# Digital image archive: Using statistical-caching techniques

Jaime Taaffe

**To achieve full implementation of the digital radiology department or of picture archive and communication systems, an efficient digital image archive must be constructed. This paper discusses the use of statistical-caching techniques within an image archive. In addition, the results of computer modeling will be used to demonstrate the advantages of this approach.**

Many radiologists believe that picture archive and communication systems (PACS) and therefore the "digital department" will soon emerge. With the increase in the percentage of digitally formatted imaging techniques, namely computed tomography, digital subtraction angiography, magnetic resonance imaging, positron emission tomography, nuclear medicine studies, ultrasonography, and especially digital radiography, it becomes natural to present the acquired images at digital workstations for diagnostic viewing and manipulation.

With the increased use of medical images, severe problems have become evident in current film-jacket–based image-management systems—mainly long delays in obtaining old films and lost or unavailable films. With a PACS, the inherent ability to electronically duplicate an image residing in the archive and then to transmit it to a workstation elegantly solves these problems. Images are quickly located through an on-line image name-and-attribute database. Images are not lost, since they never leave the archive: only copies do. Simultaneous viewing by more than one physician is achieved through quick and efficient image duplication.

Numerous reasons and arguments can be presented for the introduction of PACS into diagnostic radiology. The two reasons presented above, ie, the increased production of digital images and the problems inherent in film-jacket systems, highlight a major hurdle that the designers and implementors of PACS must solve: the creation of an electronic image archive. This archive must be capable of storing, finding, and conveying millions of images quickly and efficiently to other parts of the PACS. Simultaneously, the archive must catalog and store new images as they are acquired.

The difficulties facing the designer of a PACS and its associated archive are formidable. Existing image-management systems, albeit film-based and usually manual, are well developed and tuned. For example, at the Massachusetts General Hospital, cases to be viewed by the diagnostic radiologist are placed on alternators along with relevant comparison images by film librarians, so the radiologist need only press a button to be presented with the next exam's films and comparison films. From the point of view of the radiologist, the system is quick and efficient, and neither the patient's name, current exams, nor comparison exams need be specified. The system and the film librarians have conspired to keep the radiologists focused on their primary task, viewing cases, rather than on hunting down films or placing the films on light boxes. The response time of this system is dramatic: numerous images representing megabytes of data are displayed in less than a second. In addition, the general trend toward cost containment and reduction imposes a severe restriction on the cost of the PACS and its archive.

## Demands placed on a digital image archive

In a teaching hospital such as MGH, two distinct groups will place demands on the digital image archive: the clinical users and the research users.

Clinical requests are made on a patient-by-patient basis and consist of the current exam or exams, plus any related exams that can be used for comparison. In this paper, the collection of a current exam and its comparison exams will be called a viewing-complex.

Unlike clinical requests, which entail a number of images from a particular patient, research requests are made usually for a collection of images whose commonality may be type of image, pathologic condition, and/or type of patient.

It is reasonably assumed that the number of retrieved images for research will be smaller than that required for clinical needs. Furthermore, it is assumed that research requests need not be serviced as expediently as clinical requests. For these reasons, research requests are ignored in the considerations that follow.

Clinical demands are characterized by two temporal recall patterns, inpatient and outpatient; for inpatients, multiple requests (assumed three to six) for the viewing-complex will occur within the first 24 to 36 hours after the exam. Requests originating from a particular workstation will usually be serial progressions from the first unread case to the last. For example, all spinal CTs performed that day may first be read by the responsible resident. Later, these cases will be read, still in serial order, by a staff radiologist. The ordering of the exams is arbitrarily assigned and may be made in the order in which the exams were completed or in some other manner of priority. After the initial readings of viewing-complexes, principally by the radiology department, the images will be viewed in a nonpredictable fashion by others outside the department (assumed three to five), such as the attending and referring physicians. Before an inpatient's discharge, requests (assumed one to three) may be made to review all viewing-complexes associated with the patient. The recall rate for an inpatient's images declines rapidly after discharge.

