PERFORMANCE EVALUATION OF A PICTURE ARCHIVING AND COMMUNICATIONS SYSTEM

BY

ALEXANDRE RIOUX

Medical Physics Unit
Faculty of Medicine
McGill University, Montreal
March 2001

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements of Master of Science

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to Dr. Christopher Henri, my thesis supervisor, for his guidance throughout this work. His patience and encouragement were continuous and his help in preparing this thesis is sincerely appreciated.

I am also grateful to the members of the PACS development and support teams at the Montreal General Hospital for technical assistance and discussion. In particular, I thank Rick Rubin and Bob Cox who introduced me to many technical aspects of the PACS. I also wish to thank Roger Sharoubim and Richard Wong who shared with me their experience of dealing with the concerns of the users of the PACS.

I would like to extend my thanks to my colleagues of the Medical Physics Unit. A special thanks to Lina Zitella, Russell Ruo, Genevieve Tremblay and Katia Reyburn for their precious help in the revision of this text.

Financial support for the research presented in this thesis was provided by the Department of Diagnostic Radiology at the Montreal General Hospital and is also greatly acknowledged.

ABSTRACT

The DICOM-conformant PACS of the MGH was subjected to a performance study. The goal was to develop the tools needed to extract meaningful, quantitative measures related to performance with a view towards establishing reliable indicators that can help identify design weaknesses and performance changes over time.

Efforts focused on operational scenarios such as moving images from servers to workstations, automatic routing and prefetching tasks. The performance of each process underlying these scenarios was studied using residency timing combined with a web-based monitoring system to yield quantitative data from various low-level resources on each server.

While the critical processes and hardware on the PACS were already monitored to identify faults, the tools developed in this study now provide the means to assess overall performance issues. The results have yielded insights that have helped to identify faulty and inefficient logic in certain program constructs, and suggestions for design improvement and software enhancement.

RÉSUMÉ

Le PACS du HGM a été soumis à des études de performance. Le but était de développer des outils qui permettent d'obtenir des mesures quantitatives reliées aux performances, afin d'établir des indicateurs fiables permettant d'identifier des points faibles du design et des variations des performances.

Les efforts ont été concentrés sur des scénarios opérationnels tels que le transfer d'images des serveurs aux stations de travail, et les services automatisés de distribution d'images. La performance de chacun des procédés utilisés pour ces scénarios a été étudié en utilisant le temps de résidence combiné avec un outils de surveillance du système qui permet ainsi l'acquisition de données quantitatives pour plusieurs ressources de base de chacun des serveurs.

Bien que les principaux programmes, ainsi que les équipements du PACS, soient déjà sous surveillance, les outils developpés dans cette étude permettent d'évaluer les problèmes de performance global. Les résultats ont aidé à identifier des erreurs et des inneficacités dans les algorithmes de certains programmes, et ont permis de suggérer des améliorations au design ainsi qu'aux programmes.

Preface

The field of Picture Archiving and Communications Systems (PACS) has evolved rapidly in recent years with the evolution of related technologies and the standardization of its operations. In 1995, the Montreal General Hospital (MGH) started to integrate digital imaging with the development of a mini-PACS serving the ultrasound department. At the time of writing this thesis, the PACS of the MGH includes virtually every modality. It receives images from two computed radiography (CR), three computed tomography (CT), one magnetic resonance imaging (MR), nine ultrasound (US), and five nuclear medicine (NM) gamma cameras, and serves users throughout the hospital and in remote locations. The same PACS, developed within the MGH, is now in use in the other hospitals of the McGill University Health Center (MUHC) as well as other hospitals and clinics world-wide.

A support team is responsible for verifying the functioning of equipment and to prevent and resolve service degradation or interruption in order to maximize the satisfaction of the users. System management programs are necessary to perform these tasks properly on a distributed system. A program, already in use, monitors the PACS and issues an appropriate warning or alert when a monitored parameter exceeds its preset tolerance. Support personnel are e-mailed or paged depending on the importance of the message, and can then start their intervention.

The aim of this thesis was to develop a method for the evaluation of image transfer performance within this PACS. This has been accomplished and, although the tools developed are in their preliminary version, the results provided have helped to guide the software developers of the MGH PACS toward further software enhancements. The aim is not only to aid in the troubleshooting, but also to provide a means of examining usage patterns and quantifying data-flow and storage requirements to help identify areas that need to be refined.

This thesis describes the transfer methods used by the PACS to deliver images from all its nodes in the various situations encountered in clinical operations. Autorouting, prefetching, and manual retrieval are the three main types of image transfer. Their design as well as the methods of measurement of their operations are detailed prior to the analysis of the data-flow.

This thesis is presented in five chapters: In chapter 1, the conventional radiology department is presented with its workflow. It is shown how the introduction of the computer and its fast evolution lead to the emergence of the electronic radiology department. In order to complete this historical background, the development of the MGH PACS is introduced and previous PACS performance studies are reviewed. In chapter 2, the transmission processes are detailed, and the various levels of communication are discussed with respect to their influence on the data-flow. Chapter 3 presents the methods used for acquiring the data on image transfer within the PACS. The monitoring tools developed for this project are described. Chapter 4 presents the data volumes measured with the daily pattern in the acquisition and retrieval of images. It also shows the higher occurrence of requests for recent images, and its impact on the distributed architecture. Residency timing of each transfer process allows us to discuss the efficiency of the algorithms used. Finally, chapter 5 provides a summary of the thesis work, including the future work suggested for the improvement of transmission algorithms and the introduction of a new PACS monitoring tool.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter one

1 Historical Background

In this chapter, a historical background of medical imaging is presented. An understanding of the historical aspects of medical imaging will allow for a better appreciation of the need for Picture Archiving and Communication Systems (PACS). For many years, radiology departments have used only radiographic film for the detection and storage of medical images. However, advancements in electronic image detection and computer technology have transformed the way medical images are captured, stored and presented. The appearance of this new technology in hospitals has made the implementation and use of PACS possible. The first generation of PACS will be introduced along with the initial problems that limited its wide-scale use in medical centers that were pioneers of this technology. The recent technological progression and improvements in standards that have made PACS viable for large clinical centers will also be described. Finally, the current state-of-the-art in PACS will be described with special attention given to the development of the PACS used at the McGill University Health Center (MUHC).

1.1 Organisation of conventional radiography departments

Roentgen took the first radiograph in 1895. While some technical improvements were made in the early years, film-based imaging is well established and has remained relatively unchanged for decades.

1.1.1 Medical imaging

Since the pioneering years of radiography, many technical improvements have been made in film radiography which continues to be the principal method of imaging in use today[1]. After a century of refinement, the acquisition and production of high quality film radiographs is well understood. While the radiologist appreciates the film medium because of its image quality, some of its drawbacks include its cost, its need for considerable storage space, physical transport, handling and filing, and the fact that films physically degrade over time.

Medical imaging is a key diagnostic tool. The way films are used by radiologists influences the way radiology departments are organized.

In the case of many diseases, when a patient shows particular symptoms or several risk factors, a particular type of image could confirm the diagnosis. Medical images may also be used for the early detection of disease not clinically evident with screening in high risk population groups such as: women over 45 years for breast cancer, smokers over 50 years for lung cancer, and patients needing anesthesia for cardiac pathologies. Patients are scheduled for an appointment and a radiographic examination or film is taken as prescribed by the physician. Once the image is acquired, its quality is reviewed by the technologist to ensure there is no need for a retake while the patient is still there. Then the image is sent to the radiologist to be read and reported. Often, similar exams or films that were taken previously are also required by the radiologist for comparison. The result (i.e., report text) is then sent to the referring physician, but the report is not always sufficient; the images may be required as well.

Medical images are also used for treatment planning in surgery, radio oncology and cardiology. The images act as a guide for localizing the disease with respect to the surrounding organs. Thus, medical images for a given patient may be used for various reasons within different departments in a hospital. A system that allows each department to access and view the images while the patient is going through treatment would clearly be beneficial.

1.1.2 Film library and delivery

An extensive team of clerical personnel provides film library services throughout the hospital. They manage the film storage and retrieval, track the location of loaned films, and ensure the delivery services.

The films acquired daily require large physical storage areas. A temporary storage room is generally located directly in the department in order to manage recent image folders, and those required for scheduled follow-up visits. Older films may be stored off site, as is done at the Montreal General Hospital (MGH). Figure 1.1 shows these areas at the MGH. The storage area must be organized to easily locate any folder. Classification numbers are used as in a conventional library. Misfiled films can become lost for years. The mean time for film retrieval was estimated to be one hour at the University of Virginia Health Sciences Center[2].

a) b) c)

Figure 1.1 Storage area from the Department of Diagnostic Radiology of the Montreal General Hospital. a) Temporary film storage room and loan desk; b) & c) long-term storage areas.

1.2 Introduction of computers in the hospital environment

The invention of the computer has revolutionized many technical fields because of the large amount of mathematical or logical operations that can be done within a short period of time. In the field of medical imaging, the two major contributions of computers are: 1) improving the management of departmental information and 2) the introduction of new types of image acquisition, specifically digital modalities. These changes were a prerequisite to the appearance of PACS.

1.2.1 Information systems

The medical information system is composed of many subsystems. Each one serves a different role in information management and is suited to the special needs of various departments[3]. The Hospital Information System (HIS) provides general information that typically includes patient demographics, insurance provider information, patient location within the hospital, and sometimes a clinical summary. The Radiological Information System (RIS) is more specific to the radiology department and handles the information related to the image acquisition and interpretation.

1.2.1.1 Hospital Information Systems (HIS)

The first HIS trial was implemented during the 1950's and the early 1960's. Using punch cards for data input, the system stored the patient file number and his basic demographics[4, 5]. The use of keyboard and CRT screens improved the interaction between the computer and user, making input and output of information easier. In the 1970's and 1980's, most HIS were based on a mainframe computer serving many terminals accessible everywhere in the hospital. By the mid 1980's, the appearance of affordable stand-alone microcomputers changed the organization of HIS. Microcomputers, minicomputers and mainframe computers were integrated together by a Local Area Network (LAN).

Although HIS have evolved over the years, the clinical function of the HIS remains unchanged. The main benefit is to provide health care professionals a computer-based medical patient record system. This patient record makes possible data entry and retrieval of patient data in a centralized and up-to-date system. Quality and efficiency of communication is then vastly improved reducing the filing and handling of documents.

1.2.1.2 Radiological Information Systems (RIS)

A state-of-the-art RIS is now made up of many modules to assist the personnel of radiology departments in each of their tasks: order entry, appointment scheduling, report dictation and transcription, film management, statistical modules and billing[6].

The film file management module handles the organization of the film library. It tracks the film position in the department at any time, manages the loan of films outside the department and helps to plan the maintenance of old films. The location and retrieval accuracy rate of current RIS was reported by Lehr and Steinberg to be 98%[6], which is a great improvement over manual management. But this remaining 2% of temporary or permanent loss of film still causes problems. The loans outside the department are an important source of lost films.

These loans are also longer in time and limit the access to the film available in a single copy for a few days.

A process control module helps to follow patients during the time they spend in the radiology department procedure. It records the progress of all examinations in the department. It facilitates the flow of work, triggering actions needed in further stages of the process. For a scheduled patient, the RIS searches for previous exams and issues retrieval requests to the film room. This practice is called "prefetching" and ensures that previous relevant studies will be in the patient's folder when the radiologist examines the new study.

The RIS is generally linked with the HIS to limit information redundancy. The scheduling module organizes the time usage of each imaging room, records the information regarding the studies requested, and reduces the delay in obtaining these requested studies. The report module allows transcription, retrieval and distribution[1] of written reports.

Although RIS for film-based departments are serving their purpose, using a digital filmless system would be more practical for convenient access and improved efficiency. Technical limitations slowed down the realization of such a project. The high cost, slow speed, and low reliability for high-resolution image manipulation were the primary reasons for the delay of a suitable digital image management system module. In order to handle these challenges, independent systems were developed to improve image archiving and communication services[6].

1.2.2 Digital imaging

The introduction of computers brought new opportunities in imaging. Early imaging technology relied solely on physical and chemical processes, but the advancement in computation power allowed for the possibility of developing new imaging technologies or modalities. Many analog systems where converted into digital systems to improve their quality and offer more image treatment possibilities. Digital imaging techniques brought new

[1] Usually the reports are not written by the radiologist, they are recorded in some way and are then transcribed by a secretary. Speech recognition software has only recently begun to be used.

methods for the display, analysis, and storage of images on workstations dedicated to these digital imaging modalities.

1.2.2.1 Inherent digital imaging system

The introduction of the first clinical Computed Tomography scanner in 1972 by Hounsfield was an important breakthrough in medical imaging[7]. Computed Tomography (CT) was the first imaging modality requiring mathematical reconstruction to produce images. The tomographic reconstruction process developed was then used in many other new imaging technologies: Single Photon Emission Computed Tomography (SPECT), Positron Emission Tomography (PET), and more recently Magnetic Resonance Imaging (MRI). All of these imaging technologies required strong computational power for rapid execution of the entire reconstruction algorithm. The reconstruction of an image was only one of the roles played by the computer. The computer was also used to interface with the user (it controlled the scanner), to acquire data, and to display and store images. CT and the other digital imaging technologies introduced the computer to radiology departments and demonstrated its usefulness. Nevertheless, for viewing and storage purposes, the images were printed on film and sent to conventional film libraries.

1.2.2.2 Conversion of analogic imaging systems to digital

Other imaging modalities using electronic circuitry benefited greatly from the integration of the computers for signal acquisition and correction. In this category, we can include Ultrasound (US), Nuclear Medicine (NM), Digital Subtraction Angiography (DSA), and Digital Fluoroscopy (DF).

With the introduction of microprocessors, real-time ultrasound imaging became feasible. It facilitated dynamic evaluation of the cardiovascular system. The digitization of conventional Anger camera images in nuclear medicine allowed for new dynamic studies such as renal analysis or heart efficiency calculation. Digital Subtraction Angiography shows the veins and arteries. Fluoroscopy is one of the oldest methods of imaging as it emerged with the discovery of x-ray. This modality has benefited from digitization as well: When digitization was introduced earlier into the imaging chain the image quality improved greatly.

1.2.2.3 Dedicated workstations

Dedicated workstations were designed to perform image manipulations for all the new digital imaging systems. Standard software offered the essential tools for digital image display. More complex tools became available in specialized workstations to perform more sophisticated image analysis.

The most basic tools generally needed for viewing of softcopy images are: exam search and sort; window-level adjustment; measurement and magnification. Advanced tools are generally added to improve the visualization process, especially with multi-planar image series. For example, these might include adding cross-reference lines to help localize positions in 3D in orthogonal images. Various non-linear windowing methods can also be helpful (e.g., exponential, logarithmic, and sigmoid), and edge enhancement filters are common.

Specialized workstations provide further capabilities. 3D-reconstruction workstations allow applications like virtual endoscopy, virtual surgical simulator, and surgery guidance to be performed. Image fusion is also possible wherein two sets of images from different modalities are superimposed to combine the functional and anatomical (or other complementary) information[8].

1.3 Introduction of PACS

A Picture Archiving and Communications System (PACS) was defined early on as a digital system for acquiring, storing, moving and displaying pictures or digital image information[9]. In the early 1980's, the first attempt to manage images on a steady basis in a clinical environment showed the difficulties of such a project. Different problems limited the development of PACS. Among these were the incompatibility between vendors, the need for good digital radiography quality, a high-speed network, and affordable mass storage. Now, many large centers have pure digital radiology departments. Region and countrywide PACS will be the next step to facilitate access to patient images from any institution.

1.3.1 First trial

The first initiative for PACS came from teaching and military hospitals. Driven by the appearance of digital imaging modalities and dedicated display workstations, PACS was meant to provide basic connectivity and compatibility between existing equipment. Another objective was cost reduction by eliminating the need for film.

The first international conference dedicated to PACS was held in 1982, in Newport Beach CA, by SPIE (the International Society for Optical Engineering). The U.S. army funded one of the earliest endeavors, a teleradiology project, in 1983. The University of Washington and Georgetown University, in collaboration with Philips Medical Systems and AT&T, were involved in this project.

In 1992, after a decade of discussion on the subject, "large clinical PACS" were almost non-existent in the world. This was one of the findings of a survey published by Bauman[10, 11] in which "large clinical PACS" was arbitrarily defined as a system with at least three imaging modalities in daily clinical operation, offering images inside and outside the radiology department.

1.3.2 Introduction of the DICOM Standard

Image communication between different acquisition and display systems in a hospital was challenging due to the variety of image formats, network technologies, and communication protocols coming from different platforms, modalities, and manufacturers. In order to exchange images in such heterogeneous environments, strong standardization was needed.

In 1982, the American College of Radiology and the National Electrical Manufacturers Association (ACR-NEMA) created a committee to work on a standard to define a way to share medical images and related information within a multi-vendor environment. The first standard was published in 1985. It defined a standard for point-to-point communication and for data format dictionaries. This first standard was rapidly revised to produce a refined version, known as ACR-NEMA 2.0, in 1988. Point-to-point communication was believed to be the best way to transfer medical images at the time.

Changes were needed to transfer medical images on a regular, enterprise-based network. The fast evolution of network technology made the use of existing hardware configurations attractive. Another revision of the standard introduced more detail in data formats and network connection establishment rules using standard communication protocols. This new version was called the Digital Image Communication in Medicine (DICOM) Standard 3.0. Because DICOM covers such a broad variety of equipment, there are only some parts that specifically apply to a given modality. In order to help the consumer and to limit marketing usage of the "DICOM compliant" expression, vendors must show a "conformance statement" for each feature or service they offer with their equipment[2].

The introduction of DICOM was an important contribution to the widespread use of PACS in hospitals. With this standard, it is now possible to share images produced by different types of systems from different vendors.

1.3.3 Evolution of technology

The operations handled by PACS (e.g., viewing, storing or sending an image) are simple operations now and have been well developed over the years. A major problem in PACS is managing the large number of images and transactions done in a fully digital clinical environment. Economic factors and productivity improvements mainly led to the wide acceptance of commercial PACS. Evidently, the evolution of the technology is critical for the widespread use of PACS.

Different subsystems of PACS have evolved in the last few years to make it more accessible. Because conventional radiography is still the most widely used imaging modality, comprising 70% of all imaging studies[1], the method by which it is digitized and made accessible for a PACS is of primary importance. The explosion of network technologies and their use is also critical in the development of PACS. The massive storage capabilities and faster access have reduced this cost and made affordable the storage of the large amounts of information acquired daily.

1.3.3.1 Digital radiography

The first method used to get a digital image from conventional radiography was the scanning of film with high-resolution scanners. In the late 1980's, Computed Radiography (CR) emerged, offering the first viable system of filmless radiography. Recently, new technologies

[2] In the case of non-compliant equipment, it is always possible to adapt it using a converter. This practice is expensive and needs to be done almost on an individual basis.

grouped under the name of Direct Radiography (DR) employ a digital x-ray detector without any intermediate processor or cassette handling.

Drum scanner digitization provided the first high-resolution digital images of diagnostic quality[12]. Because of the low speed and the expertise required to operate these systems, this method was not practical for use in the clinic for daily operations. A more user-friendly "Camera-on-a-stick" type of digitizer was introduced in the mid 1980's. The camera was placed over the film that sat on a light box and the camera was set manually before the image is grabbed. The dynamic range of such a system was limited compared to the wide optical density range achievable with regular films. Laser and charged couple device (CCD) linear array film digitizers are now the state-of-the-art in digitization. Even if these types of scanners offer sufficient image quality, high speed and a competitive cost, they cannot be used on a regular basis to digitize all images of a radiology department because of the cost of film and the additional work necessary[3].

In order to go "filmless," a method was required to produce digital images directly at the acquisition step. In 1983, Fuji introduced the first photostimulable phosphor plate with its optical reader[13]. The FCR 101 introduced in 1983 needed six square meters of floor space to lodge the reader, and only scanned 45 cassettes per hour. Fuji was joined by other major competitors including Agfa, Kodak, and Lumisys. Systems introduced recently have improved a great deal and can fit in one square meter and read more than 100 cassettes per hour[14]. CR is well accepted, in part because the workflow of radiology rooms has not changed and the film handling tasks have been reduced.

Direct Radiography goes a step further, providing the digital image directly from a detector. The detector is based on an array of Thin Film Transistors (TFT)[15]. Each transistor in the array is a pixel, collecting the charge from a conversion layer. Once the exposure is done, the array is read line by line by the electronic module attached to the detector. After a fast conversion process, the image is available to the technologist on a quality assurance workstation and is then ready to be sent to the PACS. Although it is the most efficient way to

[3] It is still relevant for a filmless department to have such a scanner to digitize film from outside the department, or old film from the archive on special request.

do electronic radiology, the price and size still have to be reduced before DR is widely accepted.

1.3.3.2 High speed network

Networks provide the link between all computers that form the PACS. Their speed is important for fast image transfer. With the increased development of network applications in the 1980's, most hospitals got equipped with at least a minimal Local Area Network (LAN). The PACS could be added to the LAN already in place in the hospital or on a dedicated LAN to limit the bandwidth sharing with other tasks.

The most common technology standard used to build a LAN is Ethernet. This standard specifies hardware performances and protocols to be used for communication at a physical level. The speeds of different versions of the Ethernet standard through the last twenty years are presented in Table 1.1, which clearly shows the fast evolution of network bandwidth.

Table 1.1 Introduction of different Ethernet standards and their speed.

Year	Standard	Speed	Comments
1980	10Base5	10 Mbps	Thick wire
1995	802.3u	100 Mbps	Twisted pair or fiber optic
1997	802.3x	100 Mbps + 100 Mbps	Full duplex => could send and receive at maximum rate
1998	802.3z	1000 Mbps	Optic fiber

Over the last twenty years, Ethernet gained two orders of magnitude in speed. Usually technological improvements follow the introduction of the standard and prices are reduced with time. One of the advantages of Ethernet is the compatibility between all these different bandwidth technologies that allows networks to be configured as a function of usage and distribution of large bandwidth among users.

1.3.3.3 Mass storage

Data storage in PACS is a compromise between three important factors: price, retrieval performance and capacity. In a large PACS, the vast amounts of data that need to be stored can easily reach several terabytes[4] annually. In order to take advantage of the cost reductions in fast storage, it is wise to work with scalable systems that can grow with the need for archiving capacity. The savings from the archival of digital images have become more and more important as shown in Figure 1.2. The high costs of digital storage in the last decade slowed down the transition to digital radiology.

