Data Clustering and Other Archive Retrieval Strategies for Teleradiology and Picture Archiving and Communication Systems

Bruce Ho, Woodrew Chao, Reza Sadri, Lu Huang, Ricky Taira, and Henry Shih

A key advantage in the conversion from film-based to digital radiology is the possibility of a long-term on line electronic archival of patient studies. The popular approach based on optical disk jukeboxes for the long-term archive and magnetic disk storage for data caching is not economically attractive because of the cost of both the jukebox and the medium. Strategies for extending the archival system design with a tape jukebox have been studied. The proposed strategy calls for the use of high-ratio lossy compression together with low-cost tape storage to make long-term on line archiving more affordable. An intelligent prefetching algorithm based on hospital information system and radiologic information system triggers, which in turn are augmented by manual case preparation, can effectively overcome the longer latency of ad hoc retrievals. This longer latency is caused by both system-level bottlenecks and the sequential access constraint of the tape drive. Strategies for image clustering and tape allocation by patient classification also enhance retrieval efficiency. This archival design using image compression, prefetching, and clustering could be implemented in many of the existing teleradiology and picture archiving and communication systems.

KEY WORDS: teleradiology, picture archiving and communication systems (PACS), archive, hospital information systems (HIS), radiologic information systems (RIS), database, image management.

THE USE OF picture archiving and communication systems (PACS) is increasing with many leading institutions opting for large-scale installation rather than a phased modular approach.[1-7] Archival for digital radiologic data remains an unsolved problem because of slow retrieval speeds and sheer volume. Radiologists require immediate access to patient studies during clinical viewing. A medical center routinely generates a large volume of image data, which must be stored for many years depending on legal requirements and other factors.

The goal of this study is to evaluate a method to extend on line archive size at lower cost and acceptable speed by the use of a slower storage device coupled with several software strategies, such as patient classification, prefetching, and compression. Intelligent image-management algorithms make possible the use of alternative tape devices that are slower but more economical than the commonly used optical disks (ODs).

EXISTING PACS ARCHIVE DESIGN

Today, many existing teleradiology systems and PACS rely on OD jukeboxes for long-term on-line archival, attracted by their durability, format standardization, and random accessibility.[8] The relatively slower access of the OD (because of data seek time, robotic mounting of platters into the drive, and data transfer rate) can be offset by a magnetic disk-based cache. The cache stores the most urgent (likely to be used) data in anticipation of clinical use. After a pre-set period of time, the data is relocated to the slower OD. This downstream data movement from a faster to a slower device is called migration.[9]

The data aging approach works well for recent studies. However, it does not address efficient retrieval of comparison studies. The upstream movement of data is called image caching. In radiology, this can be best implemented based on triggering events from the radiologic information system (RIS) or the hospital information system (HIS), which effectively predict the radiologists' pending need for comparison studies and cause them to be prefetched from the OD. For example, when a returning patient checks in for a new study, the archive management software can prefetch the

From the Department of Radiological Sciences, AR-264, Center for Health Sciences, University of California, Los Angeles.

Supported by National Institutes of Health Grant No. PO1 CA51198 "PACS in Radiology"; and the Department of Radiological Sciences (B.H., R.T., W.C., L.H., H.S.) and the Department of Computer Sciences (R.S.), University of California, Los Angeles.

Address reprint requests to Bruce Ho, PhD, Department of Radiological Sciences, AR-264, Center for Health Sciences, 1083 LeConte Ave, University of California, Los Angeles, CA 90024.

0897-1889/95/0804-0001$3.00/0

comparison studies and have them at the diagnostic viewing station well before the new study is completed.

The PACS archive at the University of California Los Angeles (UCLA) is a classical example of this two-tier architecture. We use two Eastman-Kodak (Rochester, NY) Automated Disk Library (ADL) OD jukeboxes, which use 12-inch OD platters with 10 Gbyte capacity each. The design is illustrated in Fig 1. There are two archive centers, each with enough magnetic disk (MD) capacity to store approximately 2 weeks of current image data. The MDs provide fast response for the most recent clinical data. Aged data from each MD storage migrates to a ADL OD jukebox with a total capacity of 1 Tbyte. The two ADLs can accommodate about 2 years worth of images at UCLA (no compression is used at this time). After this time, fully consumed OD platters are manually removed for shelf storage to make room for new platters. The permanence of the ODs assures the availability of old patient studies for rare but critical situations such as medicolegal disputes.

However, ODs have some important disadvantages. First, the user must bear the same medium cost even though ODs only stay on line for a short part of the required archive period (which can be more than 10 years in pediatric cases). Second, at 7 cents/Mbyte, it is doubtful that the OD solution actually realizes any financial incentives compared with film storage.[8] Using fair market prices, an archival solution like UCLA's translates into almost $1,000,000 in acquisition cost and around $50,000 a year in OD supply, a draconian expense even for large hospitals. An alternative like the Exabyte (Boulder, CO) tape costs only around $1,400 a year at the same volume. Other solutions include 4-mm digital audio tape (DAT), 19-mm helical scan tape, digital linear tape (DLT),[10] and writable CD-ROM.

