the Archive

Now that we have examined the details of the storage subsystem in the workstation, it is necessary to step back up a level and reexamine the relationship of the workstation to the other components of the IMACS. Given that we now understand the storage subsystem of the workstation, we should consider higher level relationships to the archive system components of the IMACS.

Operational issues. We have chosen a strategy of prefetching images and storing them at the workstation. This strategy provides greater

opportunity to meet response time requirements at the workstation and lowers performance criteria for both the communications network and the archive(s). Now, we must consider how this will impact the operational characteristics of the archive itself.

The strategy will have minimal impact when we treat the workstation and its storage as a simple data sink. The archive is a centralized database. The archive is instructed to send specific images to the workstation, once transferred, the fact that they are available at the workstation is of no concern to the archive. Database management software operates only with knowledge of the local state of the archive. In this case the system and component software is simpler and presumably less expensive. However, a centralized database also presents a potential single point of failure.

On the other hand, the strategy will have maximal impact if we consider the workstation to be an active participant in a distributed database system. In this case, the archive must keep track of the fact that images are residing in several locations around the system, including its local disk storage. The database management may have to operate with knowledge of the global state of the entire IMACS. This requires much more sophisticated (ie, expensive) database management software and movement of data between systems to maintain necessary state information. A distributed database provides greater fault tolerance. However, this strategy would introduce a great deal of unnecessary load on the workstation, particu-

larly because we anticipate very little image update activity originating from the workstation.²¹

Certainly there are varying degrees of interaction between these two end points, in which the archive may maintain partial state information. The choice of storage strategy at the workstation impacts the choices available for operation of the archive, which in turn influences operational requirements of the workstation.

Management issues. The choices we make at the workstation also impact how the archive interacts with the workstation. For optimal image transfer we should choose to segregate image and nonimage information at the workstation. If the archive does not follow such a segregation policy, then separation of the different types of information will have to be done, either at the archive before transmission or at the workstation after receipt. In either case, the associated processing will affect the time it takes to load a case. For prefetched images this time will not interfere with the reading process. In the very rare cases in which an image is called from the archive during reading, the radiologist will have to wait for this process. Maintaining image and nonimage data separately will also make it easier for the workstation to communicate with the RIS/HIS in the situation in which the workstation has direct communications with those systems. This choice also works well for virtual database integration.³

The choice of architecture in the workstation also directly affects data management software on the workstation and indirectly affects the database management software on the archive(s). General purpose database management software will not necessarily be optimal for management of image data under the constraints of segregated, contiguous image data. The data management software on the workstation must be highly tuned to the physical configuration of the hardware. General purpose database software should only be used to manage pointers to image data and should be invisible to the user. Requiring a radiologist to enter sequential query language (SQL) queries to access images would negate optimization of the s.p.d. pathway. However, the data management software will have to understand how to communicate with the archive(s), and will essentially act as a client. The extent of this communication

will depend on whether the workstation is part of a distributed database or just a data sink. Using the workstation as an active element of the archive in a distributed database will introduce a great deal of processing overhead.

An important related, but nontechnical issue, concerns changes in the type of personnel required to support an IMACS. Implementation of computerized image handling systems, whether centralized or distributed, will require a transition from physically oriented film library clerks to higher skilled database administrators. Although the number of personnel supporting the image database may decrease, the cost may not. In addition, there may arise contention concerning which department should operate an IMACS database. This contention may arise between an information management department and the radiology department concerning which department will control the systems associated with the image handling system.

CONCLUSIONS

In this article we have examined the general and specific information environments in which a diagnostic radiology workstation will operate. We have considered the operational goals of an IMACS and it is apparent that the single most critical factor in the success of the IMACS is the ability of the system to support diagnostic reading of images. It is clear that the single most critical component of the IMACS is the workstation. We can choose operational strategies to ease the technical requirements of the archive and the communications network. These choices are based on understanding the functioning of the existing film-based process. This operational strategy of choice is preloading of images on the workstation. This strategy requires that we have about 1 GByte of storage available at the workstation.

Given the low frequency of access to films that are not preloaded, once we have moved the images to the workstation our ability to match the throughput of film-based reading is limited solely by architectural and operational features of the workstation. By examining the s.p.d. pathway from both architectural and operational viewpoints, we can see that a pipelined architecture, with the s.p.d. pathway distinct from general purpose system buses, and the use of overlapped I/O and processing will provide

the raw data rates necessary to meet our throughput goals. This architecture requires the availability of dual-ported disk controllers that can provide a sustained data rate of approximately 15 MBytes/s, and dual-ported memories that can accept data at that rate.