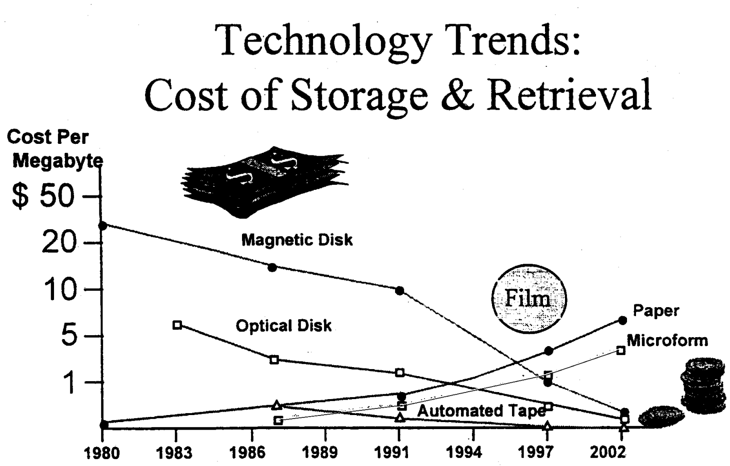

Figure 1.2 Cost evolution of storage and retrieval (figure taken from Chunn[16]).

Figure 1.2 presents three different digital storage methods: the magnetic disk, optical disk and automated tape. These three types are all automated and do not need human intervention to load or unload media. Manual systems (lacking robotics) could also be considered, but they are not suitable for anything larger than mini-PACS. Most PACS use a combination of two or three storage technologies, seeking to balance costs, automation and retrieval times.

The magnetic disk is the fastest way to store large amounts of data. It is typically used to store the most recently acquired images and is referred to as online storage. This technology is constantly improving its capacity as well as the throughput of data it can transfer. Usually, the majority of retrievals performed are for series not older than six weeks as shown previously at the MGH[17]. For this reason, compromises can be made on performance to lower the overall price for storage of older images. In this case, a tape jukebox, or any other robotic system that can automatically handle many smaller storage meds, is the solution of choice.

[4] A terabyte is equal to a million MB (10^{12} bytes)

Density of information on the medium, capacity and scalability of the physical library, the number of readers and their speed, and the efficiency of the robotic system are all factors that determine performance, capacity and price of the system.

1.3.4 Hospital-wide PACS

The first PACS were small systems that offered electronic imaging within small clusters of modalities and workstations in radiology departments. Availability of DICOM compliant imaging equipment and the evolution of technology involved with PACS made it easier to extend services hospital-wide. Supporting both digital imaging and film in the same department is expensive. The real cost savings of PACS is realized by going filmless.

A hospital may implement PACS in one of two ways. The first is to opt for a turnkey PACS project provided by a single vendor. The other approach is to build the PACS progressively by adding DICOM compatible imaging equipment as it replaces older equipment. On the user side, the workstation pool could also be built progressively, providing a station to the most important group of users and installing viewing rooms where needed. Further investments can be made in teleradiology, adding stations at the homes of radiologists or in remote clinics.

The presence of large PACS world-wide was evaluated with a survey in 1995[10, 11]. Twenty-three installations were found to correspond to a definition of a "large clinical PACS[5]." This represents twice the number of large PACS found in 1993. Now, several vendors claim to have installed PACS of various proportions in hundreds of sites. The number of systems in operation is increasing at a fast rate. However, the number of hospital-wide PACS is still low. One factor slowing the process of deploying PACS is that most radiology departments possess non-DICOM compliant equipment, making it difficult to integrate. Consequently, many may choose to use this equipment for its useful lifetime before upgrading it with DICOM compatible systems.

[5] At least 3 modalities must be in daily clinical use with service in and out of the radiology department.

1.3.5 Countrywide PACS

Because most hospitals interact with different clinics, and other hospitals in the same region, it becomes necessary for those with PACS (particularly when operating without film) to provide an electronic means of accessing images.

The American Army has initiated a project to connect all military hospitals to a countrywide PACS into a Virtual Radiology Department (VRD)[18]. All radiologists are brought together in a countrywide radiology department spread over many hospitals and remote clinics. Emergency or specialized cases to be examined are distributed among radiologists logged on the system. The objective is to overcome geographical limitations and to share the radiologist resources all over the territory. From the moment that a network with sufficient network bandwidth ensures fast enough connectivity between centers, it becomes possible to consider PACS on such an immense scale. UCLA also started to offer nation-wide and international teleradiology services[19], relying on a Wide Area Network (WAN) to offer service in Latin America, Asia Pacific and the U.S.

1.4 PACS at the Montreal General Hospital

The Montreal General Hospital (MGH) began developing its own PACS in 1995. Henri, Cox and Rubin have extensively described the implementation and the development of the system done in-house[17, 20-28]. The first experience with digital image management began with the development of an ultrasound mini-PACS in 1995. Based on the success of this system, the project was extended to Magnetic Resonance and Computed Tomography in 1997. Nuclear medicine was added one year later. The latest addition was computed radiography, which was first deployed in the Emergency department in the fall of 1999, then throughout the Radiology department in 2000.

1.4.1 Ultrasound mini-PACS

The first PACS project at the MGH was launched in 1995 within the ultrasound (US) department[23]. Nine ultrasound scanners were attached to computers with frame grabbing software. The information related to the patient and the study was entered by technicians at the start of the exam. Images were captured via a simple foot-pedal during the exam and, upon completion, the images were transmitted over the existing departmental Ethernet network to a central server. This server was equipped with a 9 GB hard drive to store the most recent exams on-line. For long-term storage, recordable compact disks (CD-R) were used[21]. A disk was burned with all the images produced each day. For patients with previous studies, new images were added to their history file to simplify the retrieval of their complete history. A web interface was later developed to access the images on the image server[24]. For retrieval from the archive, the CDs were stored in the film library, where the clerks were prompted by a computer to load the desired CDs. The start-up costs of this mini-PACS were recovered within six months[23]. The savings came from the elimination of film, a film processor and the reallocation one full-time clerk who was responsible for film management within the ultrasound department.

1.4.2 CT and MR integration

Based on the success of the US mini-PACS and the expertise gained in developing it, in 1996 a plan was devised to expand the PACS to include CT and MR. The decision was taken to support the DICOM 3.0 standard[21, 26].

In order to minimize the software development time, the DICOM Central Test Node (CTN) software developed by the Mallinckrodt Institute was used[29]. Freely available, this software offers the widest implementation of DICOM services described in the Standard. At McGill, it was adapted to work in a distributed system employing multiple servers. The Linux operating system was selected for these servers owing to its robustness, flexibility and cost. For the display of the images, eight diagnostic workstations were bought from a commercial vendor. These were the only components that were not developed in-house.

The funding for these investments came from a cash advance equivalent to the film budgets for two years in CT and MRI, plus some rebates from the film vendor. From the moment the PACS was extended to the CT and MR, no film was available for these modalities. This practice forced the extensive use of the PACS in all possible situations and highlighted weaknesses in the system. These problems included limited access to images outside the radiology department, an increasing demand for old images stored on the CD-based archive (which was a manual system), a bottleneck at the Web Server, and difficulty in documenting interesting cases for teaching files. The development team, working closely with the clinical users of the PACS, solved each of these problems.

Access to images outside the radiology department was improved by adding workstations and paper printers in critical locations throughout the hospital. One weakness of the web server was that it could only sustain the last week of images due to the limited disk space it had available. The solution was to keep the web server as the main point of entrance for the web service, and transfer image requests to the server(s) where the images resided. This required each server to become a fully capable web server. Over time, the increased volumes of data and users of the PACS placed heavy pressure on the manual CD-based archive system. The size of the online capacity was increased as a temporary solution. The problem was finally solved by the purchase of a Digital Linear Tape jukebox containing 48 tapes, each with a capacity of 35,000 MB (uncompressed). This system provided 24-hour response to retrieval requests.

1.4.3 Extension to conventional radiography

In 1999, two Computed Radiography (CR) readers were bought and integrated with the PACS. The Emergency department was first equipped with CR. Usage of CR technology was introduced gradually to other areas of the hospital, including Intensive Care Units and the main Radiology department. The number of images, their sizes, and the number of people needing them were so large that a relatively long transition period was employed to accommodate the changing of the workflows in the entire hospital. After the installation, images were accessible either in hardcopy (on film) or in electronic format through the PACS. Hardcopies were only available for a few months to allow the accumulation of enough prior exams on the PACS to minimize the need to work with film entirely.

1.4.4 Current state

Within the MGH, there are 55 diagnostic and 50 clinical workstations throughout the hospital capable of accessing any study on the PACS. Eighteen radiologists have a diagnostic workstation installed at home. The other hospitals of the MUHC are starting to implement the PACS, and have a total of 59 additional workstations. In addition to the users of these review workstations, at least 1,300 users have accessed the PACS via the web and opened an account.

A distributed architecture comprising 15 servers (as shown in Figure 1.3) forms the PACS to provide all the services needed to these users. A majority of "Modality Servers" receive and manage the images from the imaging modalities, while the others are more specialized, like the master database server and its standby mirror, the web server, and the teaching file server. Total storage capacity available online on the RAIDs attached to the Modality Servers is 430,000 MB. The Digital Linear Tape (DLT) library currently providing the long-term archiving contains 1,680,000 MB. However, a new tape library was recently installed with a starting capacity of 50 TB that could be increased by adding new modules and tapes. Hundreds of CDs remain from the first archiving system, but they still need to be manually loaded for consultation. All the information about the images archived since the installation of the PACS is centralized in the Master database.

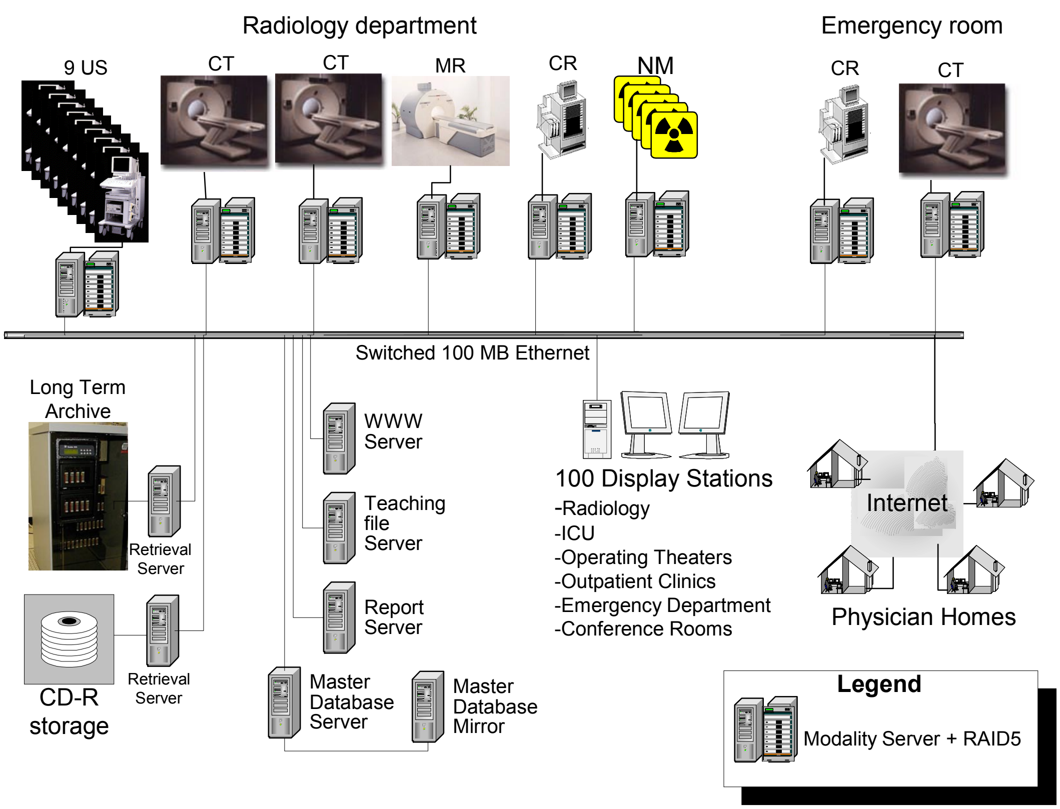

Figure 1.3 Logical schematic of the MGH PACS.

1.5 Performance evaluation

Research on PACS reflects its evolution. Simulation and modeling of different PACS configurations were studied at the design and planning stages. Once real systems were implemented, a new approach of examining PACS relied on studying different subsystems for fine-tuning and obtaining a better understanding of their performance. A more recent approach, which came with the appearance of larger systems, is the monitoring method where tools are developed to analyze the overall PACS performance.

1.5.1 Simulation and modeling

Simulation and modeling is an approach for efficient design in the planning process of networks[30-33]. Traffic engineering studies the efficiency, reliability, and capacity of different network organizations and protocols by using queuing techniques, stochastic and other mathematical methods.

Many hospitals and companies in a pre-operational phase use simulation and modeling to determine an optimal PACS configuration[34-43]. However, simulation rarely provides reliable performance indicators on more than two or three aspects of the PACS. In the early years of PACS development, simulations were used for design choices. Distributed architectures were selected over centralized systems for their reduced transfer delays when large volumes of transactions occur[39, 42] and for an expected increase in reliability and data availability. Toshiba validated a new fiber optic network with a duplicated star topology in 1990[40]. The stochastic simulation was based on a detailed model of the PACS. It assumes a typical workload and provides expected delays for image delivery as a function of the number of nodes connected. GE Corporate Research investigated leasing of T1 lines for teleradiology in digital mammography with a remote screening site[43]. T1 links provide enough bandwidth as it was evaluated on their testbed to transfer one exam per 10.5 minutes. Therefore, T1 links were judged sufficient for the mass screening program where films are read in batches at the end of the day, but additional bandwidth would be required to do remote reading in real-time.

Of the simulation and mathematical methods published, many required simplifying assumptions. Real PACS became more complex because of the variety of equipment types, vendors and usage patterns. Most pre-operational simulations do not provide any validation with data obtained on a real PACS under a clinical load. Validation with real data was performed by Masuzawa, Miyauchi et al. using the most simplified PACS segment under ideal conditions and showed discrepancies when compared to simulation results of up to 50%[36]. The simulation of large scale PACS, considering the complexity of the communication network and the image traffic, does not seem achievable as observed by Chen[44].

1.5.2 Performance study

Many groups have done performance testing on PACS installed within their centers. Performance studies have been performed on segments of PACS or partial layers of them[16, 45-55]. Storage, network and workstations are often evaluated, but in these publications particular attention has been given to large traffic volumes, the DICOM requirements, and the clinical context encountered in an electronic radiology department.

1.5.2.1 Storage

Storage is an important aspect of PACS. The large amount of data stored over time and the need for fast image access of a recent study require a hierarchical storage strategy. There is a wide range of storage technology performances and capacities. The Storage Technology Corporation[16] has compared optical disks and magnetic tapes. Performances were shown by plotting the delay for retrieval as a function of the number of media exchanges per hour. These curves clearly showed an exponential increase of the delay as the number of requests grew. AT&T Bell Laboratories has discussed different implementation strategies of efficient hierarchical storage management[56].

Redundant Arrays of Inexpensive Disks (RAIDs) are usually chosen for the fastest image retrieval times. Indeed, the presence of many disks working in parallel provides better throughput and redundancy. Many different RAID algorithms are available and parameters can be adjusted in order to optimize performance[46, 47]. Recommendations regarding concurrent or sequential retrieval are inconclusive.

1.5.2.2 Network

In a distributed architecture where several servers service workstations throughout a hospital and many other remote users, the network has an important role to play. A slow network would limit the interactivity within the PACS, slowing down all processes requiring image transfer. The physical capacity associated with each network technology is clearly defined. Prior, Niggenaber et al.[48] reported that at the application level, these theoretical limits are not attained, and an order of magnitude of loss between the signaling rate at the physical layer and the data throughput at the application layer occurs for Ethernet network using TCP/IP protocol. The loss results from additional information at each layer and delays in processing the packets from layer to layer.

Since the adoption of the DICOM standard, an additional layer in the communication stack needs to be considered. A research group from the Mallinckrodt Institute of Radiology, where the Central Testing Node (CTN) software was developed, reported on the delay for DICOM connection establishment and release using their software[49]. With TCP/IP, the association requires 0.269 seconds and the release 0.0092 seconds. These times add to the delays encountered with the other layers. This study also showed a throughput of up to 7 megabits per second (Mbps) for large CR image transfer from memory-to-memory. This throughput rate is a lot higher than reported in previous studies. It is also important to note that there was a 20% decrease of the throughput for end-to-end over memory-to-memory transfer. End-to-end transfer rate is the speed at which data is transferred from remote storage to local storage and is the parameter directly experienced by the user.

The effect of establishing a valid DICOM connection on large bandwidth networks was studied by Gohel and Whitman[50]. In order to transfer data in accordance with the DICOM standard, additional headers and processes were needed and lead to performance degradation compared with the optimal TCP transfer rate. This degradation varied as a function of the series types. Large series of small images encountered a reduction of transfer rate because each image had its own header and transfer units were not completely filled. In this particular study, the connection took a longer time to be established (0.94 sec.), but similar Ethernet throughput was achieved (7 Mbps).

1.5.2.3 Workstations

The network, storage, and server are the core components of the PACS, but the user only interacts with the PACS at the workstation. From the user's perspective, the performance of the PACS is measured by how efficiently tasks are executed on the workstation. The choice of PACS equipment has a strong influence on this performance. Better workstations improved the overall performance of the PACS[51]. Improvements by a factor of two were reported when comparing an Intel 80486 33 MHz and a Pentium 90 MHz for regular image processing time, image display time from the server, and long-term storage. It is important to note that 10:1 jpeg compression was used for this study. The improvements in display time were mainly related to the faster decompression processing abilities of the Pentium 90. Other improvements more important than decompression included changing the operating system from windows 3.1 (16 bit) to Windows 95 (32 bit). With so many variables involved, it is clear that special care must be taken when making comparisons between different workstation configurations.

In PACS, the workstation has replaced the traditional light box. Comparisons of the workflow and reporting time for softcopy viewing[6] and conventional film reading have been published[52]. Every step needed for the reading process was analyzed using Little's Law to examine the bottleneck and throughput rate for conventional and softcopy reading. In film reading, the time spent by the film room personnel pulling old examinations was found to be the bottleneck, while the second most time-consuming step was the patient examination. With softcopy viewing, the slowest step of the process was the interpretation followed by the retrieval of an old study. However, it was also noted that delays at the retrieval level could have been reduced by adopting good prefetching strategies[53]. One point that was controversial was the time savings on image reading itself. The current study reported improvements in softcopy reading, but Beard and Foley[54, 55] reported faster reading times when using conventional film.

1.5.3 System performance

PACS requires the connection of various pieces of equipment to a common network. The performance of each component provides useful information, but there is a need for developing methods that can rate the performance of the overall system. These methods of evaluating system performance help to achieve a better understanding of the PACS and allow for system improvements. Data can be acquired at one point in time or during installation for

[6] Here softcopy viewing refers to the examination of the digital image on a computer monitor.

quality acceptance testing. Such data could be used to monitor the PACS as a regular quality control tool as the system evolves and the needs of the users change. Monitoring of the results can help the manager to troubleshoot on a daily basis and to plan system improvements for the long-term.

Some early studies have looked at overall system performance[45, 57]. General system loads were given to evaluate the storage requirements and network usage. Throughputs were evaluated using the production of new images every day, the mean number of images transferred, and the number of user queries[57].

1.5.3.1 Quality assurance

Acceptance testing is critical in the purchase of high priced technology since it represents the only way to ensure that the buyer gets what was paid for. Usually, most acceptance tests in the PACS field are focused on compatibility and connectivity problems, such as those detailed by the University of Florida Radiology group[58]. However, once each unit of the system can communicate properly with each other, it is important to verify how fast the images are delivered and if they satisfy the stated requirements.

It is possible to employ simple qualitative tests mainly based on general impressions from the user[59]. However, a more systematic approach such as the one defined by the Department of Defense[60-62] is needed for a true system performance evaluation. The benchmark test procedures and equipment tests that are used were developed through a joint effort between the Government, academic institutions and private institutions. The testing is split into three phases[61]. The first part consists of a Component and Subsystem Demonstration, where each component is tested alone with appropriate input. Once all the components are tested, the system integration tests take place. In this phase, small-segmented equivalents to what is encountered in a clinical system are set up and tested as a chain. Finally, a Clinical Loading Test is performed in the presence of a background system and realistic network traffic. In total, 79 tests are defined and provide a complete benchmarking procedure.

1.5.3.2 Monitoring

Monitoring of a PACS is a high-level feature in the system. It was first studied by the pioneers of PACS at UCLA and in US military hospitals. A GE group also studied the traffic protocol with respect to the modality load on the UCLA PACS[63]. Then, the workload distribution of the Madigan Army Medical Center was analyzed at the medical entity level.

Doris T. Chen proposed many monitoring strategies for the UCLA PACS in her thesis and two related publications[19, 64]. Software was developed to monitor all the workstations and servers of the system. Traffic levels, queue lengths, and system alerts were provided. From the monitoring center, the administrator was able to troubleshoot most of the problems, even for remote centers. A feedback for the time remaining before the completion of a retrieval request was provided by using the number of images in queue before the requested images. The program developed was adapted from Netview 6000 using the SNMP protocol and object-oriented models. This kind of monitoring on real clinical PACS was more reliable for design optimization than any simulation. One of the most interesting findings was that the network traffic level did not increase the transfer delay of images for traffic levels up to 60% (0.75 MB/sec or 6 Mbps) on 10 Mbps Ethernet Network. The main factor was still the image size, but the position in the network topology also influences this delay when several bridges and routers need to relay the data to its destination.

Chapter 2

2 Image transfer in a PACS

In this chapter, the technical aspects of image transfer in a PACS are explored in detail. This material is necessary for a good understanding of the experimental chapters which describe the performance and optimization studies that were performed in this thesis.

When an image is transmitted from one device to another, it is processed by a series of applications. Each application forms a layer in the "communication stack." It is important to understand each level of the image transfer because of its influence on the efficiency and speed of the transfer. The first section of this chapter will include descriptions of image storage, the application level, the establishment of a DICOM association, and the network level of each level of the image transfer process because of its influence on the efficiency and speed of transmission.

Traditionally, film served many purposes in the organization of the conventional radiology department. The "filmless" environment now made possible by PACS has replaced traditional delivery mechanisms with electronic ones that usually include features known as autorouting, prefetching, and (manual) retrieval. The second section of this chapter will describe the design used at the MGH for these three electronic delivery mechanisms.

2.1 Communication stack

The process of transmitting images from one computer to another proceeds as follows. The images are loaded by an application from the memory support (hard disk), then a connection is established with the destination device according to DICOM rules, and the images are transmitted through the network using communication protocols. This chain of operations is what forms the communication stack illustrated in Figure 2.1. On the receiver side, images go through the same stack in reverse order. Each layer in the communication stack must be compatible with the ones above and below it in order to perform its particular task. The compatibility is also compulsory between each layer of its peer on the receiver side.