EXTENDING THE EXISTING DESIGN WITH A TAPE JUKEBOX

While offering an order of magnitude in cost savings, the tape device can not directly substitute for the ODs because of its sequential access property. A long seek time which is incurred when the tape winds to the beginning of a requested file means a severe penalty in retrieval performance and loss of valuable time by clinical staff. However, the tape jukebox is an excellent solution as a long-term on-line storage after an intermediate storage period on the OD, because the clinical need for image data after the initial few months or so drops sharply (see Results section). The longer retrieval latency would then have much less impact on clinical operation.

Theory of Operation

Adding a tape-based modality to the existing design creates a three-tier archive system. The theory of operation for the three-tier design is illustrated in Fig 2. The two types of jukeboxes can either connect by an Ethernet local area network (LAN) (as in this experiment) or share a small computer system interface (SCSI) bus for faster transfer. In teleradiology systems, the MD should be located near the radiologists' viewing station, generally far away from the archival site where the jukeboxes reside. The connection between the MD and the archival

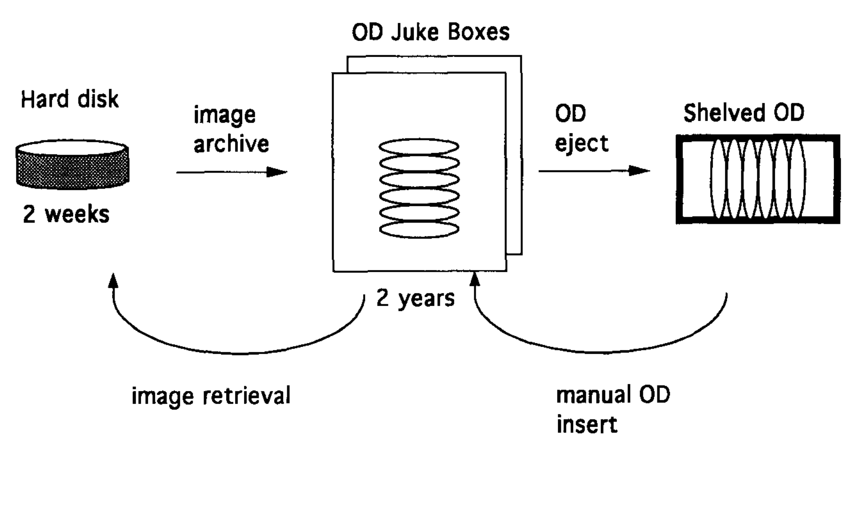

Fig 1. A common PACS archive design is based on one or more OD jukeboxes backing up some MDs. Used ODs are shelved when the jukebox overflows. Costs of both the ODs and OD jukeboxes make this design expensive.

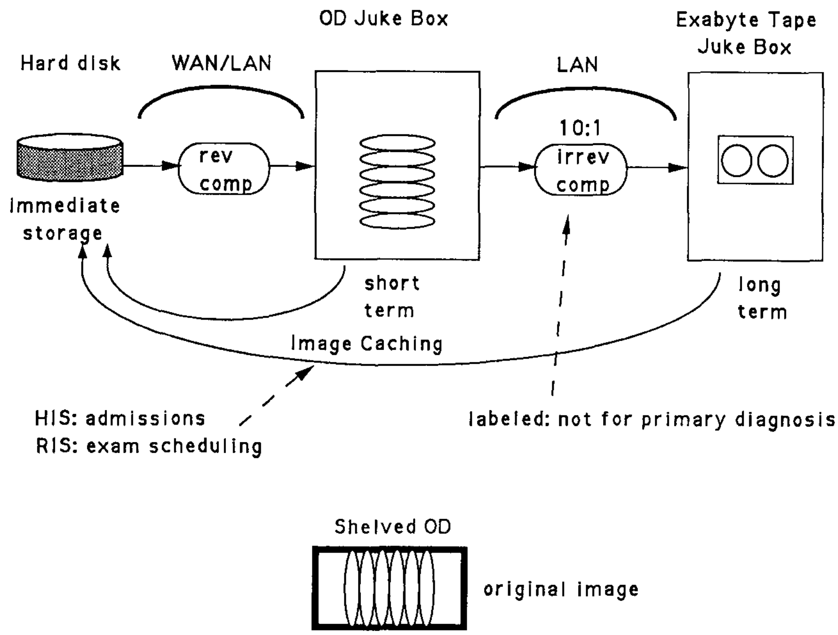

Fig 2. Proposed three-level hierarchical storage with MDs, ODs, and tapes. ODs provide short-term storage of reversibly compressed (rev comp) images for approximately 1 year, whereas the tape provides long-term storage of irreversibly compressed (irrev comp) images for the entire required archival period. This design eliminates the need for expanding the OD archive to achieve long-term on-line retrieval capability. An intelligent image prefetching algorithm is used to cache images from the OD and tape to the magnetic disk as soon as a patient has new examinations scheduled. The connection between the MD farm and the OD jukebox is a WAN for a teleradiology system or a LAN for a PACS.

site will then be via a wide area network (WAN). Reversible compression (see below) is applied to OD stored images that remain in the jukebox for a short-term period of approximately 2 years. Before an OD is ejected from the jukebox, its contents would be irreversibly compressed and stored on tape for the legally required period. The high ratio achieved by irreversible compression brings the jukebox capacity to about 5 Tbyte per jukebox, sufficient for many years of on-line storage.