*Mr Taaffe is associate director of image research, Massachusetts General Hospital, Department of Radiology, Boston.*

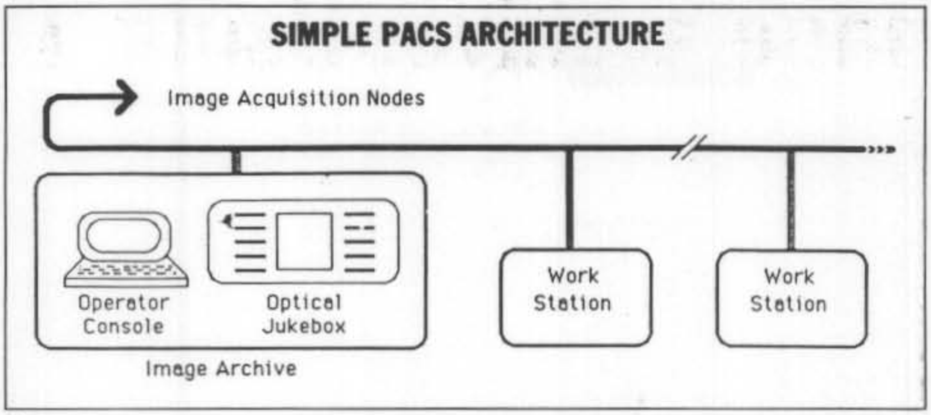


*Figure 1.*

The temporal recall pattern for an outpatient is similar to that of an inpatient in that there are intense and highly predictable initial recalls (assumed two to four) during the diagnostic period, followed by a moderate number in a nonpredictable fashion. There will be no discharge activity, nor will the long-term activity be as high as that for patients who required hospitalization.

In summary, the following assumptions can be made. The number of requests for inpatient viewing-complexes is assumed to be greater than that for outpatients, a large proportion of the requests will occur within 1.5 days after the exam is completed, and many of the initial requests will be part of a serial reading of new exams.

### Archive components

The designer of an image archive has available a wide array of components from which to construct the archive. Table 1 lists the major devices available, along with a few significant parameters. The optical disk technology listed is of the Write Once Read Many (WORM) variety.

Currently available optical media cannot be erased and reused, unlike conventional magnetic disks. The VHS digital tape unit listed is currently being used in the MGH radiology PACS lab and is manufactured by Honeywell Inc of Denver.

Capacity refers to the size of the storage device in megabytes and reflects the number of images that may be stored within. For example, the popular optical disk drives store 3,000 megabytes per disk (or platter). A 512×512 CT image 12 bits deep requires approximately 0.4 megabytes. Thus, more than 7,500 CT images can be stored on one disk. This figure could be increased by using image compression techniques. By way of comparison, if every radiologic image acquired by MGH in one year was stored on digital media, without compression, it would require 4,700,000 megabytes (4.7 terabytes).

Cost per megabyte of storage becomes of great importance, considering the enormous quantities of data. Stored on today's optical disks, the media cost for 4.7 terabytes, one year's storage only, would be $235,000. The equivalent cost of storage using VHS digital tapes would be $9,400.

Access time refers to the time elapsed from the instruction to fetch the desired data to the availability of the data. Access time directly affects how quickly a workstation displays a particular image after being instructed to do so. Data rate reflects how fast the device can transfer the image. The greater the data rate, the faster the transfer.

The task of the PACS system designer is to construct an archive with fast image access, fast transfer, enormous capacity, and low cost using various image-storage devices, processing elements (CPUs), and suitable algorithms.

### Archive architecture, demand driven

Commercially available optical "jukeboxes" capable of storing as much as one terabyte are being offered for use in PACS; such optical archive architecture is shown schematically in figure 1. In this architecture, images are transferred on demand from the jukebox to the workstations over a Local Area Network (LAN). If the image does not currently reside within the jukebox, a request is made to an operator to place the needed disk into the jukebox from off-line storage.

The architecture is acceptable with respect to access time, provided the image is on the disk currently loaded in the disk reader. It is less acceptable if the disk needs to be loaded into the reader and quite slow if the disk must be loaded into the jukebox by a human operator.