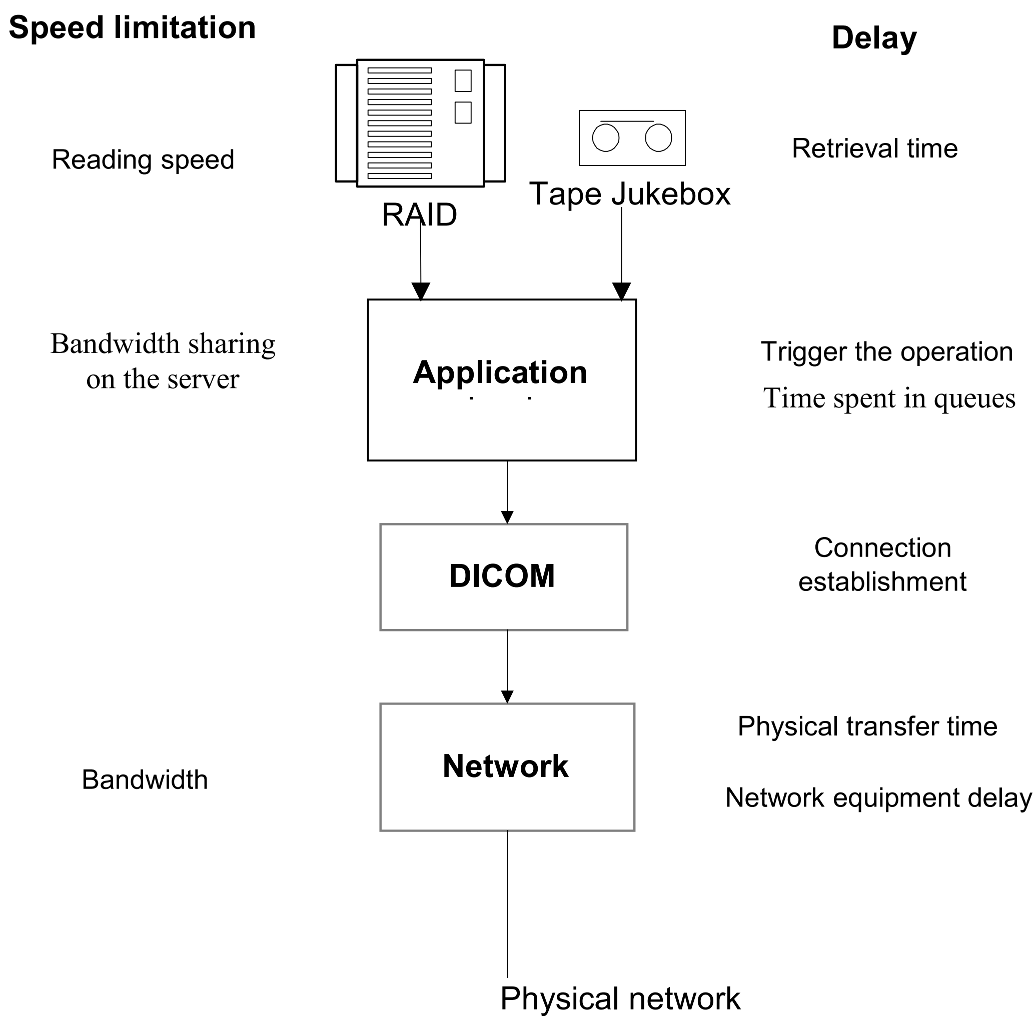

Figure 2.1 Communication stack and the expected sources of speed limitations and delays.

The performance is affected by the efficiency of each layer and the delay it induces since the images need to pass through each of these steps to be transferred. For a better understanding of the performance aspect, factors affecting the time of transfer for each layer will be presented in the following sections.

2.1.1 Storage

The storage is the source and the destination of all images transferred in a PACS. The speed at which it can read or write the information is critical in the communication stack. The transfer cannot be done faster than the speed of these two operations. For this reason, it is important to understand the factors influencing transfer speed from or to different types of storage. With proper design strategies, certain data may be stored in a more efficient and easily retrievable space while the remainder are stored in large, inexpensive long-term archives. There are a wide variety of storage technologies. They consist of conventional memory, hard drives, RAID, removable storage, and tape libraries. In order to optimize speed, the information stored on the PACS should always be retrieved from where it is the most accessible. On a PC, the cache is always verified for CPU instructions. When the information is not there, it is taken from the system memory. If the information is not in the dynamic memory, it is then read from the hard drive or from another slower storage media. In a PACS, the same logic applies to the workstation, but this is extended to the storage of the whole system. Image management provides an easy access to the most requested images, but also maintains reasonable access to all the other images produced in the department.

As mentioned above, prior studies are often desired by the radiologist for comparison with the current images. In order to respect provincial requirements[7], the department has to determine a retention schedule and keep a copy of old images for a minimum of five years in order to document the file of the patient. Mammograms are kept over ten years for the sake of research.

[7] According to the Archives Act.

2.1.1.1 Memory

The memory, also called the Random Access Memory (RAM), is the working memory used directly by the CPU to run a program and to store data temporarily while in use. This kind of storage should be quickly accessible. The response time could reach a few nanoseconds. This speed is achieved by maintaining the data at an electronic level. Since there is no mechanical reading of the data, the response time matches the speed of the other electronic components in the computer. However, one drawback is that data are lost when the power is turned off.

The speed of memory access is determined by the output bus width, the memory clock speed, and the protocol used to deliver the data on the bus. The output bus width determines the number of bits that can be read at the same time on the parallel link. The number of times the information on the bus can be updated determines the bandwidth of the memory and the amount of information transferable per second. Memory access time is on the order of a few nanoseconds (ranging from 6 to 50 nanoseconds) and cannot be used to find a steady bandwidth. Different memory technologies use different access timing methods. Most modern methods are bursted. For example, the access to the first set of data takes five clock counts and the three following data sets may only take two clock counts. For this reason, the memory access time cannot be used for absolute bandwidth measurement, but could be used as the maximum data rate achievable.

2.1.1.2 Hard drive

Over the years, the hard drive has become essential to permanently store large amount of data. The data is physically stored by reading or inducing magnetic moments on a magnetic disk. A head made from a coil can read or write the disc as it is spinning. The head is positioned on the radial axis of the disk to select the chosen region of the disk. Many disks are stacked one on top of each other to obtain larger capacities.

The density of information on the disk and the rotational speed mainly determine the physical performance of the disk. The faster-rotating disk with higher data density will have a higher output rate. In fact, disk rotation speeds of 10,000 rpm and densities of 30,000 MB/In^2 are now common.

In order to get the full picture, the hard drive must be seen as a physical reading module controlled by an interface module. The interface module transforms the raw physical signal into an intelligible signal for the parallel link to the computer. A fast cache memory acts as a buffer to ensure smooth operation between these two modules.

The performance depends on the time for the head to reach the track and on the rotation speed of the sector. This is called the seek time. The seek time on the hard drive is much longer than with memory; it takes around 10 milliseconds for a hard drive compared with memory which takes but a few nanoseconds. The longer seeking time is mainly due to the time required for the head to reach the right track and for the disk to rotate until it meets the right sector. The theoretical physical data rate is on the order of 80 Megabits per second and is found using the data density and the spinning speed of the disk. The controller also has a theoretical maximum data rate: it could reach up to 320 Megabits per second for ultra-wide SCSI controllers. But these theoretical speeds can only be obtained in ideal conditions, because there is a loss of speed due to the coordination between the physical reading and the controller.

2.1.1.3 RAID

RAID stands for Redundant Array of Inexpensive Disks. It is a way of getting larger hard drive capacity while minimizing the cost and reliability problems by using redundancy of information. RAID is obtained by using many conventional hard disks with a fast controller card managing them all with a particular algorithm. The algorithm is responsible for the level of redundancy and the efficiency of reading and writing.

There are many ways to mount a hard drive and write data to it in a RAID system. The most popular ways are called RAID levels 0 to 5. Of most interest are levels 0, 1 and 5. Level 0 is optimized for transfer rate performance, level 1 is optimized for reliability, and level 5 is a compromise between both extremes. Level 0 has no redundancy at all. The information is spread over all the disks using virtual stripes. The reading and writing performance is optimal because all the disks are used at the same time and their output rates are combined. On the opposite side, level 1 has full redundancy using a mirror of the whole data set. The reading speed is quicker because everything is duplicated and can be read from two drives but writing speed is not positively affected. The capacity is also limited, because 50% of the space is lost for mirroring. Level 5 is an efficient balance of performance and reliability. All the disks are used simultaneously in level 5 while maintaining virtual stripes as in level 0. In case of disk failure, parity information is stored on each disk to allow reconstruction of the lost data. The information needed for the reconstruction takes less space than a simple copy. For example, four disks mounted in RAID 5, required 25% of the space for the redundancy. The reading and writing performance benefit from the advantage of having multiple disks working together.

The large capacity achievable with RAID architectures combined with the performance and reliability improvements is very useful for server storage particularly in a PACS, because a large number of files are read and written concurrently.

2.1.1.4 Optical disks and magnetic tape

Many removable storage mediums are available. Their speed is limited compared with the technology we have seen so far, but the low cost of these mediums justifies their use. These mediums may also be an additional alternative for moving the data when a network link is not available.

The optical disk family consists of the CD-R, the CD-RW, magneto-optical disk (MOD) and recently the DVD-RW. The R stands for recordable, and the RW for rewritable. The disks are read by a laser beam, which is focused on the surface. Holes in the reflective coating indicate a 1 or 0 value. An optical detector receives the reflected light signal and the data stored are then reconstructed from this signal. Magnetic tapes are also available in different formats, such as quarter-inch cartridges (QIC), 4mm digital audiotape (DAT), 8mm helical scan, digital linear tape (DLT), and advanced intelligent tape (AIT).

The main distinction between these different standards is the information density and the read/write speed. Optical disks also have the advantage of having a longer life span compared with magnetic storage, which could start losing information after five years. Removable storage is relatively cheap, but requires human handling which could become a problem for a distributed system or systems running 24 hours per day.

2.1.1.5 Automatic libraries and jukeboxes

Libraries and jukeboxes are mechanical pieces of equipment that handle many tapes or disks without human intervention. A robotic arm or mechanical device brings the tape/disk to a reader (some systems accommodate more than one reader). A larger library should contain two readers or more. Multiple readers allow simultaneous reading on different tapes, and a higher overall throughput. The capacity is determined in terms of the number of tapes or disks that can fit in the library. Some systems are scalable and have practically no limit. These libraries have a huge capacity to store data. However, a disadvantage is the slower access speed. This limits their usefulness for storing frequently used information. The combination of the loading time to get the appropriate tape and the data access speed usually requires hundreds of seconds.

2.1.2 Application level

The application level takes the image from the storage and provides the interface with the users. Most users of a PACS are medical staff who examine the images and related information on the workstations.

2.1.2.1 Image viewer

The image viewer in a well-organized PACS is the only interface the ordinary user should have to work with. It provides basic image viewing and analysis tools and also links the workstation with the PACS to access all available digital images at the hospital. In a well-built PACS this link to the system would be almost transparent: This quality is called seamless integration. The user does not need to be aware of the process inside the distributed architecture of the system. He simply needs to view the images with a maximum of interactivity. Web interfaces and software native to the workstation typically include the features detailed in Table 2.1.

Web viewers have accessibility advantages over native software because they provide worldwide access and minimal setup and configuration on the client side. This advantage is limited by the transfer efficiency of the Internet compared to private networks. In order to overcome transfer inefficiencies or transfer delays, the images are often lossy compressed, which affects the image quality[8].

Table 2.1 Description of usual image viewer features.

Feature	Description
Searching	Search images on local disk, Modality Server or master database using patient's name, date of examination, ID number, or other DICOM information.
Viewing	View images in a highly configurable manner with many screen layouts. Viewing tools include: window level, zoom, probe, pan and linear measurement.
Transmission	Send an image to any Application Entity visible on the PACS.
Cross-referencing	Show reference lines to correlate the positions of image slices with scout views.
Annotating	Annotations made on the image.
Comparing	Compare two similar studies side-by-side on the same screen.
Scrapbook creation	Select important images of different studies to summarize a case or for educational purposes.

In order to study the efficiency of the sending process, one must observe the time needed by the application level to trigger the sending process, especially when there is a lot of database interaction. It may take up to a few seconds before transmission begins if polling is used.

2.1.2.2 PACS image management software

The PACS image management software runs on the servers to provide storage and communication services to the modalities and the review workstations. It manages the disk space for online access and the database for local queries and retrievals. It also provides automatic routing and prefetching to limit the number of requests needed from the users.

The efficiency of the software is difficult to analyze. The software juggles communications from several workstations and handles a large amount of images. Normally, the images spend a few seconds in queues before they are processed. As the software becomes overloaded or the server experiences equipment trouble, the queue length grows rapidly, retaining in memory the images which still require transfer.

[8] Use of lossy compression for medical images is still controversial in the PACS community. The use of lossy compression is usually limited for non-diagnostic image review.

Servers have to share their access to the network among all their tasks. When many stations request images and some are already receiving them from the same server, the speed of the transfer is reduced since the available bandwidth of the server's network card is split among many sending and receiving processes. Well-designed PACS management software will use efficient strategies to ensure intelligent choices in this distribution of the bandwidth.

2.1.3 DICOM

DICOM is a standard which specifies the way medical image exchange is performed. Using this standard allows all kinds of operating systems, imaging modalities, and system peripherals to be used together with maximum intercompatibility.

DICOM employs an object-oriented model. Objects used in the field are defined with selected attributes forming the data structure. In DICOM, the computer or the communicating host is called the Application Entity. Then, services between pairs of application entities define the action accomplished together. An application entity only performs services on selected objects. For example, a CT scanner moves, stores and prints CT images, and is not usually configured to provide the same service for other types of images. However, a printer should be able to print all the images available in a department.

From the connection point of view, DICOM regulates the data flow between the application level and the network or between the application level and the storage medium. Before any service is provided between two application entities, a connection is established and remains until the end of the transaction. Then, the transaction between the two application entities occurs using the same data format.

2.1.3.1 Establishing a DICOM Association

The purpose of the association is to negotiate the context of communications. This information forms the application context for the connection. The Service Object Pair (SOP), required for a specified type of object, is verified because it must be available on both application entities. The roles of the client and the provider are defined. The way the information is represented is chosen. The transfer syntax could be little or big endian. When using little endian, the application entity will send the most significant bits first. On the other hand, when using big endian, less significant bits are sent first. The use of compression algorithms is also specified.

Establishing an association takes some time. In order to minimize this delay, the connection is established only once for several images of the same series. Then, when the series is sent, the connection is released. The time to establish the connection depends on the latency of the network and the time needed for the roundtrip travel between the pair of communicating computers. It is important to note that as the volume of data to be sent increases, the contribution of the connection establishment and release times to the overall transmission time decreases.

2.1.3.2 Data structure

For the system to interoperate, a common data structure is required. There are three ways to analyze the data structure: the relational structure, the image header, and the data flow.

The relational structure is the logical manner in which the data are organized. It reflects the hierarchy between different types of data. The relational structure avoids redundancy in data storage and simplifies queries. In Figure 2.2, the main relational structure is shown graphically.

It may seem obvious that the patient is the object that has studies, which contains series of many images. Until it is strictly defined, any other organization could be used to classify the data of each level. In the diagram, the boxes represent the objects that contain the related information and the diamonds contain the relation between the different objects. For each link there is a field that allows one to find the related object. In Figure 2.3, the relational model has been used to design the relational database structure shown.

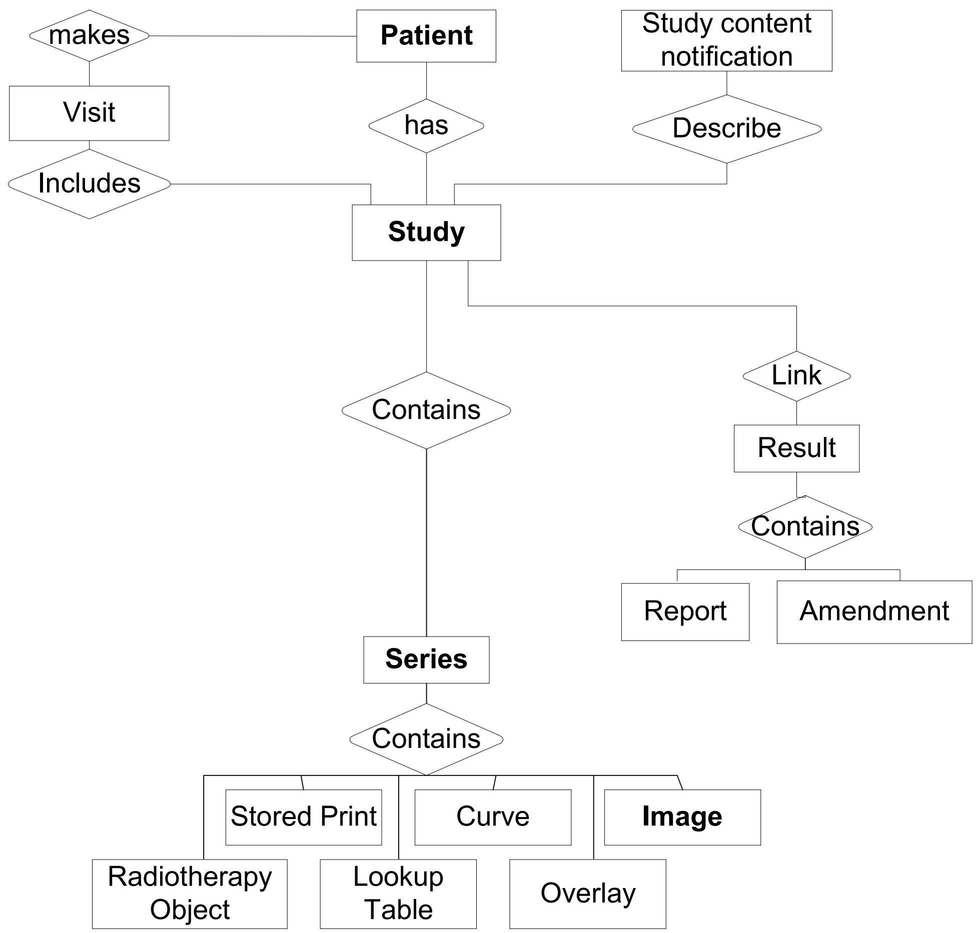

Figure 2.2 Relational model of the DICOM data structure (modified from[65]).

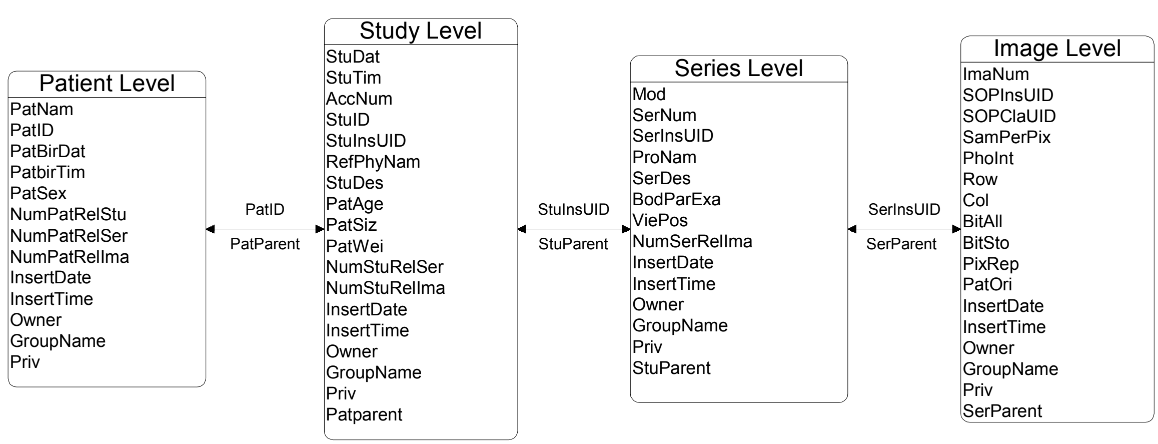

Figure 2.3 Relational database structure for DICOM.

In Figure 2.3, each rectangle is a different table in the relational database and contains its own fields. The arrow between the tables is the link. The field name under the arrow is a means to find the object related to the current one. For example, to find all the related studies for a given patient, we have to search the studies with "PatParent" fields equal to the desired patient ID. In the same way, all the series included in one study can be found. This is done to avoid redundant information for each lower level object.

A DICOM image header, like the one illustrated in Figure 2.4, is attached to each image and contains the essential data that characterize it. In fact, this example is a human readable version of the header. The real version is formed by an uninterrupted flow of data as shown in Figure 2.5. In hexadecimal base, the data fields are coded as follows: the first four digits represent the object group and the next four digits specify the object element. The eight

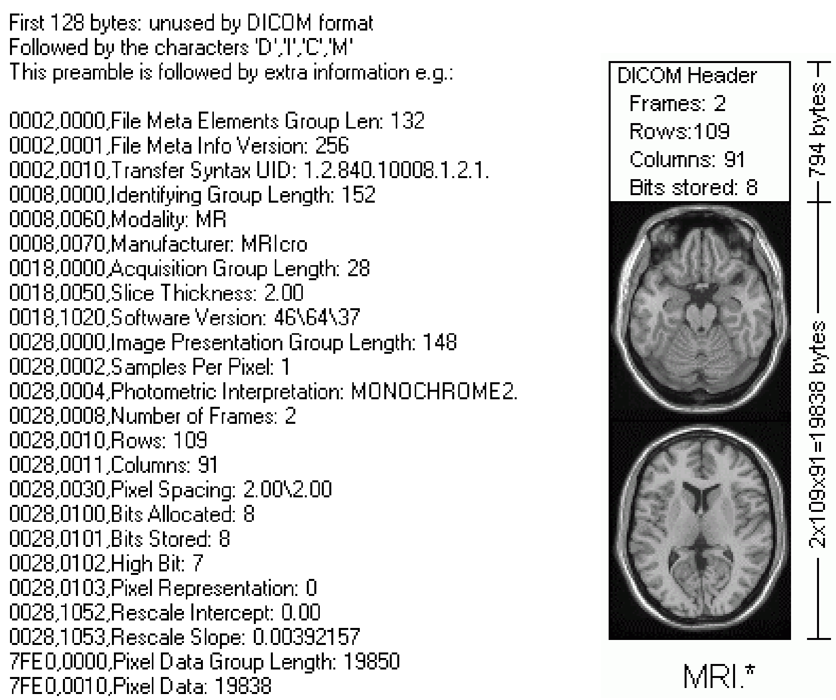

Figure 2.4 Example of DICOM image header in human readable format.

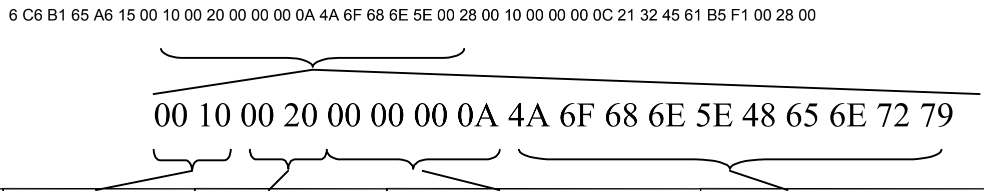

Label	Group identification	Element identification	Length of next field	Field value
Value	Patient	Name	A (hex) = 10 (decimal)	John Henry (10 digits)

Figure 2.5 Data flow in hexadecimal and its translation into human readable format.

following digits give the length for the field value. This is repeated continuously until we come to the pixel data. The way the image is numerically encoded is already part of the header information. All of the groups and elements are listed in a common dictionary included in the DICOM standard.