When old images become needed again (because of return visits of patients), they are copied from the slower tape or OD media onto the faster MD in an image-caching process.[9] Automatic image migration and prefetching algorithms are the key strategies necessary to designing a low-cost and yet efficient long-term archive. The three key technologies needed in this hierarchical design, image compression, image prefetching, and data clustering, are discussed below. Whereas, the first two topics are already well known from earlier efforts by the authors, data clustering is a new concept that only became important with the introduction of the tape archival design.

Image Compression

Two basic methods of compression are used in this hierarchical archiving strategy: reversible and irreversible. Entropy coding techniques such as arithmetic coding (AC),[11] Lempel Ziv Welch (LZW),[12] or the classical Huffman coding can be used to compress image data reversibly by approximately 2.5:1.

Previous work[13] by the authors shows that AC implemented by custom software produces the best compression result. A 2,000- × 2,000- × 10-bit computed radiography (CR) image can be compressed by a factor of 3.4 by AC, compared with a factor of 3.25 by LZW and 1.7 by Huffman coding. This reversible compression step is intended for images stored on the OD to expand the short-term archiving capacity. It will soon be implemented clinically in the UCLA teleradiology system.

Irreversible compression will be applied to images before they are migrated to tape. These irreversibly compressed images have accompanying primary diagnostic reports and will be used in the future only for comparison and research purposes. Excluded from consideration are the most quality-sensitive studies such as mammograms and bone radiographs, for which even digitization, let alone irreversible compression, has not been shown to be clinically acceptable. In practice, there is currently no consensus on the compression method or the acceptable compression ratio. A conservative approach would be using irreversible compression only for review and not for primary diagnosis.

In the past few years, two types of block artifact free compression methods were developed at UCLA as substitutes for the industrial standard Joint Photographic Experts' Group.[14] One is based on full-frame discrete cosine transform[15-17] and the other on discrete wavelet transform (DWT).

The former method was supported by receiver operating characteristics (ROC) studies, which showed that bone and chest radiographs retained diagnostic accuracy at compression ratios up to 20:1.[18,19] Although good ROC results do not necessarily make irreversible compression acceptable for primary diagnosis, they do serve as appropriate guidelines for reviewing compressed comparison studies.

The DWT method has not yet been evaluated by ROC, but has been shown to surpass FFDCT in objective quality comparisons. This is the compression method applied in these experiments. For computed tomography and magnetic resonance images, ratios around 8:1 and 4:1 have been found acceptable. The DWT algorithm is computationally efficient and requires no special hardware. Using optimized C code running on a 36-MHz SPARC 10 workstation (Sun Microsystems, Mountainview, CA), the compression and decompression process of a 2,000- × 2,000- × 10-bit CR image can be performed in 30 seconds.[13] Because the algorithm is easily made parallel, the processing time can be proportionally reduced using a multiple-processor workstation.

Image Prefetching

Prior patient images are usually needed by radiologists for comparison with new studies. When prior images are retrieved from the archive in response to an ad hoc request, delays occur from loading the storage medium into the drive, transferring data, decompressing the image, and various file-server bottlenecks and network congestion. The study of the UCLA PACS system based on network manager software also shows that the file retrieval delay from various system inefficiencies can exceed 5 minutes and can extend to hours in severe cases.[20,21] This delay compounded by sequential access of the tape makes direct retrieval from a three-tier archival system unacceptable for clinical viewing.

Comparison images can usually be accessed in a predictive fashion using an image prefetching mechanism.[22,23] The most basic prefetching rule extracts the most recent prior study of the same modality and scan setting (T1, T2, with/without contrast, etc). A more sophisticated algorithm is under development using key clinical events such as admission, patient arrival in the radiology department, scheduled examinations, therapy administration, or surgery, etc. These expert rules will continue to evolve for many more years. The event information needed for automatic prefetching is made available to the teleradiology/PACS system through the interface with HIS and RIS.[24] A patient is usually scheduled for readmission or for an examination days or hours before the study is read by radiologists. This lead time is usually sufficient for prefetching of prior images.

Given the nonuniform way radiologists perform diagnostic reading, an 100% accurate, automatic prefetching rule set that satisfies every radiologist everywhere is not likely to emerge soon. Manual prefetching by film librarians is still the most reliable method for setting up a reading session. A film librarian who used to prehang relevant films on an alternator can use the same set of guidelines to prefetch and review automatically prefetched soft-copy images. The two-year experience with manual session preparation in UCLA's pediatric radiology section showed nearly 100% success rate in predicting the comparison studies. All of our PACS stations have a graphical interface that allows image retrieval from the OD archive.

Data Clustering During Image Migration

Image migration occurs in two stages: from MD to OD and then OD to tape. Incoming images are written to the ODs in a first-come, first-serve manner. No effort is made to cluster images of a given patient. Sequential recording onto a single OD platter streamlines the archiving process at the expense of multiple platter insertions during future retrievals. This is a logical trade-off because (1) the probability of image retrieval is low compared with the certainty of archiving, and (2) ODs are random access devices, so the penalty for retrieval from multiple platter is only 15 seconds of loading time per platter.

OD to tape migration follows a different scenario. Because a tape is a sequential access device, the long seek time as well as extra cartridge mounting time must be taken into account. Therefore, all images for a single patient should be clustered onto one tape whenever possible. If a patient's images are spread over N tapes, the worst case seek time becomes N × (load time + seek time). However, if one could predict how much tape space to reserve for a patient, his future images can be recorded immediately adjacent to the current one, and yet the same tape volume can still be shared with other patients. If his entire collection of images are recorded contiguously on one tape, there will then be only one load and seek time incurred whenever this patient's folder is retrieved.