We have also considered local organization and management of the image information required to make the most of the capabilities of the s.p.d. pathway in our architecture. Here we see that contiguous storage of image data is necessary to optimize image movement from disk to memory. Segregation of image data and

nonimage data between physically and logically distinct disk storage is necessary to eliminate contention between image access and disk accesses required for routine system operation. Additionally, we can ensure the optimal use of our architectural design if we use general purpose data management software to manage pointers to image data and use highly tuned special purpose procedures to move images to and from disk. Consideration of global management has shown that we will introduce less overhead to the workstation if we operate as a simple data sink, with respect to the archive.

REFERENCES

1. Kaihara S, Watanabe R: HIS scope, in Bakker AR, Ball MJ, Scherrer JR, et al (eds): *Towards a New Hospital Information System*. North Holland, Amsterdam, 1988, pp 19-24
2. Bakker AR: Patient & medical records. *Proc IEEE Image Management Comm Patient Care: Implementation and Impact*, Washington DC, June 1989, pp 40-44
3. Liu Sheng OR, Garcia HC: Information management in hospitals: An integrating approach. *Proc IEEE Computer Society Phoenix Conference on Computers and Communications*, Phoenix, AZ, March 1990, pp 288-295
4. Attewell P, Rule JB: Computing in organizations: What we know and what we don't know. *Commun ACM*, 27:1184-1192, 1984
5. King JL: Centralized vs decentralized computing: Organizational considerations and management options. *Computing Surveys*, 15:319-349, 1983
6. McNeill KM, Odio R, Sabovik D, et al: Information Plan for the Administrative Support Staff, Department of Radiology, Arizona Health Sciences Center. Semester Project for MIS 541A (Dr E. Sue Weber) December 1989
7. Ozeki T: Difficulties in Design and Implementation of Large Information Systems in the Hospital Environment such as a Picture Archiving and Communications System: Problems and Strategic Solutions. Masters Thesis, Department of Management Information Systems, University of Arizona, April 1990
8. Rowberg AH: Clinical overview of IMACS systems. *Proc IEEE Image Management Comm Patient Care: Implementation and Impact*, Washington DC, June 1989, pp 290-291
9. Panwar RK, Wang CS, DeSoto LA, et al: UW PACS prototype performance measurements, computer model, and simulation. *Proc SPIE Medical Imaging IV: PACS System Design and Evaluation*, Newport Beach, CA, February 1990, pp 869-880
10. Seshadri SB, Arenson RL, Sprague DL: The architecture of an optical jukebox image archive. *Proc SPIE Medical Imaging IV: PACS System Design and Evaluation*. Newport Beach, CA, February 1990, pp 925-931
11. Frost MM, Honeyman JC, Staab EV: Diagnostic workstation development: Initial experience. *Proc IEEE Image Management Comm Patient Care: New Technologies for Better Patient Care*. Kyoto, Japan, April 1991, pp 368-370
12. Choi HS, Park HW, Haynor DR, et al: Development of a prototype electronic alternator for DIN/PACS environment and its evaluation. *Proc SPIE Medical Imaging IV: PACS System Design and Evaluation*. Newport Beach, CA, February 1990, pp 532-540
13. Vercillo R, Fisher III HD, Lamoreaux RD, et al: Digital image review Console. *Proc SPIE Medical Imaging*. Newport Beach, CA, February 1987, pp 708-712
14. Liu Sheng OR, Wang H, Garcia HC: IRES: image retrieval expert system. *Proc SPIE Medical Imaging*, Newport Beach, CA, February 1987, pp 832-841
15. Shannon RH: IMACS and radiology: Defining the problems. *Proc IEEE Image Management Comm Patient Care: Implementation and Impact*, Washington, DC, June 1989, pp 45-53
16. McNeill KM, Seeley GW, Maloney K, et al: Comparison of a digital workstations and a film alternator. *Proc SPIE Medical Imaging II*. Newport Beach, CA, February 1988, pp 872-876
17. McNeill KM, Seeley GW, Maloney K, et al: Comparison of digital workstations and conventional reading for evaluation of user interfaces in digital radiology. *Proc SPIE Medical Imaging II*. Newport Beach, CA, February 1988, pp 872-876
18. Fisher P, Grover B, Brauer G, et al: Digital image display station performance requirements based on physician experience with a prototype system. *J Dig Imaging* 2: 150-155, 1989
19. Kim Y, Park HW, Haynor DR: Requirements for PACS workstations. *Proc IEEE Image Management Comm Patient Care: New Technologies for Better Patient Care*, Kyoto, Japan, April 1991, pp 36-41
20. O'Malley KG, Giunta J: The alternator: Determination of its fundamental features, as a basis for design of a PACS diagnostic workstation. *Proc SPIE Medical Imaging II*, Newport Beach, CA, February 1988, pp 988-994
21. Liu Sheng OR, Garcia HC: The design of medical image databases: A distributed approach. *Proc IEEE Computer Society Phoenix Conference on Computers and Communications*. Phoenix, AZ, March 1990, pp 288-295