The probability that the expected images will be located on only one disk—the disk in the reader—can be increased by placing all new images on one disk and copying all comparison images also on this one disk. However, a very high price in dollars must be paid for this repeated copying to new disks and trashing of old disks, which must be done daily (because of the inability to rewrite optical media).

The transfer rate of the optical disk reader becomes less than ideal, considering the size of digital radiographic images (six to 25 megabytes). As fast (100 megabits per second and higher) optical LANs come into widespread use, the data-transfer speed mismatch between them and optical media may become bothersome.

### Archive architecture, knowledge driven

As shown in table 1, cost increases for storage devices with faster access times. An archive consisting solely of solid-state memory would have exceptionally fast access but astronomic cost. An archive based on VHS tapes alone would be low-cost but would have unacceptably slow access times.

Given the semipredictable pattern of image retrieval within the hospital environment and the obvious relationship between storage device cost and performance, the following question becomes conspicuous: Can the image archive be constructed so images are moved from slow, inexpensive storage to faster access storage in anticipation of their actual use, as a means of obtaining high performance at an affordable cost?

The key to achieving this goal is the implementation of algorithms that anticipate what images may be requested and when. In the process of rendering patient care, many events take place that indicate the future use of radiologic images. A few of these events will be discussed.

Scheduling a patient for radiologic exams indicates strong probability for prior exams to be recalled. Obviously, the completion of an exam indicates impending recall of the viewing-complex. If a radiologist is working through a stack of exams, then the reading of the Nth exam indicates, with high probability, the retrieval of the (N+1)th viewing-complex. Viewing of the first image of an exam consisting of many images, as in CT studies, for instance, is an indication that the others are almost certain to be immediately requested. Conversely, many events indicate the low future probability of a particular image request from the archive; inpatient discharge and exam-request inactivity are two obvious examples.

Figure 2 shows an archive segmented into progressively higher-performance elements. Starting on the left is inexpensive storage, perhaps off-line VHS tapes. The middle element may be a jukebox that allows some number of tapes to be kept on-line. The high-speed element on the right could be a high-performance magnetic disk drive. The events mentioned in the previous paragraph can be used to shuffle image data between the various elements to achieve fast average access times. Specifically, requests and anticipated requests will cause image data to move from the slower storage elements to the faster. To free up valuable space needed in the higher performance storage elements, inactivity and anticipated inactivity will cause image data to move back to the slower devices. The exact access time of an archive of this type obviously cannot be expressed as a single number but only in terms of a statistical distribution whose moments, mean, standard deviation, etc. are a complex function of the parameters of

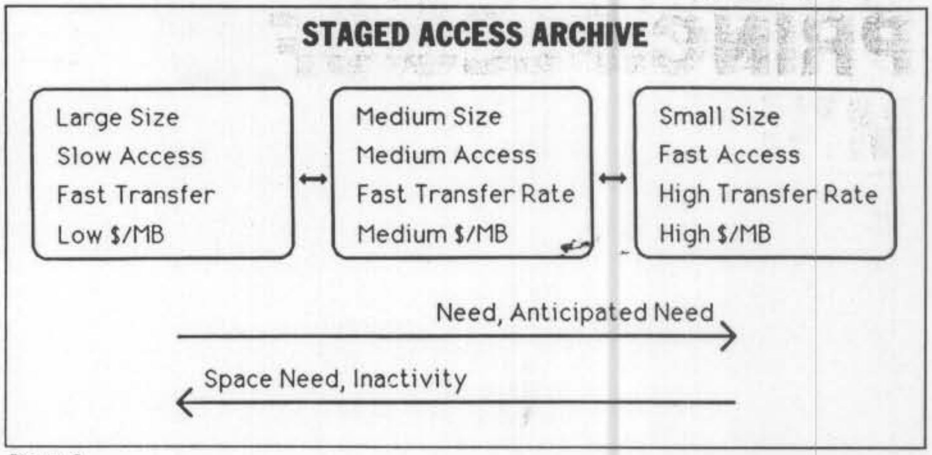


*Figure 2.*

the storage elements, algorithms used, and the pattern of image requests. The lower bound on these moments, given the image-demand distributions, can be set by the characteristics and size of the higher performance elements. This technique is usually called statistical caching, and one method of analysis for this particular architecture is to model it as a so-called Markov process.