To facilitate this prediction, an image distribution model must be developed which classifies a patient into one of three categories (A, B, or C) depending on his total collection of examinations 2 years after the first study. The underlying assumption is that a relatively sick patient, ie, one who has accumulated a larger number of examinations, is more likely to return for more examinations.

In anticipation of the returning pattern of each type of patients, category A tapes reserve less space per patient and hold more patients per tape than B or C tapes (see Fig 3). When a patient has his images transferred from OD to tape for the first time, the next tape of his category with sufficient space is selected. If he already has existing images on a tape with enough allocated space left, that tape is selected; when the allocated space is used up, another tape is selected for his new and future images. This allocation rule will tend to equalize the utilization on all the tapes in each category, thus in turn minimizing the chances for scattering any patient's images. Note that because OD data are already 2 years old by the time they are moved onto tape, someone who just had his first study is not classified until 2 years have elapsed. Therefore he represents a person in relatively good health who accumulated only a few examinations in 2 years.

From these general rules, the specific parameters appropriate for the UCLA system and patient population can be derived. In the experiments described in the next section, a data clustering model was developed based on actual patient study profiles extracted from the existing PACS database.

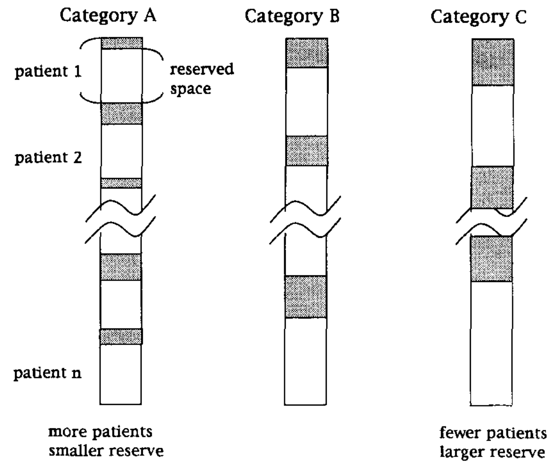

Fig 3. For tape archiving, all the data for one patient is ideally stored contiguously on one tape. This can be accomplished by reserving space immediately behind the current record for future data. The amount of space reserved is predicted based on the patient folder size accumulated over the past 2 years or so. A category A tape holds more patients per tape with smaller reserved space per patient. More space per patient is reserved for category B, and still more for category C.

MATERIALS AND METHODS

While image compression and prefetching algorithms were established in earlier works,[13,23] data clustering and the tape jukebox performance are new items that require verification.

Tape Jukebox Integration

The commercial product used in this experiment was the Exabyte tape jukebox EXB-120. The EXB-120 has 116 slots arranged in six rows for a total capacity of 580 Gbyte of uncompressed data. The helical scan technology allows the tape (5 Gbyte each) to move at a slower speed to reduce stress.[10] Two tape drives are installed with room for two more to perform read and write functions at the same time. The jukebox is attached to a Pentium-based host running the New Technology operating system (Microsoft, Redmond, WA) through a SCSI connector. Communication with the jukebox is implemented through a set of software library calls including functions that control the robot arm to move a single cartridge between a drive and the slots. A database table keeps track of the tape cartridge volume number for each image file and the location of each cartridge. Unique bar-coded labels are attached to the cartridges for identification.

Profiling of PACS Database

The entire PACS database (45,000 patients) was analyzed to find the range of study folder sizes collected over the past

5 years. These patient data were extracted manually using structured query language script instructions. The clinical situation was simulated by copying about 300 of these patient cases to an experimental database. Selected patient images were compressed and loaded onto tape using a classification scheme derived from the patient study profile. RIS triggers were simulated manually to initiate data prefetching from the tape.

Using the experimental image data subset, three software algorithms were integrated and tested: image compression, clustering, and prefetching. The overall objective was to test whether a slower less expensive tape drive modality could provide inexpensive, longer-term on-line archiving.

RESULTS

EXB-120 Performance

The cartridge transport and loading time, fast forward speed, and data transfer rate of the EXB-120 tape jukebox and drive were fully measured (Table 1).

The tape retrieval time consists of moving a cartridge from a slot, inserting it into a drive, fast forwarding it to the beginning of file, and performing the actual data transfer. The overhead time before data transfer starts is equal to (20 + 30 + 0 to 180) = 50 to 130 seconds. This overhead minus the 15-second loading time of an OD in the Eastman-Kodak ADL jukebox is the expected difference in average retrieval times from the two types of media.

Patient Classification

Only a tiny amount of patient data was accumulated in 1990. For the purpose of this analysis, 1991 is considered the first year of the full-scale archive. The database survey result shows that the data volume collected per year increased steadily since 1991, until reaching 0.7 Tbyte in 1994. At this rate, one tape jukebox can easily accommodate up to 5 years (5.8 Tbyte) of image data with an average compression ratio of 10:1.