The DICOM interface between the application and the network ensures a good conversion between the execution of a process and the data flow.

2.1.4 Network

The network provides the link to transmit data to other computers. Despite the simplicity of this task, there is an incredible amount of complexity due to the large number of protocols and standards defining transfer, routing, addressing and hardware. It is not the purpose of this thesis to study the details of network operations and protocols. Generally, in a PACS implementation, the network is managed by an Information Services department, which is responsible for its use throughout the entire hospital[66]. At the MGH, only the Transmission Control Protocol and the Internet Protocol (TCP/IP) are used over fast Ethernet (full or half duplex, 100 MB switched). Although Ethernet and TCP/IP are most commonly used, Asynchronous Transfer Mode (ATM) and Fiber Distributed Data Interface (FDDI) may also constitute the physical layer of Local Area Networks (LAN) used for PACS[50, 64, 67-69]. The open systems interconnection (OSI) model from the Internationnal Standard Organization (ISO) is an alternative to TCP/IP[69], but is less widely used.

2.1.4.1 Communication protocols

TCP/IP is defined by the Internet Engineering Task Force (IETF) of the Internet Advisory Board. TCP/IP works in layers and converts the data flow from an application into a series of packets that can be transmitted from the network card, routed to the destination and reconverted into the original data flow on the receiving side. The TCP layer is responsible for forming the packets and ensuring their correct delivery in proper order. In case of data lost in the transfer, TCP will trigger re-transmission until it is successful. The IP layer provides the address information to the packets and ensures that the packets are routed to the right computer host through the Internet. The IP address is a unique four bytes that identifies all computers connected to the Internet or an IP-based intranet.

2.1.4.2 Physical network

The LAN usually extends to the size of a large building or a few smaller neighbors. There are many different ways to connect all the computers within this area. In the beginning, almost every company offered its own networking technology with a proprietary hardware configuration. Ethernet standards have defined several generations of physical network, its hardware and its operation using a Carrier Sensitive Multiple Access with Collision Detection (CSMA/CD) communication strategy for transmitting packets efficiently. With this strategy, any node connected to the network can transmit its packets at any time, but it must detect collisions with other computer communications. When collisions occur, the packets are sent again in such a way that chances of a second collision remain low. Fast Ethernet still uses this principal with a star topology where each host is connected to a switch with Unshielded Twisted Pairs (UTP) of copper wire. The switch controls packet traffic between all of its ports and is linked with a larger bandwidth link to the backbone of the network.

2.2 Types of image transfer

The different types of image transfers used in a PACS are adapted to the workflows of the radiologist and other clinical users. Autorouting is a mechanism that sends newly acquired images to where they will first be interpreted. Prefetching refers to the automatic retrieval and delivery of prior relevant studies to allow side-by-side comparisons to be made during diagnosis. Autorouting and prefetching try to satisfy the needs of the radiologist by sending the images directly to their workstation before they have to request them. But, it is also important for the user to be able to randomly retrieve any image he may want. In the following sections, the three main types of transfers are described with their rationale and the process used on the MGH PACS.

2.2.1 Autorouting

Autorouting is applied to all new images produced on the PACS, to ensure that a physician or a radiologist reviews each study at least once. The destination of these new images is selected from a routing table. This routing table describes a series of tests which are applied to every newly acquired image. The image may be routed according to the examination type, the name of the referring physician, the time of acquisition, or any other information included in the image header. The tests in a routing table depend on the organization of the radiology department. Large departments with specialized radiologists may route the image according to anatomical specialization. Smaller centers with fewer radiologists might plan the routing table so that all images are sent to radiologists according to their schedule, and vary by date, time and physical location.

The images could also be sent to more than one review workstation at the same time. In light of this, it could be tempting to send all the images to every workstation. However, this practice would have a negative effect on the network, particularly on the network connection to the server. This unjustified load would slow down more important transfers and unnecessarily fill the hard disks on the workstations.

Therefore, the perfect routing table would send the image to the workstation with the highest probability of the image being reported. The advantage of this practice is that the waiting time is minimized since the radiologist does not need to download the image from a remote server. When the image is stored locally, it can be displayed almost instantly.

Performance evaluation of the autorouting process can be measured by the perception of the user, the report turnaround time and the network load. The user will perceive the time needed to load the image. For emergency cases, the regular flow of operation could be interrupted for one particular patient. In these cases, the transfer time will be noticeable since the images are expected to be read soon after the arrival of the patient. The transfer time does not only depend on the PACS performance itself. It also relies on the ability of the acquisition system to output the images.

The turnaround time of the report is the time from the arrival of the patient at the radiology department to the time the final report is made available. This time is critical, especially in the case of hospitalized patients who may be waiting for the results in order to move to the next step of their treatment. An unreliable autorouting process could extend the turnaround time. In normal situations, the transfer time should not influence the turnaround time because the few minutes required are negligible in comparison with the other processes needed for the completion of the report: image acquisition, interpretation of the image, report dictation and transcription.

The autorouting process developed and used at the MGH is shown in detail in Figure 2.6. The program "image_server" handles the DICOM connection as well as the reception and storage of images. Image_server also triggers several operations related to the arrival of the new image.

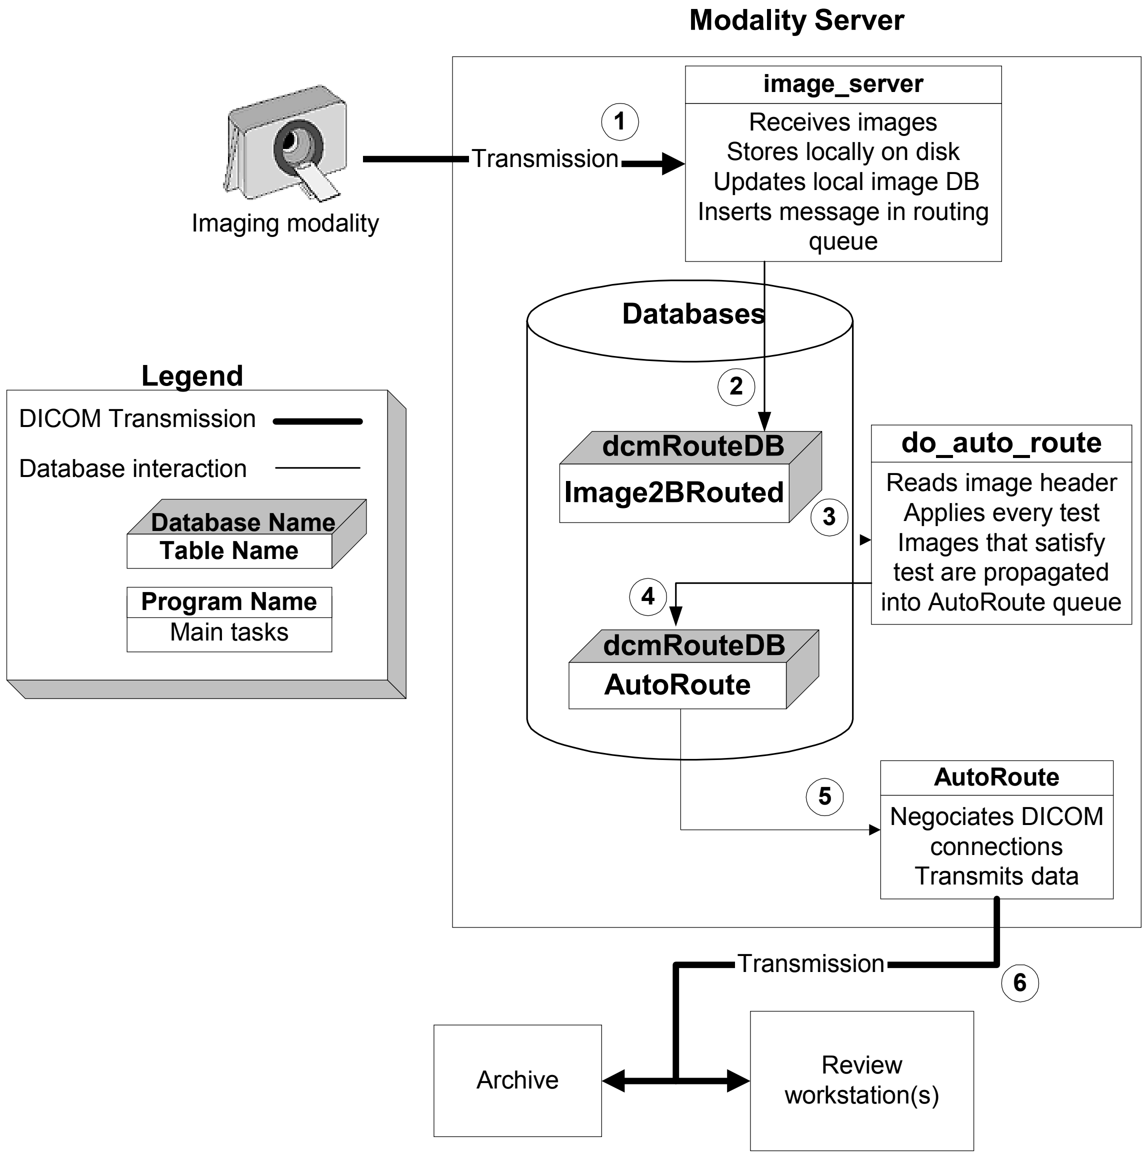

Figure 2.6 Autorouting process and its main steps. (1) The program "image_server" running on the Modality Server receives the images from the imaging modality; (2) "image_server" inserts the image identification information in the table "Image2BRouted"; (3) The program "do_auto_route" reads the DICOM header of the image in "Image2BRouted" and applies the tests described in the routing table; (4) Images that satisfy routing tests are inserted in the table "AutoRoute" along with the intended destination(s); (5) The program "AutoRoute" reads the "AutoRoute" table to select images for each destination; (6) A DICOM connection is established, and the images are transmitted.

We will see that the process of receiving an image is complex, as we cover the different tasks performed by "image_server". For the moment, we will focus on the operation related to the autorouting process.

Database tables are used as queues for the transfer process since they are more reliable than temporary memory allocation. Should equipment fail or transmission be cut-off, the system can recover and start where it left off in the queue. Once a request is inserted into the table, it will not be deleted until the image is transferred or a "time-out" period has elapsed.

The "AutoRoute" program sends images to different destinations at the same time, but only one connection at a time is established with a given destination device. The different connections share the bandwidth available on the server. In most cases, each image is sent to a few relevant workstations and to the archive for long-term storage.

2.2.2 Prefetching

The prefetching process is a common practice in most conventional radiology departments. When a patient is scheduled for an exam, the film library pulls in advance his master film bag, containing the films from previous studies. This strategy has been adapted to PACS in order to keep and improve the same efficient workflow. Automated prefetching increases productivity and does not preclude manual retrieval when desired.

Prefetching consists of pushing the prior relevant studies where they are needed before they are actually requested. Pushing these studies prevents the user from waiting for a real-time retrieval request that could take particularly long if the data must come from the archive. Based upon the preference of the radiologists, it is possible to anticipate what types of studies certain users need. Ideally, prefetching strategies should be used to satisfy these (predictable) needs.

Many criteria can be used to select which studies are of interest, but the most common is based on anatomical region. This facilitates comparisons to see the evolution of a disease or the efficiency of a treatment.

Outpatient consultations at the clinics are usually scheduled in advance. Some clinics refer more often to previous images than others and benefit from a similar form of prefetching. In this case, the trigger event is based on the schedule of patient appointments at the clinic, not the acquisition of new images or the schedule in Radiology.

Another common practice is to put all of the available images of a patient online while they are hospitalized. However, this practice is a less efficient mode of prefetching since the studies are not sent to a workstation, and they may include prior studies that have no relevance to the current reason for hospitalization. This mode only saves on the retrieval time from offline storage. This kind of approach requires a level of integration between the PACS and the HIS to keep track of admissions and discharges.

In autorouting, the real transfer time is not important for the user because it is not perceived. The most important aspect of the process is the efficiency in providing to the users what they need at their workstation before they even have to ask for it. Again, the administrator could be tempted to send almost everything, everywhere. Choices must be made to avoid overloading the system, the network, the server and the user workstations.

The prefetching processes are responsible for pushing prior relevant studies to workstations where comparative studies might be performed. In the MGH PACS, there are two ways to do this. The first occurs automatically, in parallel with autorouting and is triggered by the acquisition of a new study. The second is triggered by an external scheduling system that organizes patient visits.

The first method, called "auto-prefetching" is illustrated in Figure 2.7. It is very similar to the autorouting scheme describe earlier. Auto-prefetching uses a system of database queues and different programs to select and deliver images to their destinations. This auto-prefetching feature only manages the previous relevant studies related to the current study that is being autorouted.

The number of previous relevant studies will depend on the clinical history of the patient. Usually, one or two of the most recent studies are selected then sent to the same destinations where the current study is being autorouted. The prior studies are requested by the program

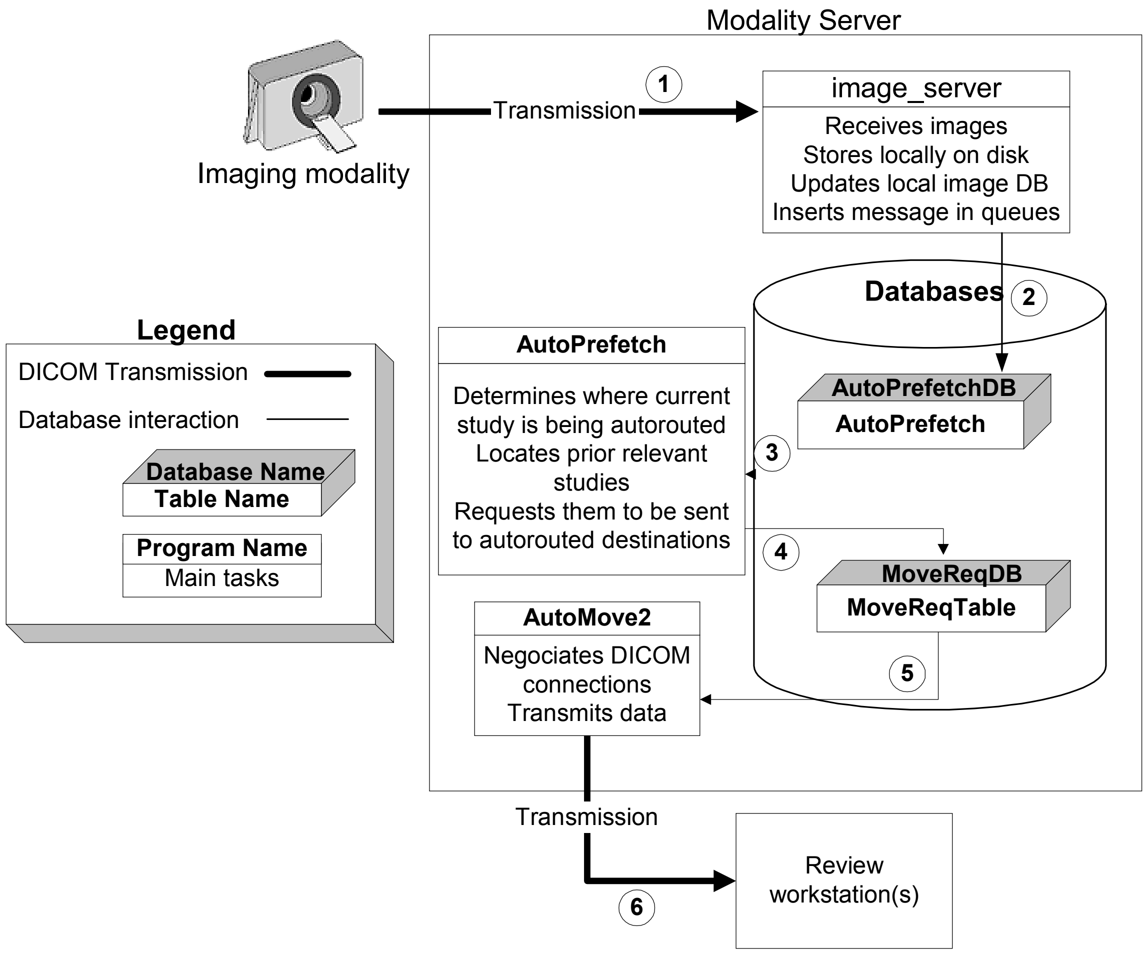

Figure 2.7 The auto-prefetch process and its main steps. (1) The program "image_server" running on the Modality Server receives an image from an imaging modality; (2) "image_server" inserts the image identification information in the table "AutoPrefetch"; (3) The program "AutoPrefetch" reads the DICOM header of the image to determine where the image is being autorouted, then determines if there are any previous relevant studies for the given patient; (4) Prior relevant studies are located, and AutoPrefetch requests them to be sent to the current routing destinations by inserting requests into the local (or remote) database "MoveReqDB"; (5) The program"AutoMove2" reads "MoveReqDB" to select images for a destination; (6) Then a DICOM connection is established and images are transmitted.

“AutoPrefetch” by inserting a move command into the transmission queue, “MoveReqTable,” that is processed by the program “AutoMove2.” Note that the prior relevant studies may reside on multiple servers, in which case move commands would be inserted in the “MoveReqTable” on each server. Note also that “MoveReqTable” is not only used for prefetching; it is the general queue for most of the retrieval processes on the Modality Server.

A second method of prefetching is triggered by certain messages received from the HIS. These messages contain scheduling information for patients who are visiting certain clinics in the hospital sometime in the future. The HIS sends the PACS the name of the clinic; the patient's name and ID and the date of the upcoming visit (plus a flag indicating if this visit is being cancelled). The PACS maintains a table that lists all known DICOM workstations residing in each clinic. Overnight before the scheduled visit, the PACS retrieves the patient's entire digital folder and sends the images to each workstation in the clinic. The goal, like the prefetching described above, is to deliver the previous images to the users in advance to minimize the need for real-time (slower) retrievals.

The third method is found on the PACS browser with the interface shown in Figure 2.8. This method has the disadvantage of requiring manual operation. In fact, it is a utility that allows users to make “post-dated” moves (i.e., to specify retrievals to be performed at some time in the future). It is used primarily by clinics that have scheduled patient visits but have no means of communicating this information to the PACS. In order to minimize the network load, transfers are done overnight, while the PACS activity is low. When the request is submitted, the information about the image, the destination and the date selected are inserted in the “MoveReqTable” which is processed by AutoMove2. When the patient arrives at the clinic, his folder is already at the workstation. However, for most cases at the MGH, patient schedules are integrated with the RIS, so no human intervention is necessary.

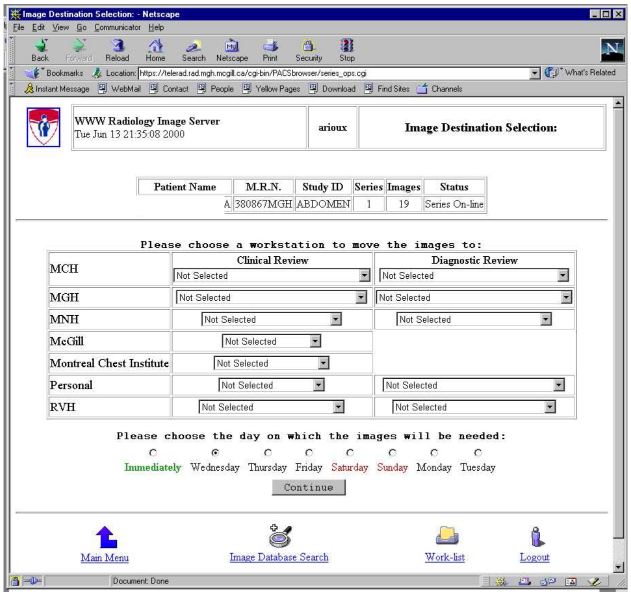

Figure 2.8 Interface of the PACS browser for manual prefetching. After selecting the desired study, the user selects the destination and the day the images are needed.

2.2.3 Manual retrievals

It is not always possible to foresee the needs of the users. A PACS must allow users the option to retrieve any study they want. It is impossible to work in a radiology department without a retrieval mechanism of this sort. In fact, the other transfer processes help improve the performance by limiting the frequency of manual retrievals. Unfortunately, they do not eliminate them.

Manual retrieval is a critical process in terms of image transfer. Because the user is initiating the transfer, he expects quick feedback to know that the transfer is really beginning. Most users do not tolerate a long delay for complete transfer of the study requested. Ideally, any image should be accessible anywhere, at any time, in a few seconds. But in reality, there are technological and economic limitations that make this ideal hard to achieve. Better overall performance is achieved by using intelligent strategies for image transfer and storage.

The transfer process for manual retrievals is more complex than other processes. There are multiple ways to request the image from different types of storage locations. In the distributed architecture used by the MGH PACS, images can be stored on any Modality Server, or in one or more tape libraries. In order to keep track of study locations and information related to each study, a master database is used. In order to illustrate this process, a simple situation is given as an example, and this will become more complex as the example is developed further.

The simplest retrieval occurs when the user needs an image from a specific Modality Server. This type of retrieval is illustrated in Figure 2.9. The user selects the server using the retrieval facility of his review workstation. Usually, the user needs to search the server for the desired data. Filters may be used to narrow the scope of the search. Efficient searches might use fields like the date, the name of the patient, the ID of the patient, and the referring physician. These fields are the most popular, but other information included in the DICOM header of the image might also be used. Once the query yields a result, the user selects the data he needs and submits a retrieval request.

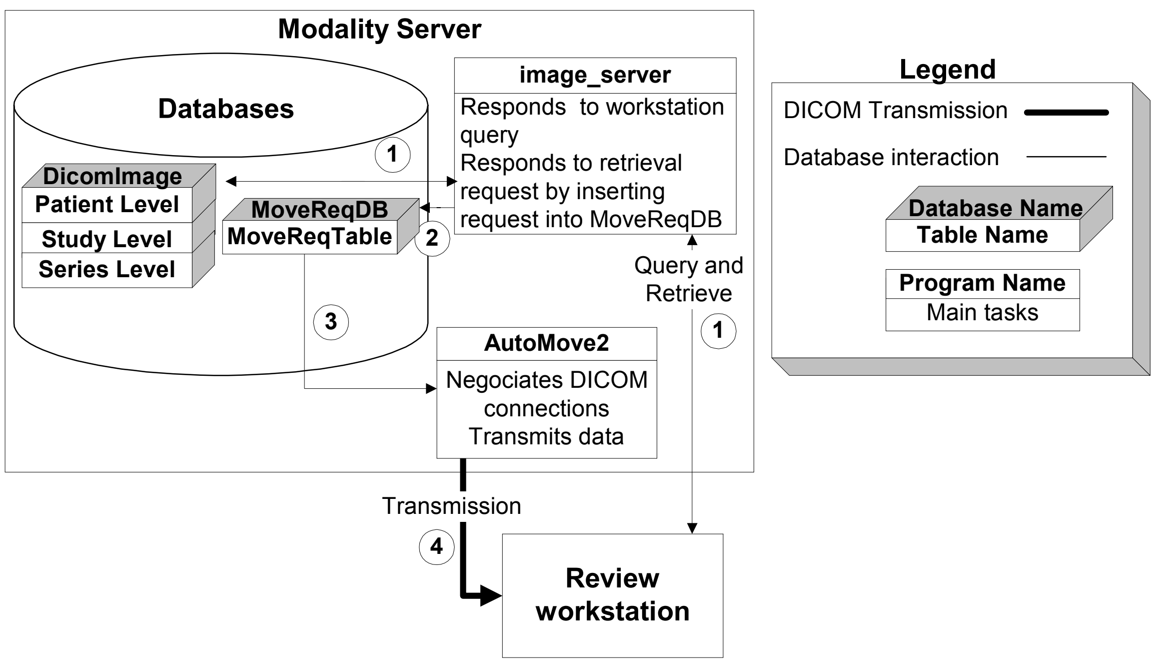

Figure 2.9 Process for a manual retrieval of an image from a Modality Server selected by the user. (1) The program "image_server" running on the Modality Server handles the connection with the workstation and brokers the query and retrieve dialogue with the local DicomImage database; (2) "image_server" inserts the image identification into the "MoveReqTable" table; (3) The program "AutoMove2" reads the "MoveReqTable" table to process move requests for a destination; (4) A DICOM connection is established and images are transmitted.