## Architecture simulation

Designing a PACS is difficult in itself; designing one using distributed statistical caching techniques is more difficult because no analytic expressions describe its behavior. It is well known that computer systems, LANs, and storage devices exhibit nonlinear responses with respect to load. This nonlinearity, coupled with the inadvisability of expressing the physician's demands on the system as simple periodic requests, requires that the entire PACS be modeled and simulated on a computer.

NETWORK II is a simulation program (CACI Inc, Los Angeles) that takes a user-specified description of the computer hardware and software, then provides measures of hardware utilization, software utilization, and process execution times. Specifically, in NETWORK II the various processor elements, along with their instruction repertories, are described. Storage elements representing disks, tapes, buffers, etc, are specified in terms of access times, sizes, and bandwidths. The interconnecting buses linking processors and storage elements are described in terms of bandwidths and protocols. The software modules to be simulated are described by using the instructions defined for the various processor elements. In addition, the events and conditions that cause various software modules to execute are specified. Accurate measures of system performance can be obtained by allowing the program to simulate many transactions, workstation requests, and new-image filings. (This program was run on the radiology department's VAX computer.)

## A simulation example

In distinction from the architecture depicted in figure 1, where the archive is a self-contained node connected to the LAN, figure 3 shows an architecture in which the archive starts from off-line media stored on shelves and progresses to small high-speed solid-state image buffers within each workstation. In this architecture, ten workstations are connected to the image archive by a high-

| Device capacity (Mb) | Cost per Mb | Access time (msec) | Data rate (Mb/sec) |
|---|---|---|---|
| **Fixed media devices** | | | |
| Solid-state memory based on random-access memory chips | | | |
| 1 to 32 | $200 to $500 | < 0.001 | 5 to 120 |
| Parallel-transfer Winchester disk drive (video or PTD disk) | | | |
| 670 | $30 to $40 | 40 to 80 | 10 |
| Single-head fixed Winchester disk drive | | | |
| 100 to 1,000 | $5 to $20 | 40 to 80 | 1.2 to 2.0 |
| **Removable media devices** | | | |
| Optical WORM disk | | | |
| 3,000 | $0.05 | 50 to 100 | < 1 |
| VHS digital magnetic tape | | | |
| 7,500 | $0.002 | 5,000 to 120,000 | 4 |
| Standard 12-in reel 6250 BPI magnetic tape | | | |
| 50 | $0.20 | minutes | 1 |

MB, megabyte (1,000,000 bytes); PTD, parallel transfer drive

*Table 1.* ***Possible image-storage building blocks: Performance and cost***

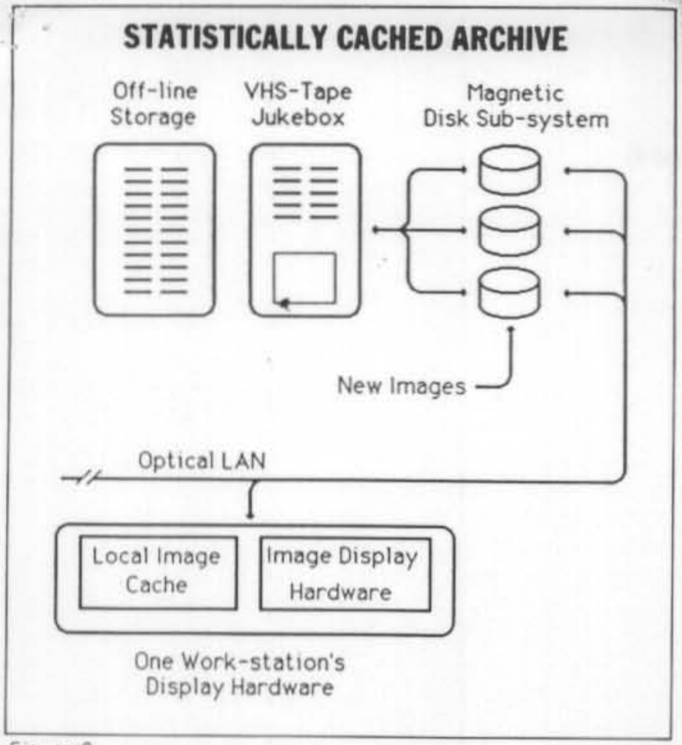


*Figure 3.*

speed optical LAN. Each workstation contains local high-speed solid-state memory being used as a local image cache. The image server consists of the following series of storage elements:

- Off-line VHS magnetic tapes
- VHS-tape jukebox
- Medium-performance magnetic disk subsystem
- High-speed, solid-state image cache

For this particular simulation, the gross attributes of the various elements are listed in table 2. Note the ability to express access times in terms of statistical distributions. Table 3 describes the load placed on the system because of users requesting viewing-complexes, new images entering the system as they are acquired, inactive images moving to slow storage (VHS magnetic tape), and images moving to faster storage elements in anticipation of their use.

The software model is one in which a user periodically requests a viewing-complex. On receiving the request, each workstation is assumed to be able to service 40% of the requests from its local image cache, due to statistical caching. For the 60% of requests that cannot be serviced locally, the workstation then sends a message via the LAN to the image server requesting the viewing-complex. The image server determines where the viewing-complex currently resides. It is assumed that 30% of the requests will be serviced from the server's image cache and the remaining 70% from the magnetic disk subsystem. On locating the viewing-complex, the server transfers the data into the LAN controller and initiates their transmission, via the LAN, to the requesting workstation.

As a background task, the image server is accepting, cataloging, and storing new images as they are acquired. Last, also as a background task, the server is moving images from one storage device to another, including the workstations, in an attempt to anticipate their use.

(This simulation illustrates one of many architectures modeled by us. The details of the model given here are brief and simplistic because to fully describe the details of the model and the simulation software, many pages would be required. This particular model alone required more than 40 pages of simulation language to describe.)

The simulation was run for 30 minutes, in the target system's time frame, and yielded the following values for waiting time: average 81 msec, minimum 30 msec, maximum 340 msec, and a standard deviation of 75 msec. These times reflect the total time to display an image, including processing, lookup, and transmission times.

Workstation's processor instruction rate = 4 MIPS
Image server's processor instruction rate = 5.7 MIPS
Image server's magnetic disk subsystem
  Access time = 20 to 30 msec (uniform distribution)
  Transfer rate = 6 MB/sec
Image server's solid-state cache
  Access time = 0.1 msec
  Transfer rate = 40 MB/sec
Image server's internal bus rate = 40 MB/sec
LAN data transfer rate = 100 megabits/sec

MIPS, millions of instructions per second; MB/sec, megabyte/second; LAN, local area network.

*Table 2.* ***Simulated system's hardware performance***

**Because of the users**
Ten workstations, each requesting a viewing-complex every three to five minutes (random uniform distribution), where
- —a viewing-complex contains 1 to 32 images
- —an image consists of 1,024 × 1,024 × 12 / 3.1 = 4 megabits
- —the factor of 3.1 is due to image compression.

**Because of caching**
Each workstation receives an image every ten to 15 seconds.
The image server transfers an image from the disk to the image cache every one to two seconds.
The image server transfers an image from the image cache to the LAN controller every one to two seconds.
The image server transfers an image from the disk to slower storage every one to two seconds.

**Because of new images being acquired**
Incoming new images are transferred onto the disk every one to two seconds.

*Table 3.* ***Simulated system's work load***

## Conclusion

From the results of the simple model presented, it is evident that there is much to be gained from statistical caching. Statistical caching enables the use of seemingly slower and significantly less expensive storage devices, augmented with small, high-performance caches to achieve a high-performance image archive at reasonable cost.

Simulation provides a method for checking the design and a means by which numerous architectures may be rapidly evaluated. This ability to "mock up" a proposed system aids insight and reveals behaviors and interactions within the system that could not otherwise be predicted. Thus, simulation of system architecture, software, and algorithms is a powerful tool for the PACS designer. **AR**


APPLIED RADIOLOGY, January 1988