Figure 4 shows the number of total and new

Table 1. Measured Performance Characteristics of the EXB-120

Action/Characteristic	Performance
Move/relocate cartridge	20 s
Load (insert) a cartridge	30 s (maximum)
Unload a cartridge	20 s
Rewind	180 s (maximum)
Fast forward to beginning of file	180 s (maximum)
Sustained transfer rate	500 Kbyte/s

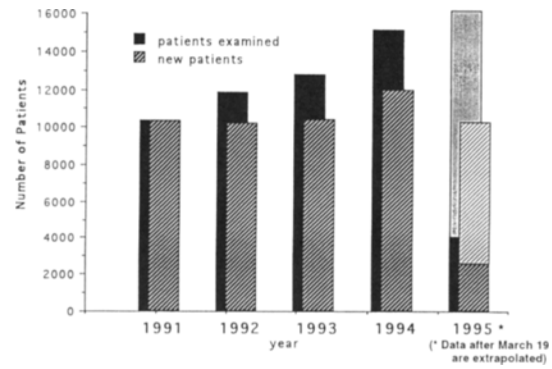

Fig 4. The number of patients per year up to March 1995. Data beyond March 1995 is extrapolated linearly. The 1990 data is disproportionally low because data archiving was not fully implemented yet that year. Both total and new patients are plotted. Around 20% of the patients examined each year are returning patients.

patients examined each year from January 1991 to March 1995. The 4-year trend from 1991 through 1994 shows an 8% to 18% increase in patients examined yearly, only 20% of which are returning patients. By extrapolating the data to the end of 1995, the 5-year patient total can be expected to reach around 53,100 (assuming linear distribution).

In Fig 5 and 6, patient data are placed into bins based on the number of studies in their folders. The bins form finer divisions of the patient population from which the large partitions of categories A, B, and C can be defined. Figure 5 shows the number of patients in each bin, whereas Fig 6 shows their total data volume. The classification is set at 1 through 10, 11 through 30, and more than 30 images for catego-

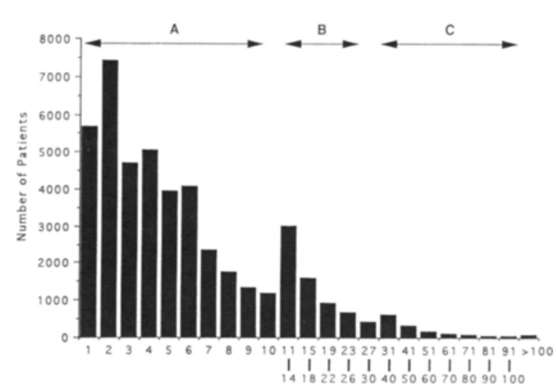

Fig 5. Patient data distribution from our PACS database. The number of studies per patients are plotted. The single largest group of patients had two studies each. People with a large number of studies are rare.

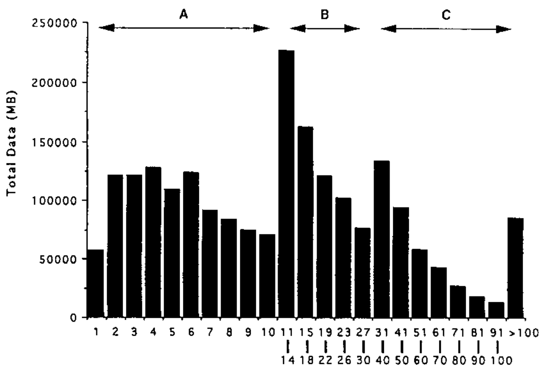

Fig 6. Total image data for each folder size bin are plotted. Categories A, B, and C consisting of patients with 1 through 10, 11 through 30, and more than 30 studies make up 45%, 32%, and 23%, respectively, of total data.

ries A, B, and C patients, making up 45%, 32%, and 23% of the total patient data volume, respectively.

Based on this patient study distribution, the total number of patients (in 5 years) can be set to 43,692, 7,653, and 1,578 for category A, B, and C, respectively (see Table 2). Out of a total of 116 tape slots, 64, 34, and 22 are assigned for the three classes based on the approximate percentage of total data volume each class represents. This allocation scheme then provides 7.0, 22.8, and 71.1 Mbyte of tape space for compressed image data in each class. The PACS database survey shows that in only 5%, 4%, and 4.6% of cases in classes A, B, and C, respectively, did a patient's data exceed the storage space allocation with spillover to a second tape.

Intermediate Storage Period

The hierarchical storage model calls for shorter-term storage on the random access OD and migration to tape afterwards for long-term archiving. From existing data in the PACS database, a computer analysis can simulate the impact of a 1-year shorter-term archive. Figure 7 shows that the chances of returning within the first year is about 12.7%, after which the chance of return is only about 4% a year. Note that the chance of return after the first year is plotted in a cumulative fashion. This result shows that the 2-year archive currently provided by UCLA's OD jukebox is actually excessive, and that migration to tape could take place after just 1 year without degradation in retrieval efficiency. Naturally, certain caution must be exercised in making this interpretation because the PACS database is only three and a half years old. The classification and slot assignments can be readjusted as the patient population's clinical histories evolve over time. Other more sophisticated predictive methods based on patient return and examination scheduling patterns will be studied in the future.