In this situation, it is assumed that the user knows the source from which to retrieve the study. This scheme is only practical when there is a limited number of well-trained users to work with a PACS containing a small number of servers. It becomes obvious that to be efficient, this distributed architecture must be transparent to the user. In order to achieve this, a master database server provides access to information from all the servers. Then, the user needs only to query one server. Once he issues a retrieval request, the order is communicated to the server deemed best suited to respond most quickly, as shown in Figure 2.10. The master database allows the users to have the same interface as with a single Modality Server while accessing any study held in the whole system (even those that are archived).

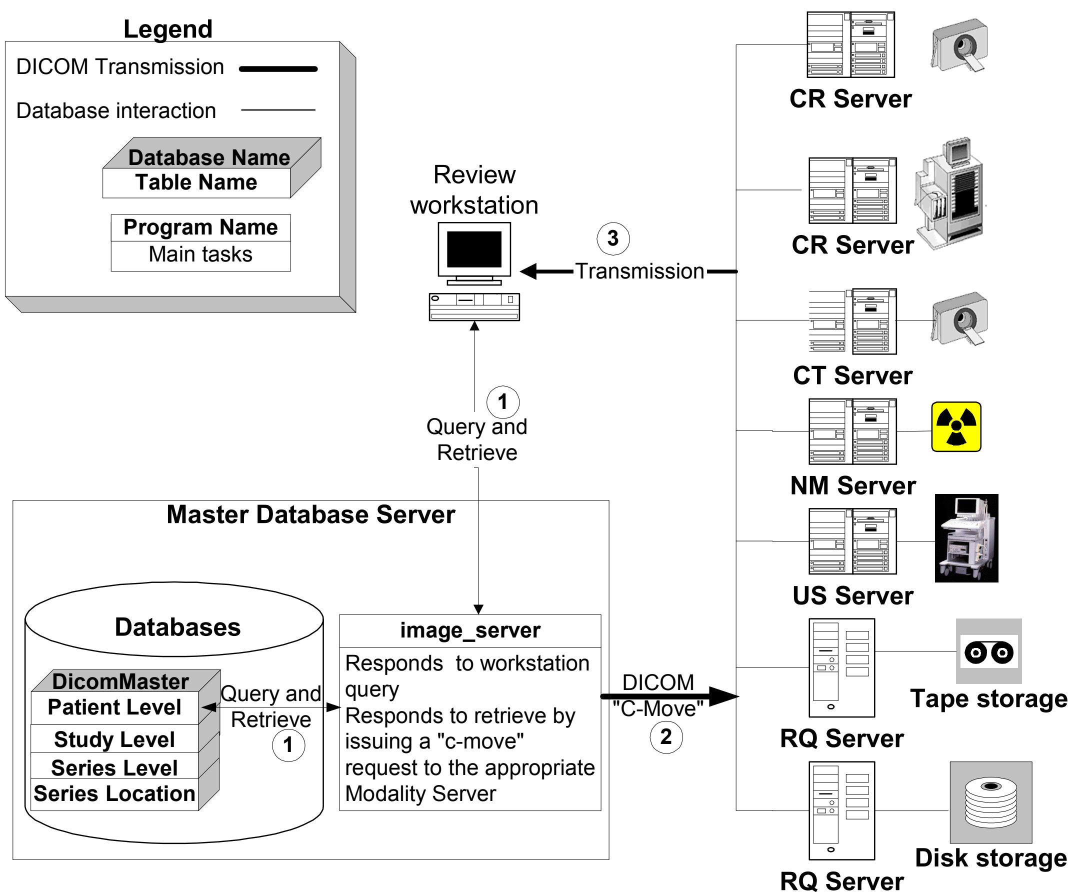

Figure 2.10 Master Database interaction with other servers. (1) The program "image_server" running on the Master Database Server handles communications with the workstation and brokers the DICOM query and retrieve dialogue with the Master Database (DicomMaster); (2) "image_server" issues a DICOM "c-move" request to the appropriate server; (3) A DICOM connection is established and images are transmitted by the program"AutoMove2" from the selected server.

The Master Database Server is directly queried by the user's workstation. The server translates the request in SQL language and gets the results from the database. These results are then transmitted in a reply to the workstation. The dialogue with the database is usually a multi-step process where the specific patient, study and series of images are selected. The user then submits the retrieval request. The program "image_server" determines the best source from which to obtain the requested study (or studies). This information is stored in the SeriesLocation table in the DicomMaster database. The images may be stored in multiple locations. Faster storage media such as a RAID attached to a Modality Server will be preferred over an archive medium like tape. The request is communicated to the appropriate Modality Server(s) via a DICOM "c-move" request. The program "image_server" on each Modality Server receives this request and inserts it into its local copy of the MoveReqDB database. Automove2, as described earlier, processes the queue and transmits the images.

The process of retrieving images described above was for a user working from a workstation directly on the PACS. The same process is possible from any computer with Internet access. The PACS browser allows users to move images to any review workstation on the PACS or to view the images within the browser. Because the Internet is generally slower than internal networks, the images, when viewed with the browser, are stored by default in a compressed format (JPEG) to minimize the transmission delays.

The retrieval process for the PACS browser is similar to the one described above. The web application communicates with the same databases to issue study retrieval requests. A central web server services most of the user queries. When the images are requested to be viewed on the Internet, the user is automatically redirected to the right Modality Server. This server then sends JPEG images from its RAID to the user with the same web interface and capability as the web server. This practice allows all the images stored on a RAID attached to a Modality Server to be available, and several web users to be distributed on different servers.

2.2.4 Retrieval from long-term storage

When the Modality Server RAID nears full, it automatically discards the oldest studies after verifying that they have been properly archived. All the prefetching and manual retrieval methods described in the previous section are able to access these studies stored offline in the long-term storage. Additional steps in the retrieval processes are invisible to the user but result in noticeable delays.

Several long-term storage technologies were used through the different phases of the implementation of PACS at the MGH. First, the images were stored on recordable CDs that were loaded manually into the CDROM drive of the Retrieval Server upon request. In June 1998, a DLT library was installed and has since been filled. Most recently, a Sony Digital Tape Format (DTF) library of larger capacity was installed (in March 2000). All of these archive systems (CD-R, DLT and DTF) are still in operation, but only DTF is used for ongoing archival. The master database orients retrieval requests to the proper archive system as it does with all the Modality Servers, but retrieval requests are only passed to the archives if the studies cannot be found elsewhere on the RAID of a Modality Server or Retrieval Server.

On an archive/retrieval system, the retrieval request is put in a special queue named DCM_REQUEST_DB. This database is read by software running on each archive server which subsequently processes the request. Tape libraries are equipped with a robotic arm that loads the required tape into an available reader drive. The tape is then initialized and positioned at the beginning of the requested series. After that, the series is read and written to the RAID of the Retrieval Server. In the case of CD retrieval, a clerk in the film room is prompted with the required CD number, and loads it into the CD drive of the retrieval computer. Once the series are on disk, the conventional transfer process is initiated with the insertion of the retrieval information in the MoveReqDB queue, and the program AutoMove2 finally sends the images to the workstation.

2.3 Summary

This chapter has described the communication stack and the factors that influence communication speeds. Also shown was how the DICOM standard faced compatibility problems between devices that exchange medical images by providing standard data structure and communication protocols. PACS not only allows access to archived images, it also includes management applications that automatically deliver studies as part of the workflow in a hospital. The image transfer mechanisms specifically used in the PACS of the MGH were reviewed.

Chapter 3

3 Methods

Performance monitoring is a vast topic because of the wide number of variables affecting PACS. In order to make system improvements, it is important to begin with a general approach and then focus on specific improvements where they are relevant and achievable. In order to get this general overview of the system, the load was evaluated at the input level. Then a residency timing analysis was used to measure the time spent in each step of the relevant processes. In order to have a quantitative analysis of the PACS, performance tools were developed to find the waiting time distributions.

3.1 Monitoring of the PACS load

The PACS activity level is difficult to evaluate due to the distributed architecture of the system. There is no single point where all the activity can be measured. The evaluation of the load is expected to provide valuable insights that should help predict the levels of activity on other systems. The PACS "load" is basically the data volume handled by the system in the form of the new data (images) entering the system, and user-initiated data transmissions. The activity must also include the tasks performed during automatic routing of images and the prefetching of prior studies.

The load on the system is defined by the technical requirements of the PACS as illustrated in Figure 3.1. Shortly after their acquisition, the new images need to undergo a quality control check, be reported by the radiologist, possibly be re-examined by a referring physician, be archived and kept available online. Users may interact in many ways with the system, including image database queries, information updates, report or teaching file editing, and system management. Requests for image transfers deserve special attention because of the long time taken for the transmission of large images as opposed to text and other operations that are fast and interactive.

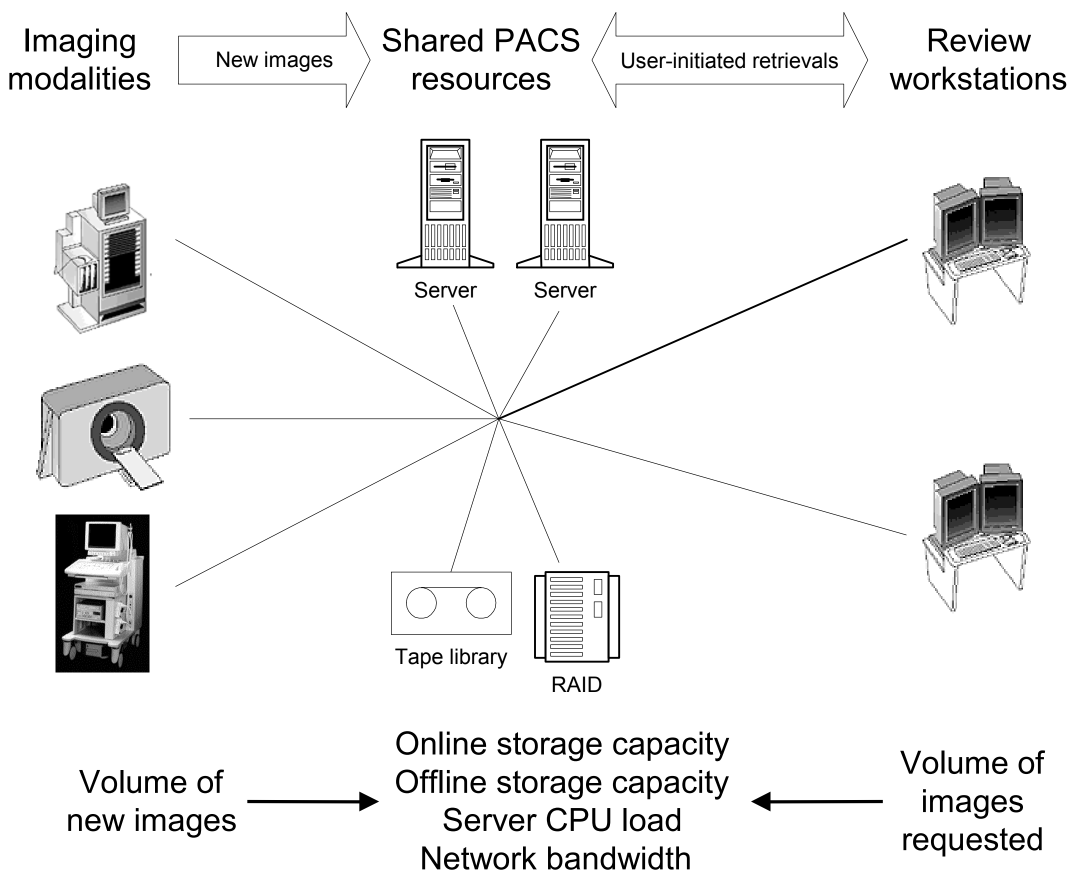

Figure 3.1 The load on the PACS comes from the volume of new images acquired and from user-initiated requests.

New image acquisition and image retrieval are the two main sources of workload, and both processes utilize the shared and limited resources available on the system. The evaluation of the activity on the whole system helps to determine the optimal size of the system for a particular clinical setup. This is particularly important in a phased project, where new imaging modalities and review workstations are constantly added to the PACS. More users than ever have easy access throughout the hospital to an ever-increasing volume of images. Overloaded or underused equipment should be avoided and upgrades planned more effectively.

Two methods of monitoring PACS load were developed. The first one, based on query-retrieve operations allows us to study the load on the system retrospectively. The second method provides, in real-time, the current state of some system parameters of interest.

3.1.1 Retroactive load study

The Master Database contains all of the information pertaining to images acquired on the PACS since its implementation. In order to minimize stress on the database while gathering enough data to ensure meaningful statistics, a period of one week was selected for this study. Only those fields shown in Table 3.1 were extracted for this period.

Table 3.1 Fields used from three different tables in the DicomMaster database.

Table Name	Field Name	Description
SeriesLevel	SerInsUID StuParent Mod InsertDate InsertTime	Unique Series identification number Unique Study number in which the series is included Modality of the image Date of insertion of the image on a DICOM server Time of insertion of the image on a DICOM server
StudyLevel	StuInsUID Owner	Unique Study identification number Server that first registered the study
SeriesLocation	SerInsUID NumSerRelIma	Unique Series identification number Number of images in the series

These data were used to study the input of new images, their size distribution, the interval of time between insertion of new series, and the distribution of data volumes among the different imaging modalities or servers.

Calculations of series size were based on a constant file size for each modality and the number of images per series. Image file sizes were assumed to be constant, equal to the value the most commonly encountered for each imaging modality (MR, US, NM, CT, CR). The file sizes shown in Table 3.2 were determined from the number of rows and columns forming the image and the number of bytes needed for each pixel. File compression was taken into account separately when the evaluation of storage capacity requirement was done. An average compression ratio was used and was taken from a representative population of images. The load evaluation was done relative to the file size since it has a closer relationship with storage capacity and transfer delays than the number of images or series managed.

Table 3.2 Image size most commonly encountered for each imaging modality used at the MGH.

Image type	Rows	Columns	# of bytes/pixel	Image size (MB)
CR	3730	3062	2	22.843
CT	512	512	2	0.524
MR	256	256	2	0.131
NM	256	256	1	0.066
US	480	640	1	0.308

Another database that is part of the PACS records all the main actions performed by the user via the web interface of the PACS. These records were used to study the number of requests for image transfer from servers to workstations. Again, a period of one week was selected for this study. The weekly usage and the distribution by modality were examined.

In order to study the efficiency of the hierarchical storage architecture, the age of the series requested was also examined. The frequency of requests of recent images compared to old images is useful in examining the effectiveness of the hierarchical storage scheme, and helps to find an optimal combination of storage technology and capacity.

These studies were done retrospectively from the master database and the PACSbrowser database, both of which maintain all the information since the beginning of operation. These data also help to illustrate the current state of the system, historical problems, and trends in the PACS.

3.1.2 Real-time web-based PACS Watcher

The PACS Watcher is a tool developed for this thesis to show the state of the Modality and Retrieval Servers via a graphical Web interface. Parameters describing the state of the equipment are frequently updated and provide almost real-time information. The PACS

Watcher is based on software called the Multi Router Traffic Grapher[70] (MRTG). It triggers requests for the state of all equipment and compiles the results in different time scale graphics. The results are shown graphically in web pages. Some HTML editing is necessary to ensure a simplified browsing through all the information gathered. Summaries of the activity levels are available with graphs of the means for the past few days, weeks, months and years. In order to have constant sized logfiles, the data are averaged as they become older and finally rejected after 13 months.

The query trigger used by MRTG could be one of many types. MRTG simply starts an executable and waits for a numerical answer. In the case of network monitoring, because the request goes through the network, it is important to keep the amount of data transfer as low as possible. The Simple Network Management Protocol[71] (SNMP) was specially created for this purpose. It is a standard way to retrieve operational parameters about computers on the network. SNMP agents that answer information queries are installed in most computers and peripherals that could be attached to a network. The CMU-SNMP[72] agent for Linux was installed to monitor the status of PACS servers from their operating system. The agent on the remote device listens for incoming SNMP requests from the monitoring computer. The agent keeps a small database with SNMP object values previously selected. Upon request, the agent refers to these values and responds using the same communication protocol. This method is efficient, and not overly demanding on the network and the observed computer.

SNMP was used in this thesis with MRTG to monitor incoming and outgoing network traffic, the CPU usage, the number of processes running, disk space, and the memory used. In order to compute the network transfer rate of the server, a counter of bits is sent in the SNMP answer. MRTG refers to the value of the counter at the latest acquisition and its time stamp to find the incoming and outgoing rate. For the other SNMP queries, the value at the time of the request is simply sent and added to the MRTG graph. In the case of the CPU, it is the percent of usage over the last minute. Disk space monitoring was only done for the RAID that receives the images. The boot disk that stores all the applications was not monitored.

The parameters that can be obtained via SNMP requests are mainly low-level technical data and do not give information about incoming image transfers or requests. In order to get information at that higher level, other methods must be used. Database queries were employed for values already available in the database. Special attention was given to design efficient queries. Queries of each server database provide the information about arrival rates of new images by counting the number of images inserted in the local database (DicomImage) since the last update of the MRTG graphs. In a similar way, the number of images waiting in transmission queues read by the two programs, AutoRoute and AutoMove2, were also examined.

A web page was designed to allow browsing through the data on a per server or per parameter basis, as shown in Figure 3.2. The first level gives an overview of the daily data for all the servers and all the parameters. A second level of details is obtained when one chooses a daily graph of interest. It shows the summary for the last week, month, and year. A login and password limits the access to this data.

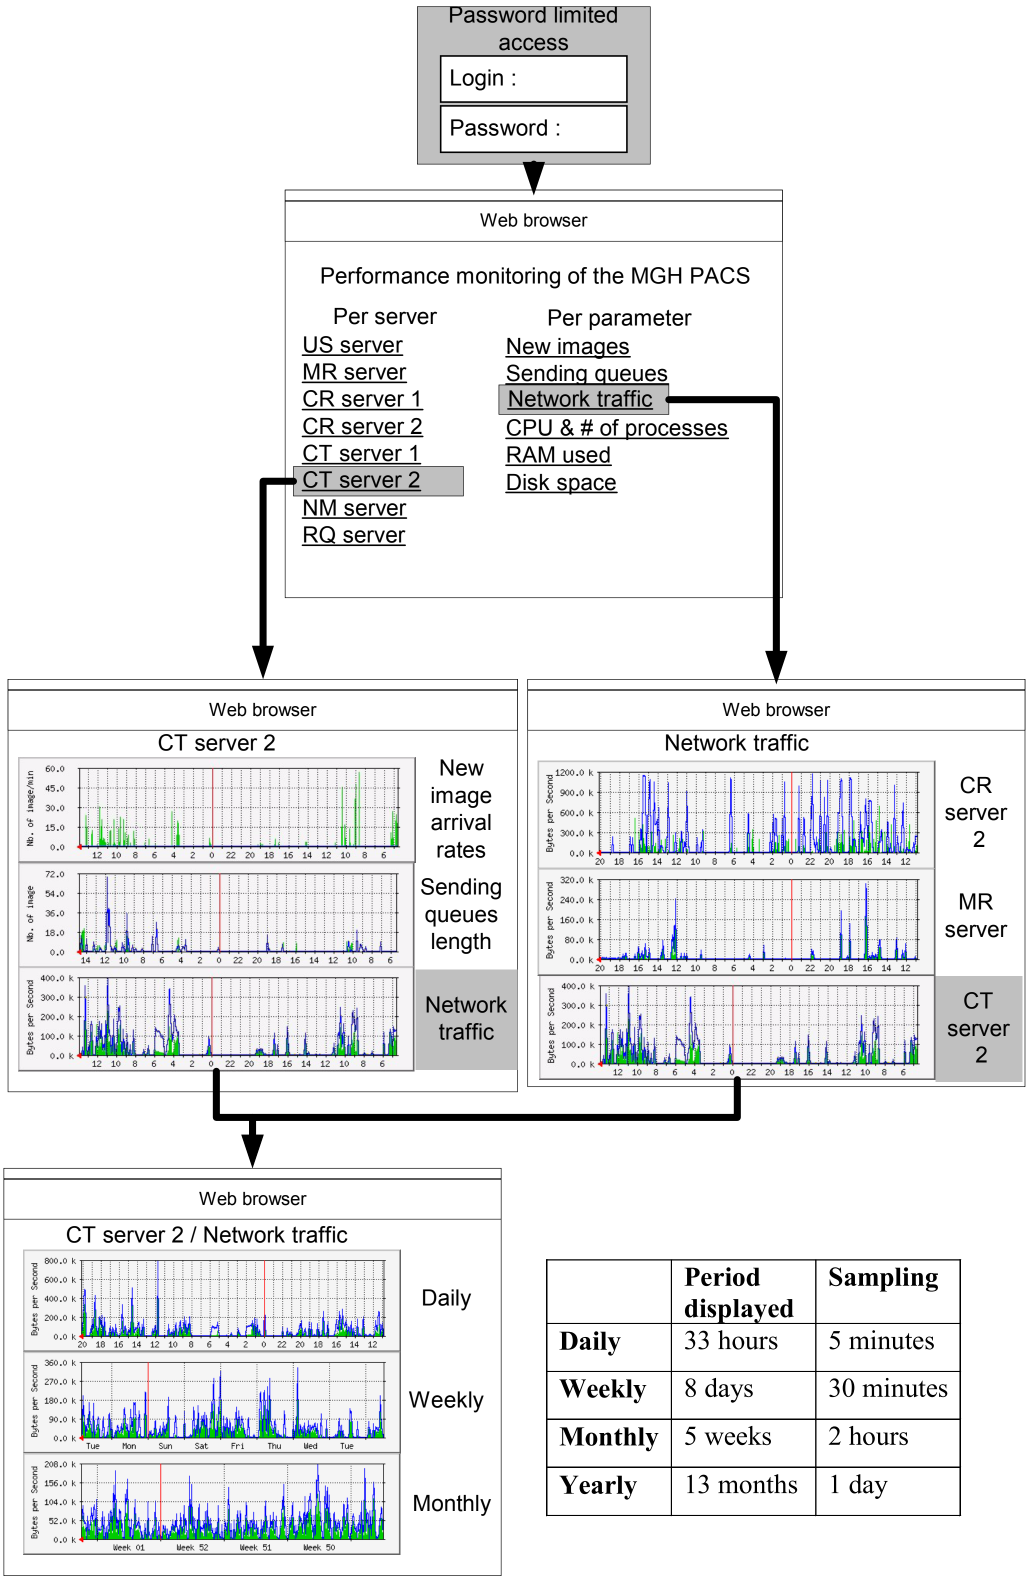

	Period displayed	Sampling
Daily	33 hours	5 minutes
Weekly	8 days	30 minutes
Monthly	5 weeks	2 hours
Yearly	13 months	1 day

Figure 3.2 Web site map for the real-time monitoring tool.