Table 2. Classification Parameters for a Tape Jukebox

	Classification		
	A	B	C
No. of patients in 5 years	43,692	7,653	1,578
No. of tapes in one jukebox	60	34	22
No. of patients per tape	728	225	72
Space per patient (Mbyte)	7.0	22.8	71.1

DISCUSSION

Additional Strategies in Managing the Tape Archive

In the medical imaging field thus far, the tape jukebox has been overlooked as an archiving solution because of a number of concerns including reliability, speed (sequential access constraint and slow cartridge insertion), lack of vendor support, and shortcomings of standard archiving software offerings. However, the Exabyte tape offers data transfer rates and capacities comparable to ODs. Given the inherent performance characteristics of the EXB-120, these user concerns can be overcome using custom archiving software incorporating some key data management strategies. The main user

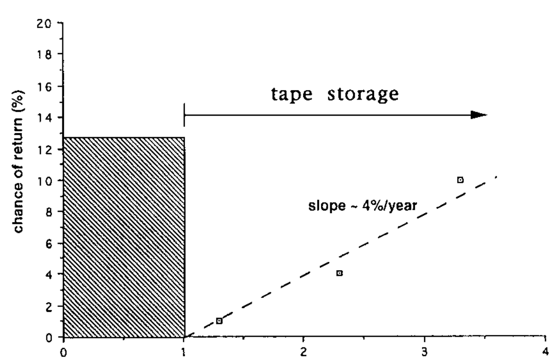

Fig 7. The five years of data show that a recently examined patient has, on average, a 12.7% chance of returning within 1 year. After 1 year, his cumulative chance of returning is only about 4% per year. In other words, there is a 4% chance of returning by year 2, an 8% chance by year 3, a 12% chance by year 4, and so on. This figure indicates that a 1-year short-term archive on OD is actually sufficient. The current 2-year storage time in the OD jukebox would not be necessary once the tape archive is implemented.

Table 3. Management Strategies for a Tape Archive

Concerns	Specification	Strategies
Reliability	1,500 uses (maximum)	Redundancy, migration to new tape, environment, tape exercising and random checking
Long seek time	120 s (maximum)	Prefetching
Long loading time	100 s	Image clustering
Sub Tbyte capacity	580 Gbyte (maximum)	> 10:1 compression

concerns and the proposed solutions for them are summarized in Table 3.

Data integrity on tape generally deteriorates because of wear and tear from extended use instead of idle shelf life. The tape manufacturer specifies more than 10 years shelf life in optimal storage conditions. To prevent data loss caused by excessive wearing of a tape, the access history of each tape in a PACS database should be kept. Based on recommendation by the manufacturer, data on tapes used more than 1500 times for either recording or reading should be transferred onto a fresh tape. Further data protection is achieved by redundant storage on two separate tapes. Some form of random checking of tape content from time to time is another means of quality assurance.

The delay caused by the long seek and loading time of the tape medium can be minimized by three strategies: (1) successful prefetching algorithms based on HIS/RIS triggers, (2) manual prefetching by film librarians who set up the reading sessions for the radiologists, and (3) data clustering by patient classification so that one patient's data are stored contiguously on a single tape.

Finally, the large capacity needed for long-term archiving can be satisfied by the tape jukebox using irreversible compression that, on average, can achieve around 10:1 in compression ratio while retaining high image fidelity.

Although legal requirements for medical records differ by state and disease category, there is always a time beyond which the image data can be purged. A more conservative approach that may result in a much longer period of storage is purging data 5 years after a patient's death. Purging can balance new records entering the system, except for any expansion in image volume serviced by the system. A stable pattern of data placement will then emerge from the cycle of data entry and purging.

Data reorganization is most conveniently performed during purging. A small number of fresh tapes are needed for buffer space. All images for each patient are transferred from one or more used tapes onto contiguous blocks of a new tape, eliminating expired data in the process. Old tapes can be either rewritten or discarded depending on its history of wear. This reorganization process can take place before data puring if fragmentation of patient record, unbalanced allocation within and between patient categories, or overuse of some tapes become evident.

Cost Analysis

The most basic problem in hierarchical storage is the trade-off between cost and speed. Cost of OD storage continues to be a disincentive for digital archiving in radiology.[8] A low-cost tape based solution, if successfully implemented in a teleradiology system or PACS, can result in significant savings. A cost analysis for the commonly available storage media illustrates the significant price difference in storage devices.

Table 4 shows the common storage media for low terabyte applications. Data for a MD are also included for comparison. Table 5 shows cost and capacity of the various types of jukeboxes.

It is clear that OD solutions rank the highest both in the cost of jukebox and medium. CD-ROMs are more reasonably priced than ODs because of their role as a mass distribution medium. However, they are less attractive than the other candidates because of the very slow

Table 4. Common Storage Media

	DAT	Exabyte Tape	Kodak OK	MOD	CD-ROM	MD
Transfer speed (Kbyte/s)	300-600	500	500	2,000	300	2,000
Capacity/disk (Mbyte)	5,000	5,000	10,000	1,300	680	1,300
Price/capacity (cent/Mbyte)	0.6	0.2	7	5.5	1.8	34.6
Access type	Sequential	Sequential	Random	Random	Random	Random

Table 5. Price and Capacity of Jukeboxes

Juke Boxes	No. Slots	Diameter/Capacity	Total Capacity	Approximate Price
Eastman-Kodak ADL	100	14 in/10 Gbyte	1 Tbyte	$400,000
MOD	908	5 1/4 in/1.3 Gbyte	1.1 Tbyte	$140,000
	144	5 1/4 in/1.3 Gbyte	186 Gbyte	$ 35,000
	32	5 1/4 in/1.3 Gbyte	41.6 Gbyte	$ 14,000
	16	5 1/4 in/1.3 Gbyte	20 Gbyte	$ 6,000
CD-ROM	300	5 1/4 in/640 Mbyte	192 Gbyte	$ 28,000
Exabyte	116	112-mm/5 Gbyte	580 Gbyte	$ 29,000
	10	160-mm/5.5 Gbyte	55 Gbyte	$ 9,000

access speed (150 Kbyte/s) and relatively small capacity (680 Mbyte). This may change as new CD-ROM drives become faster.[25] Between tapes and ODs, the tape media have a factor of 35 in price advantage. A large price gap also differentiates their jukeboxes: approximately $400,000 for the 14-inch OD jukebox[8] with 1-Tbyte capacity as opposed to $28,000 for a tape jukebox (Exabyte) with 0.5-Tbyte capacity (see Table 5).

Alternative Three-Tier Archive Design Using Magneto-Optical Disks

In systems that do not already have an OD jukebox in place, a rewritable magneto-optical disks (MOD)-based design (see Fig 8) should be considered for the important benefit of eliminating sustained OD medium cost. The rewritable capability of the MOD means that after data has been transferred to tape, a used MOD platter can be recycled for new data. Various sizes of MOD jukeboxes are listed in Table 5. The smaller capacity per jukebox offers a more modular design for new sites and is more suitable for lower-volume applications including many teleradiology practices.

In systems with MOD jukeboxes, the prefetching process involves secondary caching on the MODs. It is cost effective to retrieve more data from the tape, because the MOD is much cheaper than the MD, but far outperforms the tape in load and seek time. Using a hierarchical caching concept, MOD can cache 100% of the returning patient's prior data while moving only the most likely to be used 10% or so onto the MD. Such a provision minimizes the penalty of imprecise prefetching algorithms discussed earlier. WAN-based systems should cache more images to the MD during the prefetch than LAN-based ones because a more severe transmission delay is imposed for retrieving missing images later.

Because the MOD records are not permanent, the user must rely on the tape copy of patient images for long-term storage. Redundant copies can be kept on tape as an extra precaution with very little cost impact because of the low cost of tape. Because there is no longer a permanent record of the reversibly compressed image, the second copy should be reversibly compressed, but should be managed off line to eliminate additional jukebox cost. Although the irreversibly compressed MOD

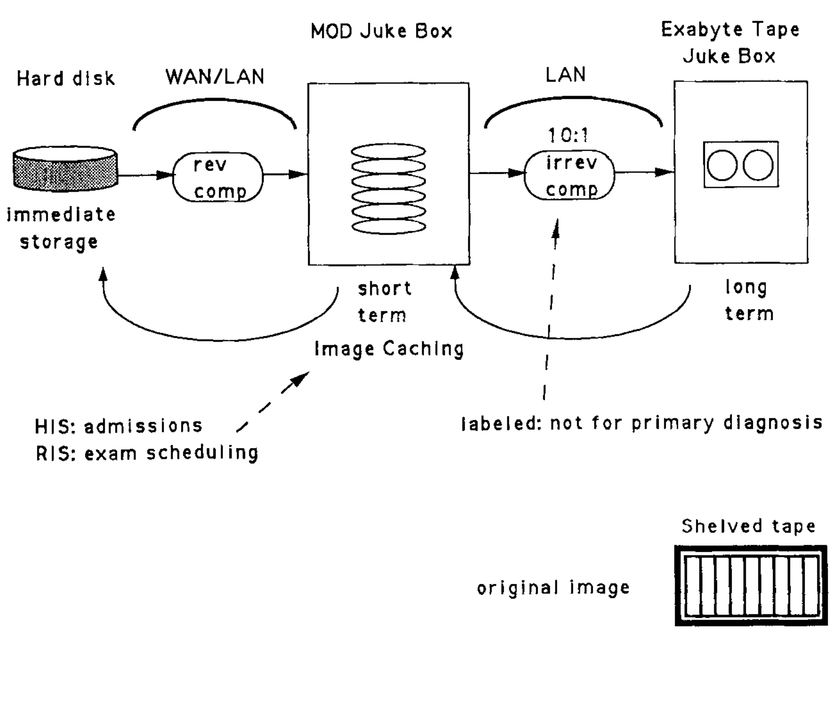

Fig 8. The most cost-effective hierarchical archive design would use erasable MODs. MOD platters are reused, thus eliminating sustained media cost. The image-caching process first transfers a large portion of patient data from tape to the MOD, then a smaller portion is transferred to the MD. Because the MODs are not used for permanent storage, an irreversibly compressed copy of the image should be stored on low-cost tapes.

copy remains available on line for review and research purposes, the off-line copy provides a redundant back up to protect against data loss as well as absolute fidelity for rare events such as medicolegal disputes or highly quality-sensitive studies.

Other Archiving Devices

The choices for storage solutions are rapidly growing with increasing popularity of 19-mm tape (E Mass, Garland, TX) and DLT, and the announcement of optical tape technology capable of storing 1 Tbyte per reel.[26] Continued downward spiral of storage and transmission costs will eventually make PACS archiving an easier problem to solve. It is important to design today's PACS and teleradiology systems based on an expandable network computing architecture so that new technologies can be easily incorporated when they become available.