3.2 Storage strategy evaluation

In order to provide fast access to the most recent images acquired, each Modality Server is equipped with large capacity hard disks configured in a RAID 5 configuration. Each RAID keeps available the newest images sent to the server. The older ones are deleted as the RAID fills on a first-in-first-out basis. On first reception of the images, they are autorouted to several review workstations and to the long-term storage for archival. This long-term storage is a tape library controlled by a few servers dedicated to reading and writing data to the tapes. Each archive server is equipped with its own RAID which allows image retrieval from the tape library to remain available online for a few weeks, and it buffers the data that needs to be written while the tape library is busy or unavailable. This hierarchical storage structure is illustrated in Figure 3.3.

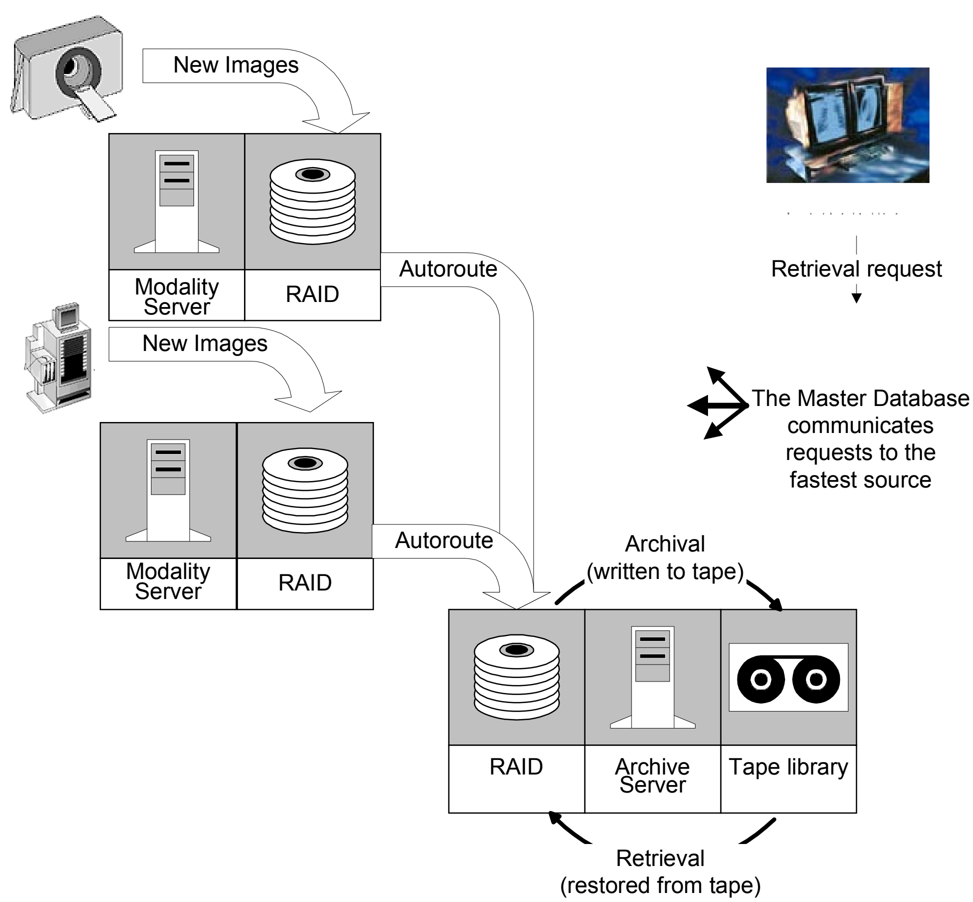

Figure 3.3 Hierarchical storage structure used at the MGH.

The time a study remains available online can be found for each Modality Server. First, the daily averaged incoming data volume was determined using the retroactive load study method. Then, using the storage capacity (in MB) of the RAID attached to the Modality Server, and considering the average compression rate of its images, an estimate of the time the images stay on the RAID was computed. This time of availability online represents the time an image is usually kept on the Modality Server before it is deleted.

The delay between acquisition and retrieval was collected to evaluate which ratio of requests is satisfied online for a given time of availability online. These delays were obtained from the database logging the action of the user via the web interface of the PACS. Only manual retrievals via the web were considered, and these were assumed to be representative of the other retrievals done directly from the review workstations. The delay was found by subtracting the acquisition date of each image from its date of retrieval. The data were collected for one week, and then plotted in a cumulative histogram.

3.3 Residency timing

Residency timing is a way to study a process using the time the process takes to reach each step. This study was performed for processes where the delays were apparent to the user; that is, the different image transfer types.

The detailed description of image transfer processes in the second chapter shows all of the steps a request goes through. In order to get the residency timing, timing information is needed. Software probes were put into the programs managing the images to get the time at which particular steps begin and end. These software probes provided the information about the steps in the image transfer process in a file. These additions to the software that are providing valuable timing information about the flow of image in the system will be refer further in the text as the timing probes.

The data produced by the timing probes were entered in a database and organized to make analysis easy. Once the complete history of the images was reconstituted, the evolution of the processes for multiple images was easily shown. These data provide information about the throughput of different steps of the processes under clinical loads. Within a typical day of operation, stable operational values are readily visualized, and discontinuities better analyzed. It is then possible to evaluate where and how the time is spent in normal operation and where certain problems may be coming from.

3.3.1 Residency timing in the autorouting process

The timing probes were inserted in the software involved in the autorouting process. Their locations are shown in Figure 3.4.

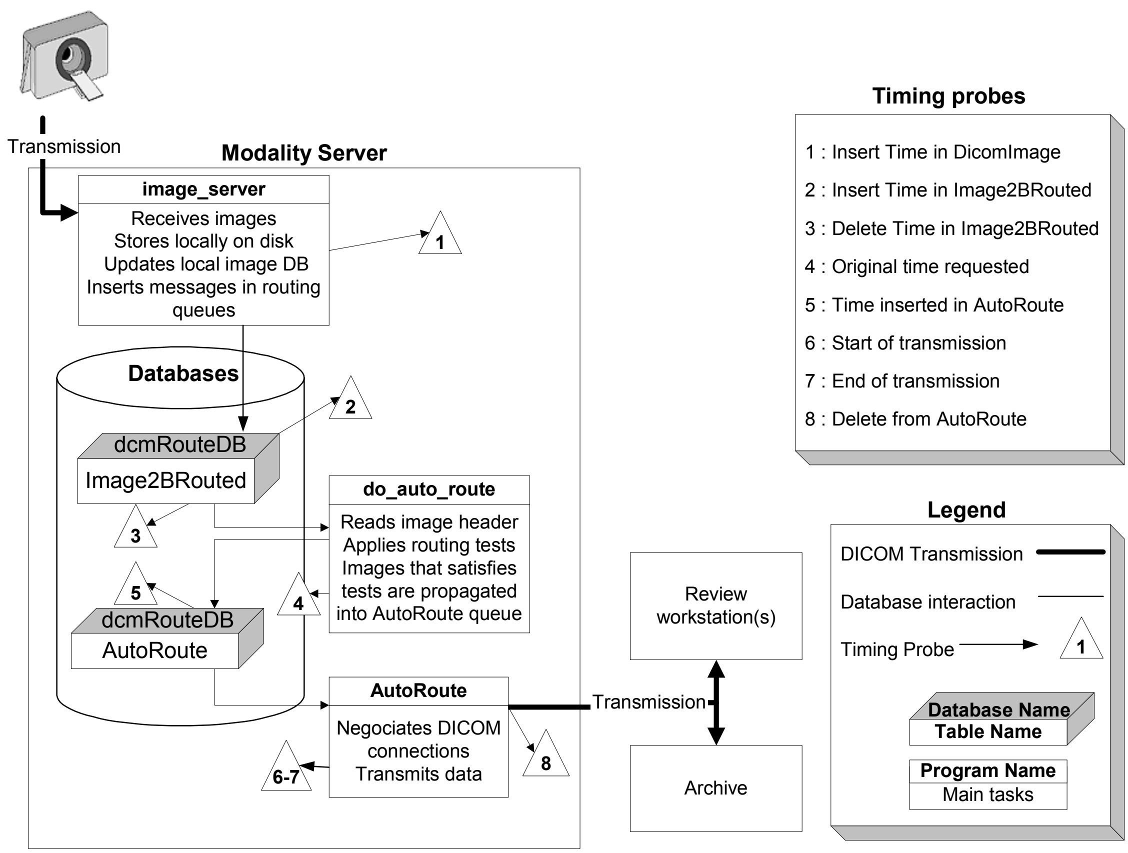

Figure 3.4 Distribution of the timing probes throughout the autorouting process.

The data obtained from the timing probes were gathered over several days. The data were grouped by server to enable observation of their stability of operations when receiving their respective load of incoming images. This also allowed the calculation of the mean transfer time and the frequency of usage of each modality. The values of most interest were found to be the total time spent between the reception of the image at the Modality Server and the end of the transmission. Also of interest was the time between the beginning and the end of the transmission.

3.3.2 Residency timing of online retrieval

Requests for retrieval of images stored online on a RAID are handled by the program AutoMove2. Timing probes were inserted into this program in order to evaluate its efficiency. Only successful DICOM associations and image transmissions were logged with their starting and ending times. In addition to these two timestamps, the file name, the name of the user who made the request, the destination, the version of the sending program, and the time at which the decompression of the image file started and ended were included for each image sent.

For DICOM association negotiation, only the name of the destination computer was available in addition to the starting and ending time. DICOM associations were established at the beginning of each transmission with a new destination. Then the images to be sent were sequentially decompressed and transmitted.

3.3.3 Residency timing of retrieval from long-term archive

Two automatic long-term archives are currently in use at the MGH. A Scalar 458 DLT library from ADIC (Advanced Digital Information Corporation, Redmond, WA, USA) has been in use since June of 1998, but is now filled and used only for retrieval. This jukebox, shown in Figure 3.5 (a), contains 48 Digital Linear Tapes (DLT) for a total capacity of 1,680 MB. A new tape library was installed to satisfy increasing storage needs. A Sony (Sony Corporation, Tokyo, Japan) PetaSite DMS 8800 DTF2 tape library was chosen. This particular model, shown in Figure 3.5 (b), is expandable to accommodate growing storage requirements. The smallest unit consists of two modules installed with a total capacity of 26,000 MB. The system is expandable to 11,200,000 MB.

Figure 3.5 Tape library for long-term storage: a) ADIC Scalar 458; b) Sony PetaSite DMS 8800.

Dedicated servers manage image archival and retrieval from these libraries. In addition to the usual software configuration used on the Modality Servers, a separate program controls the robotics and tape drives. Timing probes were inserted in this program to examine each step of the retrieval process: insertion in queue, robotic moves, tape positioning, tape reading, and request of the transfer process.

Care was taken in the comparison of the two tape libraries. The configuration of the server, the robotics, the tape drives, and the tapes themselves are different for each system. Firstly, the Sony PetaSite has four servers: One retrieval job dispatcher that also controls the robotic movements, and one retrieval server connected via SCSI cable to each of the three PetaSite tape drives. In contrast, the ADIC is connected to only one server that coordinates all of these functions. Secondly, the PetaSite tape drive caches tape positioning (forward or reverse) requests while the ADIC does not. This made it virtually impossible to extract any useful information from the tape-positioning timing probes for the PetaSite. Instead, the time is included in the reading time. Finally, since the ADIC is already full it does not do any data archival, while the Sony PetaSite must do both archival and retrieval.

3.4 Summary

This chapter has presented the methods used in this thesis to evaluate the PACS of the MGH. The three methods presented were designed to provide more information about the transmission of images from the distributed backbone of the PACS throughout the hospital. Queries of PACS databases and continuous monitoring of system parameters provide valuable information on the usage of the system and its impact on system parameters. A second study was designed to look at the time of availability of the study on-line and to allow the evaluation of the storage strategy. Finally, the insertion of timing probes in the applications responsible for transmission generates useful information about the data-flow. Each data collection method provides insights into the operation of the system while it is under clinical loads.

Chapter 4

4 Results

This chapter presents the results of the studies described in Chapter 3. The determination of the PACS response to a user request in a clinical environment must take into consideration two important factors. First, the additional load put on the system by the request must be accounted for and, second, the influence of such loads on each service provided by the PACS must be considered. Arguably, the most relevant performance factor for the evaluation of a PACS is the time delay before the requested image can be displayed on the viewing screen, as perceived by the user. Several processes deliver images and thus require timing studies under clinical load conditions to evaluate their respective throughput and stability. The distribution of these time delays provides information about the overall system behavior. In addition, it also provides information about the distribution of the load on the system, the efficiency of the transfer process, the optimal organization of the hierarchical storage, and the influence of computer specifications on the display time of images.

4.1 Daily load

4.1.1 Average daily distribution of data volume

Figure 4.1 shows the daily distribution, per hour, of the acquisition of new images. The data show the number of series and total MB produced per hour to illustrate the resource requirements (storage capacity and network bandwidth).

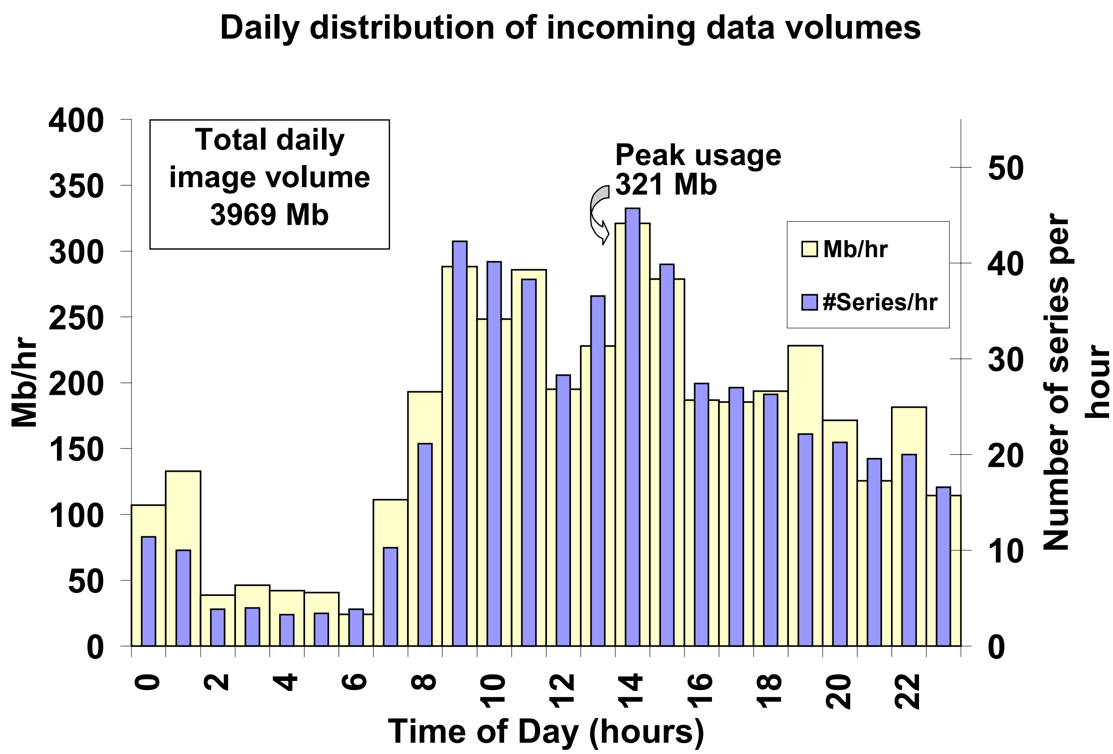

Figure 4.1 Distribution of the daily load of new images averaged over one week.

The acquisition of new images is related to working hours. Most of the acquisitions are performed daily within business hours (08:00 to 16:00). The evening has a lower load and minimum activity is reached from 02:00 to 07:00. The peak hour is between 14:00 and 15:00 with 8.1%[9] (321 MB/hr) of the 3969 MB acquired daily. Averaging the results for each day over a week minimizes the variations from day to day. Saturdays and Sundays are included in these results, which allows calculation with 365 days per year, but reduces the daily load compared to weekdays.

Each new image acquired by the PACS is first saved on the RAID attached to the server. The autorouting process is then initiated by a database update and the image is sent to one or more workstations where it is expected to be reported. The archiving of these images in the autorouting process is done by including the long-term storage server in the list of destinations for all image types. The mean number of destinations for autorouting was observed to be 3.87. Since each image received is sent, on average, to four other destinations, the network traffic due to new images can be evaluated by multiplying the data in Figure 4.1 by 5 to include the initial transmission of the image from the imaging modality to the server.

Apart from autorouting, the prefetching process, creation of JPEG copies for web access, and Master Database updates are initiated when the new image is saved on the Modality Server. These operations involve a lot of network bandwidth, disk access, database operations, and CPU time. In order to avoid an overload of the server during the daytime, compression, which is the most CPU intensive process, is delayed until nighttime.

The other major contributions to the PACS load are the user-initiated image retrievals and transmission. In Figure 4.2, the volume of images requested via the web interface is shown in a daily histogram averaged over one week.

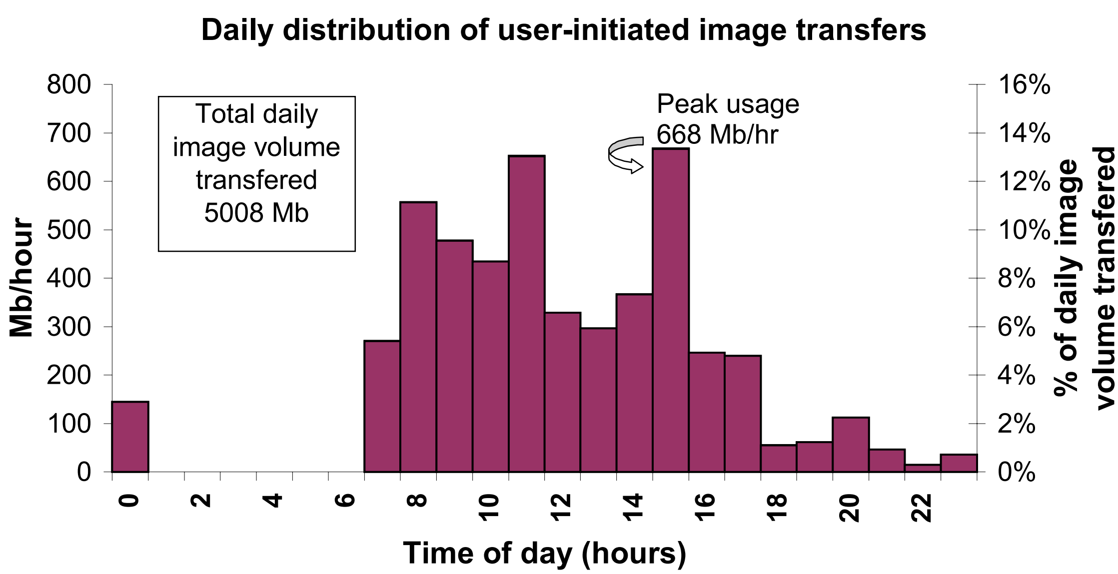

Figure 4.2 Distribution of the daily load of requests issued via the Web interface averaged over one week.

---

9 Similar relative results were documented at the Madigan Army Medical Center[57].

This distribution has a shape similar to the distribution seen in Figure 4.1 (incoming data volumes) but is busier in the daytime and completely unused at night. The peak load occurs between 15:00 and 16:00 with 13.3% (665 MB) of the 5,008 MB moved daily via the web. In comparison to the acquisition of new images, this process is more straightforward since it does not involve other operations beside the sending process.

Image retrieval via the web is only one method of three available to retrieve an image. DICOM-conformant review workstations are also able to retrieve images from the archive. Prefetching of previous studies for comparison purposes is the third retrieval process. These other transmission processes contribute significantly to the volume of images transferred daily, but information about these transfers is not saved in a database. Therefore, it is not possible to record all of these events as it was for the transactions performed via the web. However, it was possible to quantify these data volumes using the log files produced by timing probes. Over one day of operation, the volume of data retrieved via workstations reaches 98% of the volume of data requested via the web. For prefetching, it reaches 87%. The daily distribution is expected to be similar for requests made from workstations, but not for prefetching because the latter occurs primarily in the evening, once per day. Despite the existing difference in daily distributions, the overall daily volume of images transferred can still be estimated by multiplying the volume of image transfers from the web (5,008 MB) by 2.85 (100% via the web plus 98% via workstations, plus 87% via prefetching) to give 14,273 MB per day.

The network traffic in the peak hour of usage can be determined by adding the volumes of images found for each transfer type. The highest volume of new images per hour is 1,605 MB (321 MB x 5) plus 1,904 MB (668 MB x 2.85) of archived images retrieved via either the web, review workstations, or prefetching. The total image volume comes to 3,509 MB for the busiest hour or 0.97 MB per second (7.80 Mbits per second since there are 8 bits per byte). This transfer rate can be compared with the 100 Mbits per second signaling rate of the switched Ethernet network installed at the MGH.

The incoming data and retrieval volumes are shown in Figure 4.3 for each server. These results show how the different servers are used as well as the distribution of the data volumes within the PACS.