CONCLUSION

The three key concepts in this paper are purposely presented in a generalized fashion to make them applicable to most digital radiology environments. They are image compression, prefetching and clustering. Because the design is intended for a multiyear cycle, many features of the architecture will take a long time to verify and optimize. The strategies presented here are being implemented for a teleradiology project involving an imaging center in Florida and subspecialty radiologists at the UCLA Department of Radiological Sciences.

The studies of algorithms and the simulations performed with a lower-cost tape jukebox suggest that the tape jukebox is a more economical alternative than an OD or a MOD jukebox as a long-term archive. The sequential access constraint of the tape can be effectively overcome with the use of intelligent prefetching algorithms backed up by manual case preparation in almost all cases. Where ad hoc retrievals are inevitable, tape access delay can be minimized by an image clustering algorithm. Furthermore, the tape medium is relatively efficient in light of the inherently slow dearchiving process and other system bottlenecks.

ACKNOWLEDGMENT

We thank Dr John Curran for editing this manuscript.

REFERENCES

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

2. Irie G: Concept of Hokkaido University picture archiving communication system. J Digit Imaging 4:3-5, 1991 (suppl 1)

3. Mosser H, Mandl A, Urban M, et al: The Vienna project SMZO in pricture archiving and communication systems (PACS). NATO ASI Series F, vol 74. Berlin, Germany, Springer-Verlag, 1991, pp 249-250, 1991

4. Glass HI, Slark NA: PACS and related research in the United Kindom in pricture archiving and communication systems (PACS). NATO ASI Series F, vol 74, Berlin, Germany, Springer-Verlag, 1991, pp 319-324

5. Georinger F: Medical diagnostic imaging support systems for military medicine in pricture archiving and communication systems (PACS). NATO ASI Series F, vol. 74. Berlin, Germany, Springer-Verlag, 1991, pp 213-230

6. Goldberg M, Rosenthal D, Chew F, et al: New high-resolution teleradiology system: Perspective study of diagnostic accuracy in 685 transmitted clinical cases. Radiology, 186:429-434, 1993

7. Weatherburn GC, Bryan S, Cocks R: Evaluation of PACS at Hammersmith Hospital: Assessment of radiology performance in the accident and emergency department. Proc SPIE, PACS Design and Evaluation 1899:452-458, 1993

8. Tighe LC: Digital archiving. Med Imaging 9:52-57, 1994

9. Miller, SW: Review of IEEE-CS MSS reference model. Digest of Papers, Tenth IEEE Symposium on Mass Storage Systems. Crisis in Mass Storage (Catalogue No. 90CH2844-9). Monterey, CA, IEEE Computer Society Press, 1990, p 3-5

10. Gerchar T: 8 m and DLT: Apples to apples? Comput Technol Rev 1994, p 34

11. Witten I, Neal RM, Cleary JG: Arithmetic coding for data compression. Commun ACM 30:520-540, 1987

12. Welch TA: A technique for high-performance data compression. Computer 17:8-19, 1984

13. Saipetch P, Ho BKT, Panwar R, et al: Application of wavelet transform with arithmetic coding in radiological image compression. IEEE Eng Med Biol 1995 (in press)

14. Wallace GK: The JPEG still picture compression standard. Commun ACM 34:31-44, 1991

15. Ho BKT, Huang HK: Specialized module for full frame radiological image compression. Opt Eng 30:544-550, 1991

16. Ho BKT, Chao J, Wu CS, et al: Full frame image compression module for radiological and industrial applications. Mach Vis Appl 3:89-96, 1991

17. Ho BKT, Chao J, Zhu P, et al: Design and implementation of full frame bit allocation image compression hardware module. Radiology 179:563-567, 1991

18. Sayre J, Ho BKT, Boechat MI, et al: Subperiosteal resorption: The effect of full-frame image compression of hand radiographs on diagnostic accuracy. J Invest Radiol 185:599-604, 1993

19. Aberle D, Gleeson F, Sayre J, et al: The effect of irreversible image compression on diagnostic accuracy in thoracic imaging. Invest Radiol 28:398-403, 1993

20. Chen DT, Ho BKT, Chao W, et al: PACS intelligence: Automatic feedback of PACS retrieval requests. Proc SPIE, PACS Design and Evaluation 2435:452-461, 1995

21. Chen D, Ho BKT, Chao W, et al: Image flow management in a PACS network. Proc SPIE, PACS Design and Evaluation 2165:233-240, 1994

22. Sheng ORL, Hu PJH, Wei CP: Acquisition and validation of image retrieval knowledge. Proc SPIE, PACS Design and Evaluation 2165:137-147, 1994

23. Huang HK, Taira RK: Infrastructure design of a picture archiving and communication system. Am J Radiol 158:743-749, 1992

24. Breant C, Taira RK, Huang HK: Interfacing aspects between the picture archiving and communication systems, radiology information systems, and hospital information systems. J Digit Imaging 6:88-94, 1993

25. Hartigan J: Desktop compact disc recording: What imaging pros need to know. Advanced Imaging 9:40-42, 1994

26. Frost MM: Honeyman JC, Staab E: Image archival technologies. PACS mini refresher course. RadioGraphics 12:339-343, 1992