Modality Server	New image acquisition	Web query
CR One reader in the emergency room Ten autorouting destinations (on average)	**1,259 MB/Day** 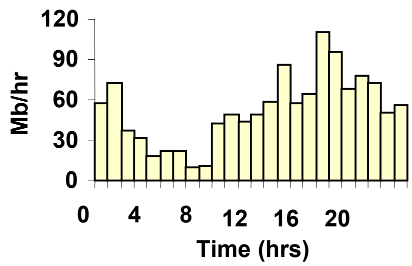	**116 MB/Day** 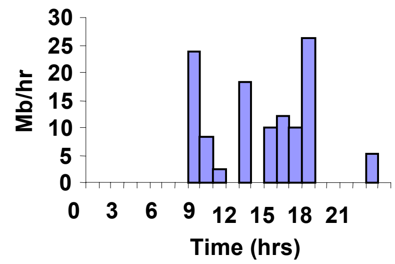
CT 1 Two CT on the radiology floor Four autorouting destinations (on average)	**1,069 MB/Day** 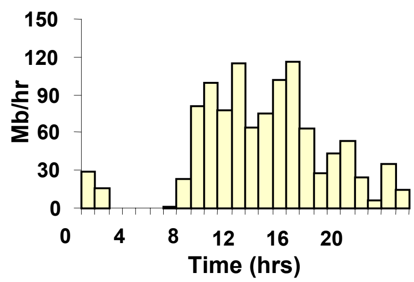	**728 MB/Day** 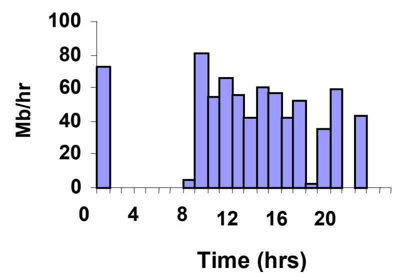
CT 2 One CT scanner in the emergency room Five autorouting destinations (on average)	**874 MB/Day** 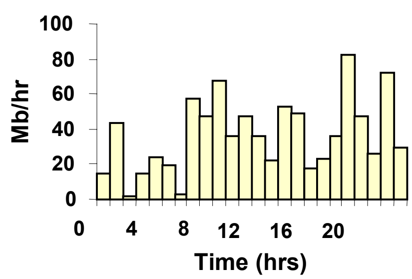	**1,221 MB/Day** 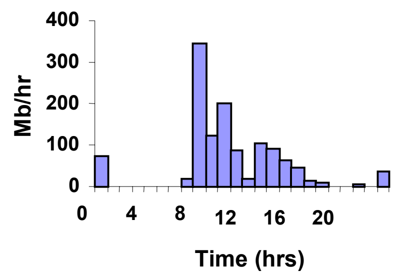
NM Five Gamma cameras Three autorouting destinations (on average)	**9 MB/Day** 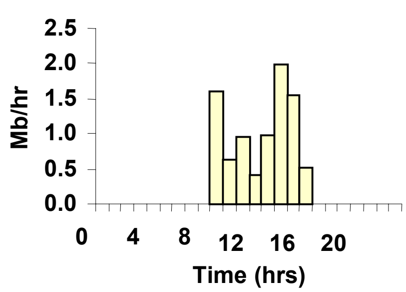	**1 MB/Day** 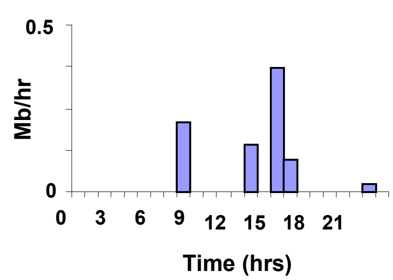

Figure 4.3 Network data volumes due to the acquisition of new images and user retrievals for each Modality Server.

Modality Server	New image acquisition	Web query
MR One MR in operation From 8 to 23hrs Six autorouting destinations (on average)	**394 MB/Day** 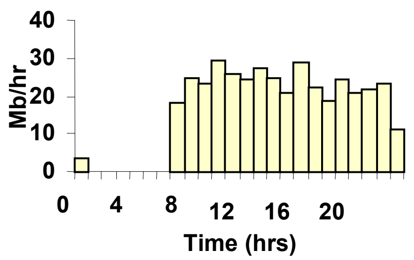	**242 MB/Day** 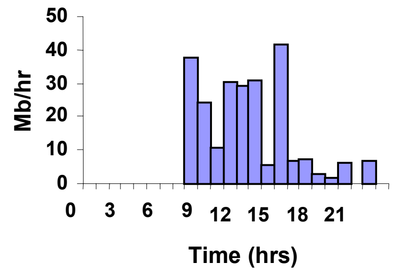
US Nine US Two autorouting destinations (on average)	**366 MB/Day** 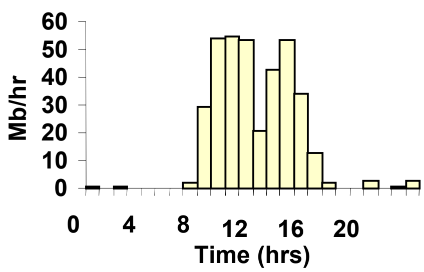	**50 MB/Day** 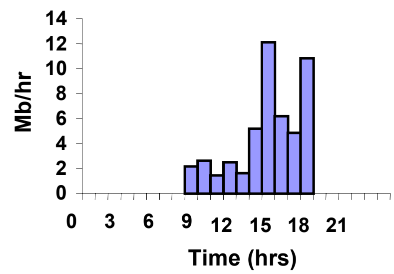

Figure 4.3 Network data volumes due to the acquisition of new images and user retrievals for each Modality Server (continued).

4.1.2 System parameters

In order to evaluate what level of load each server can handle and how it is affected, the usage of the resources of each server must be examined. CPU load, disk space, network traffic, and RAM usage were monitored to see how intensively they were used. The incoming rate of new images and the queue lengths (AutoRoute and MoveReqTable) were also monitored. A graph for each monitored parameter on one of the CT servers is shown in Figure 4.4 as an example. Graphs of those parameters were obtained using MRTG with data sampling every 5 minutes and were observed over 33 hours.

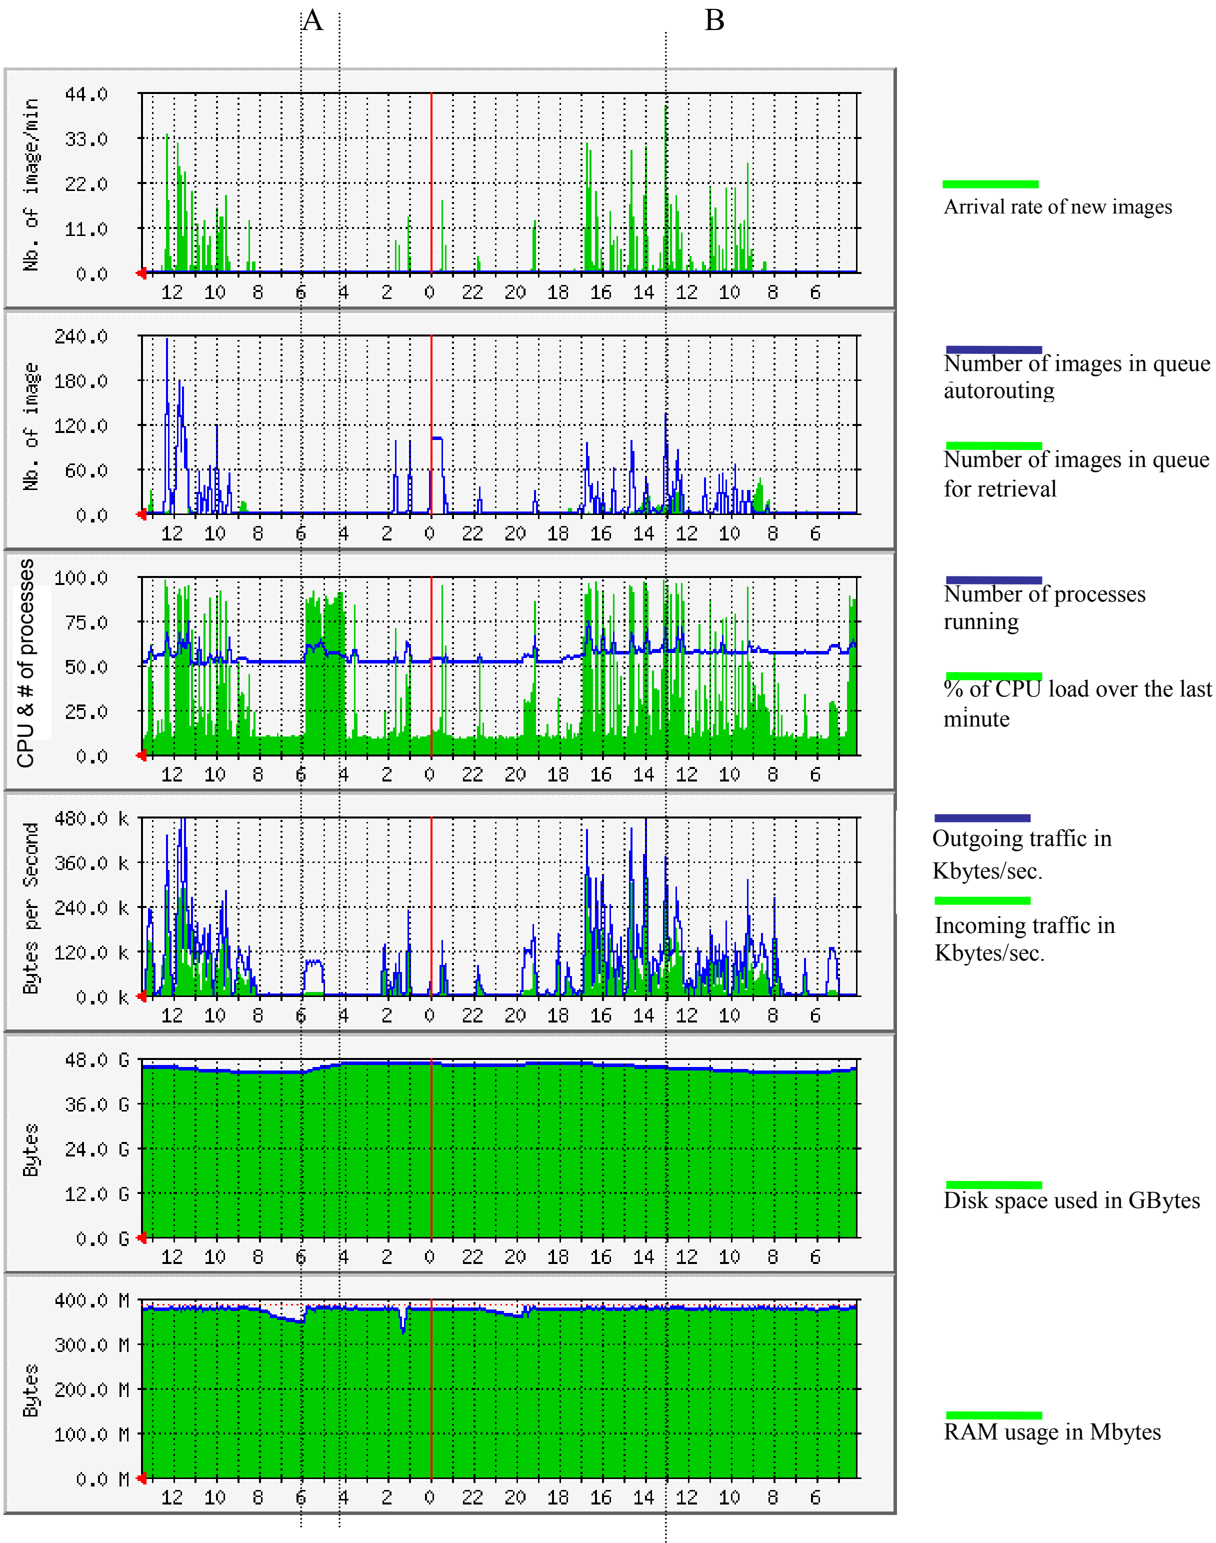

Figure 4.4 Monitoring of a server receiving only CT images over 33 hours. "A" Shows the compression done overnight for all the images acquired during the previous day. "B" is an example of a burst of incoming images and its effect on the various system parameters.

The monitoring of low-level system parameters allows easy comparison of each server. These comparisons are summarized in Tables 4.1 and 4.2 for the network traffic and CPU usage, respectively.

Table 4.1 Ranking of the Modality Servers as a function of their network traffic.

Network traffic (averaged over one week)		
Server	Outgoing traffic	Incoming traffic
CR server	223.5 kB/s	58.9 kB/s
CT server	51.9 kB/s	34.5 kB/s
MR server	12.7 kB/s	9.82 kB/s
US server	7.23 kB/s	7.34 kB/s
NM server	0.5 kB/s	1.5 kB/s

Table 4.2 Ranking of the Modality Servers as a function of their CPU usage.

CPU usage (averaged over one week)	
Server	CPU (%)
CT server	35%
MR server	16%
CR server	12%
US server	9%
NM server	8%

4.1.3 Discussion

The average daily distributions of the data volume, shown in Figures 4.1, 4.2 and 4.3, provides general information on the usage of the PACS. Relations of the data volume with daily workload are easier to observe on a server-by-server basis as is illustrated in Figure 4.3. Each imaging modality has different regular working hours and various levels of usage. From these graphs, the load introduced by new images is more important than retrieval when taking into account the operations initiated upon reception of new images (autorouting, archiving, JPEG creation, database updates, and prefetching of prior studies).

The graphs from Figure 4.3 illustrate the distribution of the load among all Modality Servers. This distribution of the data volume among servers is not well balanced. The NM server is practically unused when compared to the other servers.

Average daily distribution graphs are useful to determine the distribution of the load among servers, the minimal bandwidth required for each server, long-term storage needs, and the ideal time for system management or servicing. However, network data volumes do not give an indication of what is considered acceptable for the computing capacity of a given server. Further information is needed to know how the resources of servers are used and how they influence the transmission of images.

Resource usage of the servers, illustrated through the web monitoring tool shown in Figure 4.4 demonstrates how the monitoring tools may be used to relate high-level PACS operations with the lower level system parameters. The graph in Figure 4.4 that illustrates the arrival of new images over a day is similar to the daily distribution shown in Figure 4.1. Hourly sampling and averaging over one week smoothes out results and shows an almost constant acquisition rate over daytime hours. The monitoring graph, updated every five minutes, instead reveals discreet events occurring with some idle periods or some peaks above the normal level. The average load found when examining database queries is more useful for the calculation of long-term storage needs. But the monitoring method is certainly more appropriate to determine peak usage and its effect. It is important to realize that a 5-minute sampling rate still needs averaging. Continuous monitoring or higher sampling rates would probably reveal even more details.

The arrival of new images on the server triggers many processes. Image arrivals significantly modify the network incoming and outgoing traffic and CPU usage. Minor changes are observed in the number of processes running and the number of images in autorouting and retrieval queues. However, no significant changes are observed on the RAM usage or the hard disk space.

After an image is received on the Modality Server, it is sent to several review workstations as well as the archive. The network traffic (incoming and outgoing) have the same shape during the day, but the outbound traffic is expected to be higher since images are sent to more than one destination through autorouting. Part of the CPU activity is related to the reception and sending processes throughout the network, but most of the database queries and updates are done upon the reception of the image. The local and remote (master) databases are updated with information for each new image. In order to get the appropriate viewing parameters (window width and level), the image header values are read. The prefetching program looks into the database for relevant studies performed previously on the same patient. These operations also require processing time as can be seen on the CPU usage graph. Many processes fork from the image management program to perform these tasks. The number of processes running was monitored to see if it could reach unacceptable levels under different conditions, but this was never observed. The number of forked processes needed to manage the images remains small compared to the total number of processes that run constantly.

In normal situations, images are sent rapidly to their destination and do not stay long in the queues. Since the monitoring tools only sample the queue length every five minutes, most images can be inserted and removed from the queues between the measurements. Queues are most visible to the monitoring software when a workstation does not receive the images properly or when large numbers of images accumulate rapidly.

No influence is shown on the hard disk space used even when a large volume of images are received. The volume of the new images, even in large numbers, is not significant when compared to the total space available on the hard disk, which is typically tens of gigabytes. The only variation observed was the reduction of the space used while the images of the previous day were compressed overnight. Monitoring the hard disk space used in this manner has little practical value. In the context of a PACS, it would be more interesting to show the number of days the data are available online and show an estimate of the time needed to fill the space remaining.

When ranking the servers as a function of their network traffic (Table 4.1) and their CPU usage (Table 4.2), the order is surprisingly different. The CR server is the highest network user but its processor is one of the least used. CR images managed on the CR server have file sizes up to 32 MB, which is much larger than the usual file size handled with other modalities, ranging from 0.066 to 0.512 MB. Even with network traffic that is four times larger than any other server, the CR server handles fewer images than the CT server, retrieval server and MR server. This would imply that the number of images manipulated is the main influence the CPU usage. High incoming and outgoing network traffic is possible without a large effect on CPU. The large number of database updates and queries related to the arrival and transfer of each image may require more computation than the transmission itself.

Even though it is possible to monitor the low-level system parameters, it is still difficult to determine the maximum acceptable load on a server before image transfers and the other services provided to users are affected. The presence of peaks in the daytime, resulting in nearly 100% CPU usage, may affect normal operation and reduce overall transfer time. The width of such peaks is important. Reaching the limit of CPU capacity on occasion, when simultaneous events occur, can be tolerated. However, in critical clinical operations, it is not tolerable to have a build-up of a large queue of operations. Large queues result mainly from a busy PACS system that is not capable of processing the tasks required more quickly than they accumulate. In order to get a more objective method to identify unacceptable loads, image transmission time should be monitored and correlated with the low-level system parameters. Unstable or increasing transfer times should be correlated to server usage to help determine more objectively when equipment upgrades are needed.

On the other hand, monitoring of the network shows a relatively low overall usage. For a 5-minute average, the highest level of transfer registered was 9,600 kilobits per second, which represents 9.6% of the network signaling rate. Care must be taken when comparing the throughput and the signaling rate. As reported by Prior, Frey, and Horii, the maximum throughput is considerably less than the signaling rate[48, 68, 73]. Additional headers from the protocol and delays between each communication layer of the TCP/IP protocol could reduce the throughput to as low as a tenth of the signaling rate.

4.2 Efficiency of Hierarchical Storage (Online vs offline)

The storage of data hierarchically is a strategy used to optimize retrieval speed while balancing hardware cost. The balance between different storage technologies in the system is a compromise between speed of accessibility, capacity, and price.

4.2.1 Results

Generally speaking, recent images have more clinical interest. Physicians request more images acquired in the last few days than older images. Figure 4.5 illustrates this. It was generated by plotting the number of days between the initial acquisition of an image series and when it was retrieved versus the cumulative number of series retrieved. In Figure 4.5, over 42% of the images were retrieved in the first 20 days following acquisition and only 15% of retrieval requests were for images older than a year. A good hierarchical storage design should make use of the information contained in this graph.

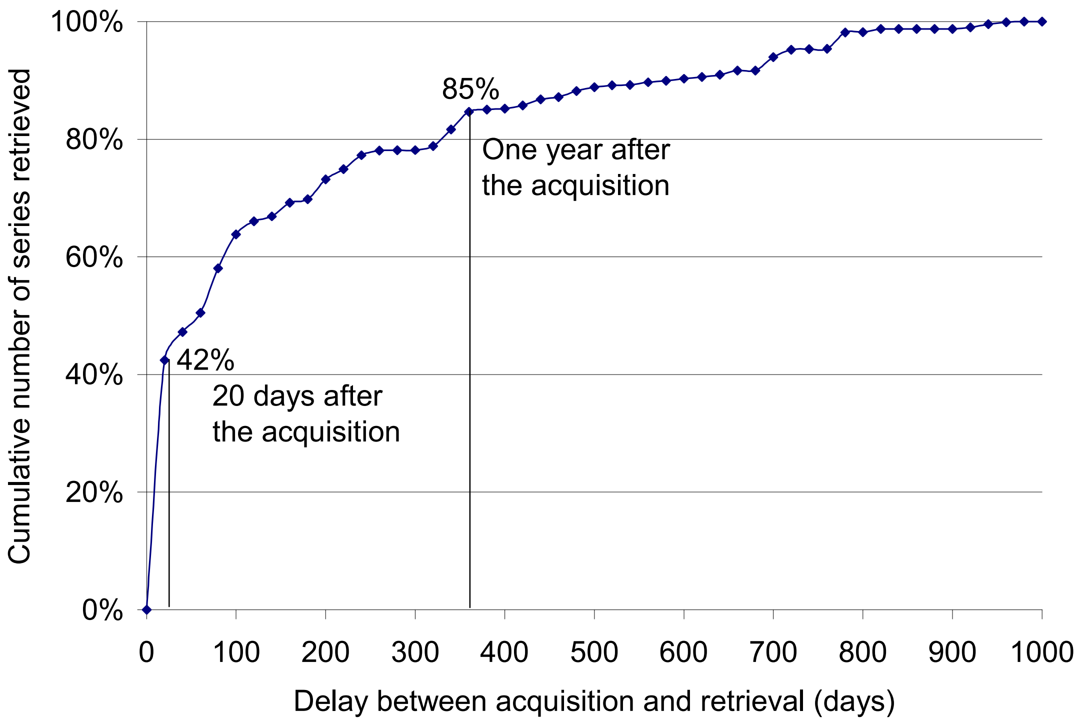

Figure 4.5 Delay between acquisition and retrieval plotted versus the cumulative number of series retrieved per date.

Based on the volume of incoming data found for each server and on the size of its respective RAID, the capacity of the RAID was estimated, measured in terms of the number of days before the storage of the RAID is exceeded. The results are presented for each Modality Server in Table 4.3. An average compression ratio of 54% was assumed.

Table 4.3 Availability of data online for each modality.

Modality Server	Capacity (MB)	Acquisition load (MB/day)	Online time (days)
US	68,000	366	344
CT 1	50,000	1,069	87
CT 2	50,000	874	105
MR	34,000	394	159
CR	89,000	1,259	130
NM	25,000	9	14 years

4.2.2 Discussion

Having a distributed system represents a great advantage in terms of load reduction on each server, but the load needs to be balanced across the whole system to obtain the greatest benefit. Table 4.3 clearly shows imbalance in the storage capacity of each modality. This imbalance was noticed with the distribution of incoming data volumes. The imbalance here in terms of days of availability online is also influenced by the size of the RAID attached to each Modality Server. Nuclear medicine and ultrasound servers clearly have more storage capacity than needed. The excess storage could be better used for CT images since they have the lowest time of availability online. If all the storage capacity were put together, a total of 316,000 MB of space could store up to 147 days of data for all modalities. With the distribution of the requests plotted in Figure 4.4, the 147 days of availability would satisfy 70% of the requests. Doubling this total online storage capacity would only raise the ratio of requests satisfied to 80%.

The consolidation of the storage in one location allows equal sharing of space for all modalities, but eliminates the distribution of the load on several servers and their respective network connections. The best of both worlds could be achieved by adding an intermediate step in the hierarchical structure. A very large and centralized storage could keep the images from each Modality Server for an additional period of time. Since all the images would be managed with the same first-in-first-out rule on this intermediate central storage, there would be no imbalance. The storage capacity on the Modality Server could be reduced to get 20 days of availability online and still satisfy 42% of the retrieval requests. Further improvements could be made by shortening the online availability for particular types of exams that are unnecessary after a given time or are known to be rarely requested.

In the case of long-term archiving, the number of tapes needed to store the images produced daily is constantly increasing. The actual policy for archiving of film radiographs allows the hospital to destroy the film five years after its acquisition. In fact, the MGH only destroys the film folder of a patient when its most recent film is five years old. Policies for the management of old digital images have not yet been determined at the MGH. Depending on the criteria of rejection, a large number of images could be removed when they become too old. The data do not necessarily need to be destroyed. Special tapes could be written with complete patient image folders and be kept off-line for special requests. After several years, supporting a legacy of old storage media could become difficult. Magnetic tapes have a limited lifetime of 30 years[10] when stored under ideal atmospheric conditions and properly maintained. New standards of tapes and readers are not always retroactively compatible[68].

4.3 Time spent in the sending process

Each step of the image transmission process was examined to assess its efficiency. The evaluation was performed using the residency timing method described in section 3.3. Because we were unable to insert timing probes in the software peripheral to the PACS (e.g., on the modalities themselves, or on the review workstations), we examined the rate at which new images are received by the servers from the modalities. The time spent before autorouting to the review workstations was also analyzed. Of most importance in terms of transfer time is the time needed to transfer images requested manually by users. Whether the images requested are available online or not, both processes have been studied to show their relative efficiency.

4.3.1 Receiving new images

The time used for determining the arrival rate of images on each Modality Server is the time between the local database insertion of the first and last images of the same series. These transfer times are shown in Figure 4.6. This arrival rate reflects the throughput at the highest level. It does not represent the transmission rate. The arrival rate is characteristic of each imaging modality and reveals the way it is integrated into the PACS. Even if there is no possibility to modify this arrival rate, it is important to understand the flow from the beginning of the path since it influences the overall throughput.

[10] ADIC specify a minimum Tape Shelf Life of 30 years if the (DLT) tape are stored at 20ºC and 40% of relative humidity.

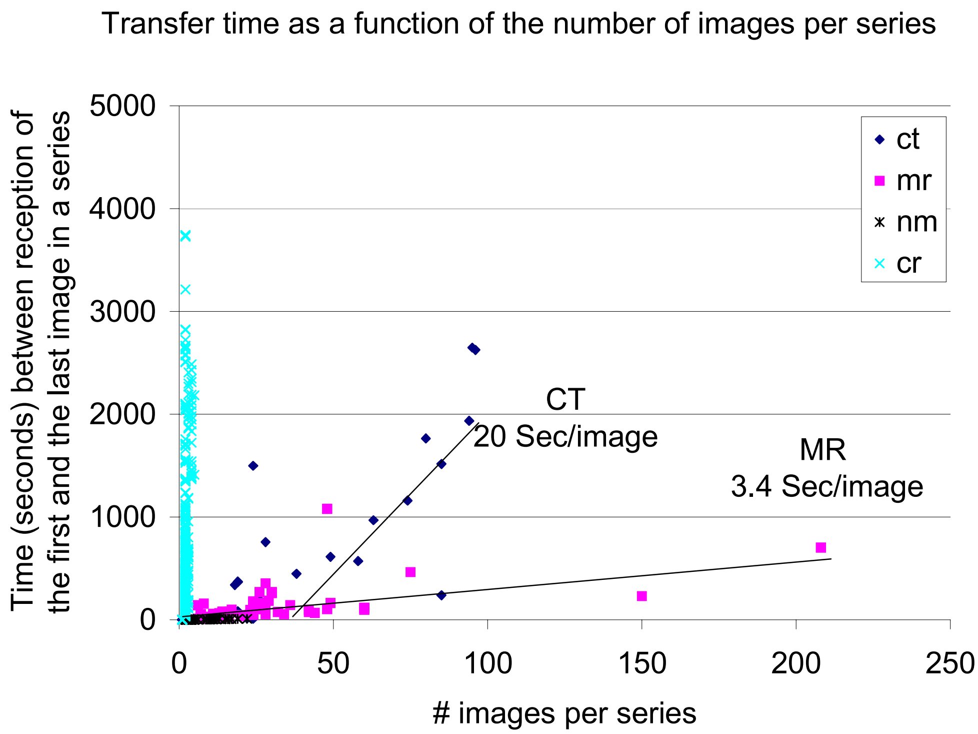

Figure 4.6 Time spent for the reception of images per series from the imaging modality to its server.

The calculation of the minimal slope gives the transfer rate for each modality and was found to be 0.4 seconds per image for nuclear medicine, 3.4 seconds for MR, and 20 seconds per image for the oldest of the three CT scanners. For the CR, the slope is impossible to calculate since the series are small (often only a single image) and there is a large variation in the time to complete the series transfer. These rates are influenced by the image size and how and when the software on the modalities chooses to transmit the images. Most MR and CT scans have to undergo reconstruction to obtain the images and send them as they become available. In the case of CR, apart from the large file size, the reading time (i.e., the time taken to scan the CR plate to generate the digital image) is significant. It will become important to consider this modality's limitation in the analysis of further steps in the transfer process.

4.3.2 Time needed for the autorouting process

The timing information on the autorouting process was obtained using the residency timing method. As is shown below, the time spent for one of the CT servers is representative of most other Modality Servers and of the special case encountered with CR.

4.3.2.1 Timing of image autorouting on a CT server

The first example is the autorouting of all the images received by the CT server supporting the Emergency Department through a normal day of operation. The total times taken from the reception of the image at the Modality Server to the completion of its transmission to the review workstations are plotted in Figure 4.7.

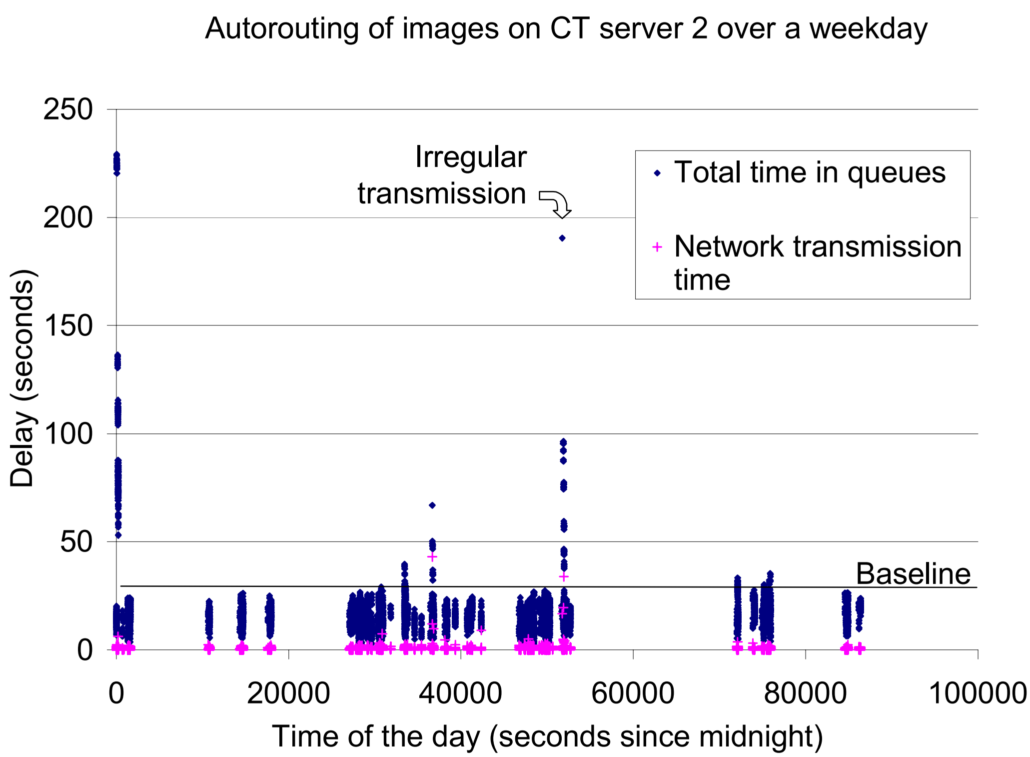

Figure 4.7 Time spent in autorouting process over a day on the server that supports the CT server in the Emergency Department. Most transmissions are completed less than 25 seconds after their insertion in the transmission queue with the exception of a few irregular transmissions.

The observation of the height of each peak shows that in most situations, the longest time to transfer an image is below 25 seconds. In some situations it takes longer to transfer the images. In the example shown in Figure 4.7, we can identify four such (unusual) peaks. For the most part, we can consider the autorouting process as stable because of the reproducibility of each transfer and the small number of abnormal delays. A normal process is illustrated in Figure 4.8. In Figure 4.9 it is enlarged to better understand the efficiency of each step. An abnormal transfer is illustrated in Figures 4.10 and 4.11.

Figures 4.8 and 4.9 clearly show how slowly the images are fed to the Modality Server by the modality itself. Then, even if the transfer is fast from the Modality Server to the workstation, it does not change the perceived transfer speed. In this case, the transfer rates to the workstations are sufficient to keep the pace of the slow receiving rate. The autorouting process has only added an additional delay of 23.4 seconds to the overall transfer of the 72 images that took 228 seconds to be completed. But in the case of a faster imaging modality (and faster network connection to the modality), the transfer rate could go faster to improve the interactivity of the PACS. With a closer look at the sequence of events for a few images, it is possible to evaluate the efficiency of the sending process.

The four timing probes used in Figure 4.9 refer to the most important events in the image transfer process. From Figure 4.9, it is clear that the network transfer does not contribute significantly to the overall transfer time. The average transfer time for these 72 images (network only) is 0.5 seconds and 95% of the transfers are done within 1.2 seconds. Most of the time is spent in the queue and brings the overall image transfer from the time of reception up to 14.2 seconds on average with 95% below 19.9 seconds. This shows inefficiency in the programs involved (do_auto_route and AutoRoute). Even with all the processing and database access needed, it cannot justify such delays. The insertion of the first image is done almost instantaneously, but for the other images there is a long delay (up to 10 seconds) after which they are all inserted in the transmission queue. The loop algorithm used is responsible for this behavior and will not look at new incoming images while the computer is busy sending the first image. For small images this result in a stepwise instead of a continuous flow of images. In the case of autorouting, this level of performance is acceptable since it has no impact on the transfer rate because it is limited by the imaging modality transfer rate. But in

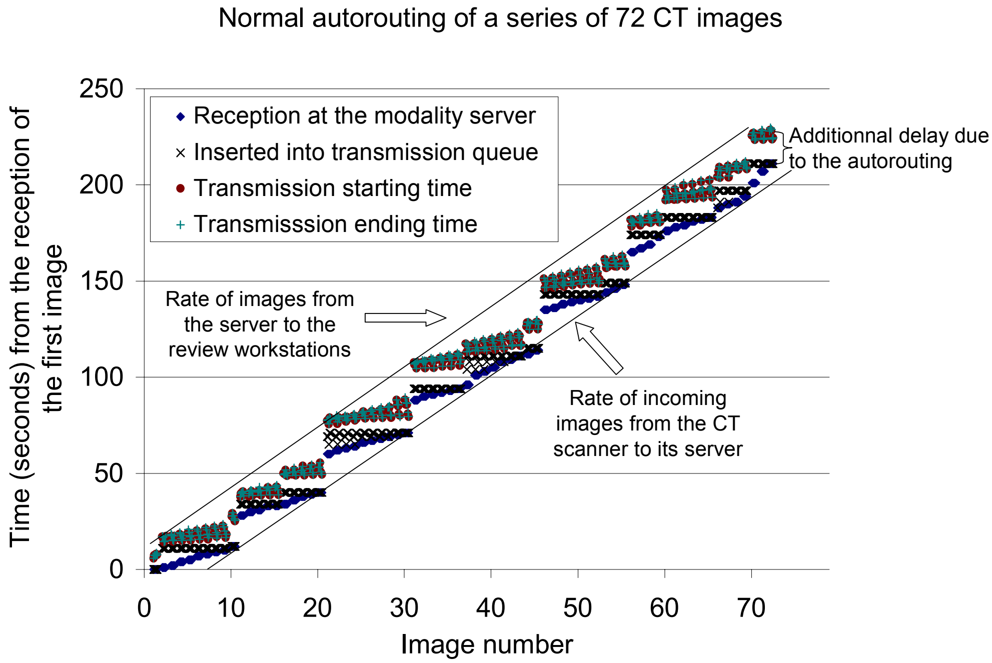

Figure 4.8 Normal autorouting transfer of a series of 72 CT images series. It shows a similar transfer rate from the imaging modality to the workstation. A delay of less than 25 seconds was added to the maximum transfer rate achievable with this modality.

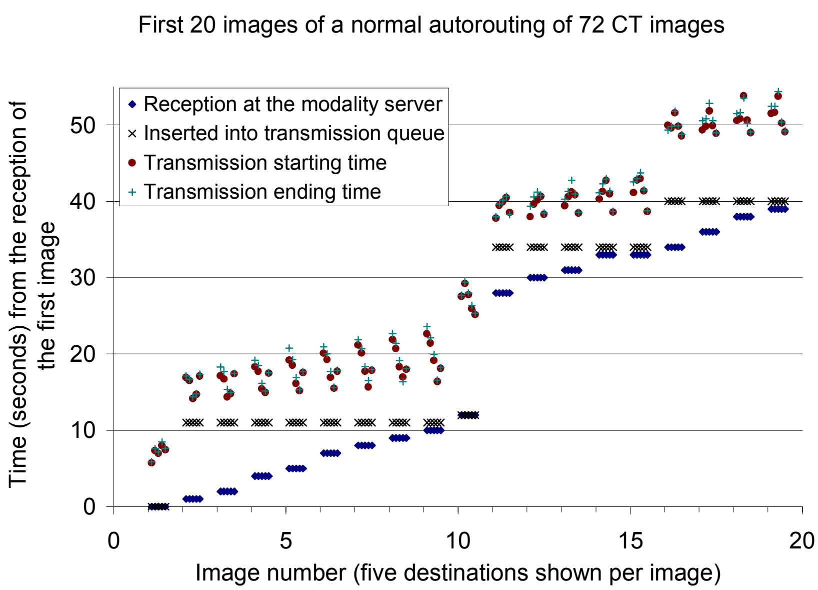

Figure 4.9 Internal processes timing for the first 20 images of a normal autorouting of CT images (enlarged from Figure 4.8). Once each image is received at the Modality Server the insertion in queue and the transmission is then shown for the five destinations selected by the autorouting rules.

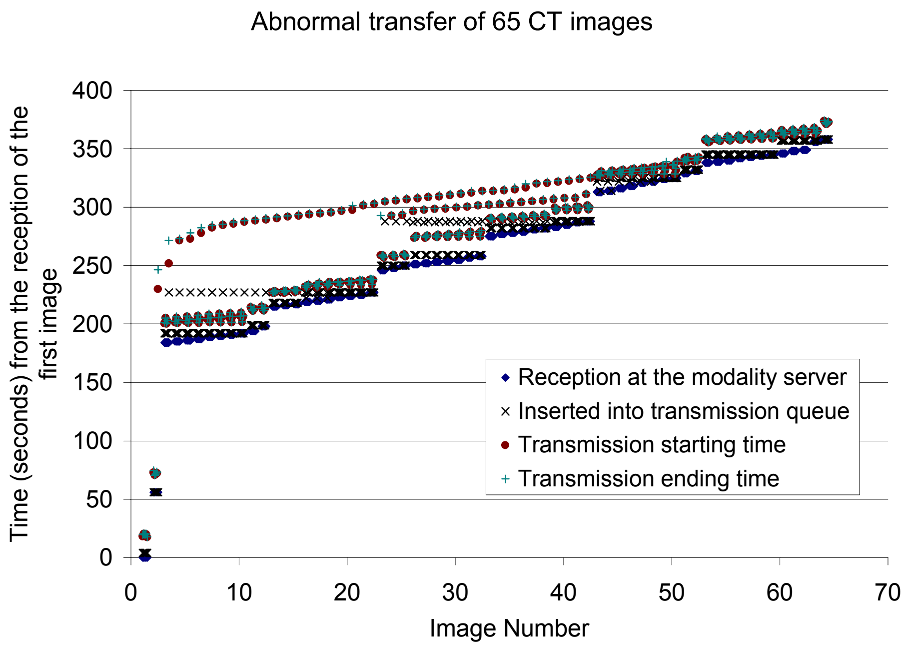

Figure 4.10 Autorouting timing of CT images in an abnormal situation.

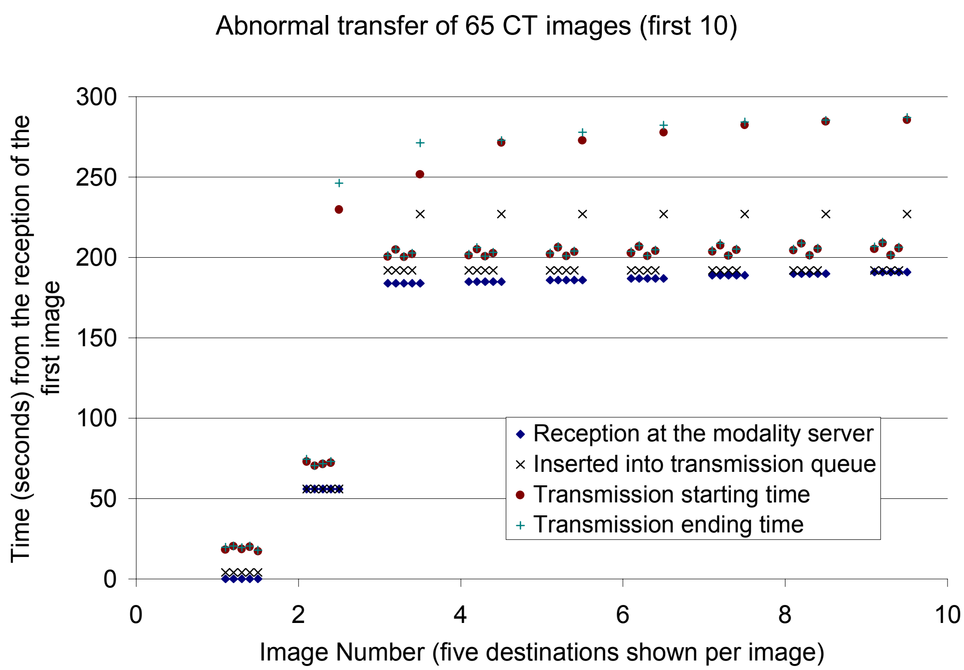

Figure 4.11 Abnormal transmission of CT images (enlarged from Figure 4.10). Once each image is received at the Modality Server the insertion in queue and the transmission is then shown for the five destinations selected by the autorouting rules.

the cases where the images are already on the server, there will be no such limitation and better performance is expected, especially when a user directly requests the transfer.

Applying the same analysis to transfers with abnormal delays, it is possible to look at the source of this delay. The largest peak in Figure 4.7 occurred around 52,000 seconds after midnight. Figures 4.10 and 4.11 show this peak in more detail.

In comparison with the normal transfer process, some delay occurs between the time the images are first inserted and completely sent. This delay appears at the second or third image and is slowly reduced as the images are transferred. Visibly, a problem happens with the transfer of the first image. After this, new images are sent faster than they come into the system. Again, at the end of the process, the transfer was mainly delayed by the time taken by the imaging modality to send its images to the server, and only 15 seconds are added to the overall procedure due to the processing on the server.

Figure 4.11 allows one to see that for one destination (the fifth), the transfer of the second image was delayed for 174 seconds before its transfer was even started. Once the sending was initiated, it took 16.7 seconds to complete it. This is an abnormal situation that delayed the transfer of the whole series. The third image was also has been abnormally delayed. After that, the process continued normally. But since the images are sent one after the other to a given destination, the other images in the series had to wait in queue longer than usual to be sent to the fifth destination. It is hard to find what really caused the delay in sending the second and third images to the fifth destination. Because all the other destinations received the images normally, the problem is either on the receiving workstation or the link that connects it to the network. The punctual and brief aspect of this abnormality indicates that we are dealing with temporary trouble and not a throughput limitation on any part of the hardware.

4.3.2.2 Timing of images autorouting on a CR server

In the case of CR, one typically sees small series, between one and three images with large file sizes from 4 MB to 32 MB. Using the same method as for CT, we show in Figure 4.12 the transfer time of each image over a working day.

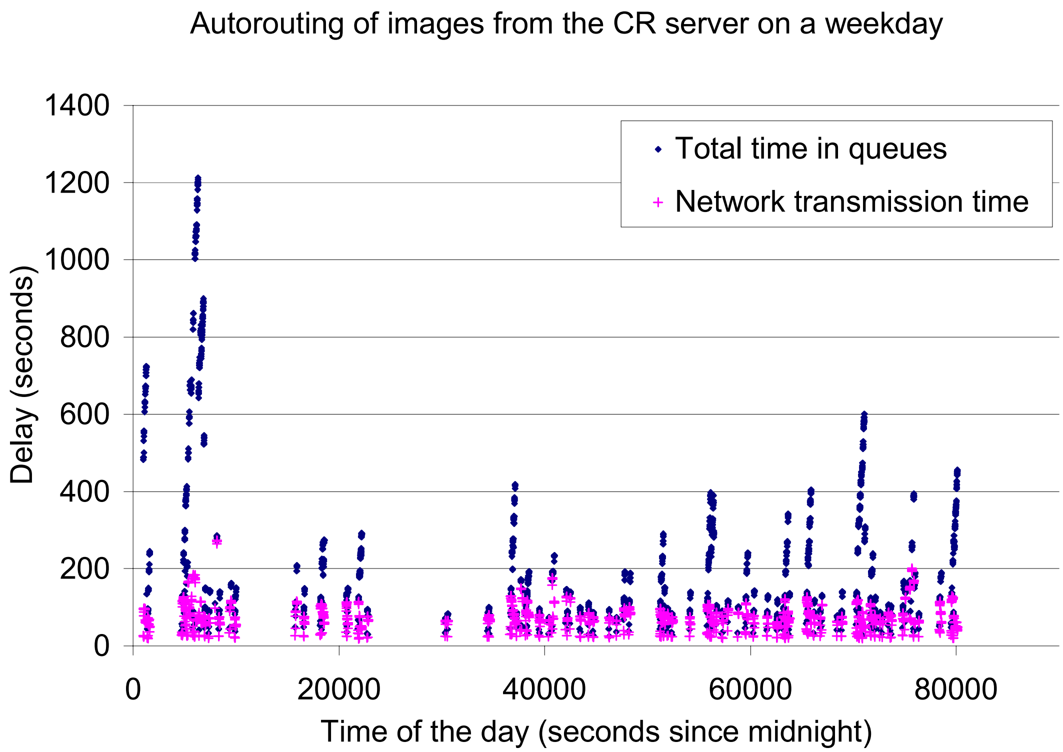

Figure 4.12 Time spent in autorouting process over a day on the CR Modality Server.

As with CT, we can establish a normal measure for the transfer and determine the stability of the process on this server by looking at the number exceeding the limit. In this case, it is harder to determine what is normal since there is no clear baseline for the transfer times. Considering that the median for the entire autorouting process is 137 seconds, we will assume that the normal limit is put slightly over this; say, 200 seconds. A large number of transfers lie above this arbitrary limit. The examination of transfer sequences below this limit, in Figure 4.13, and above it, in Figure 4.14, is shown on the following page.

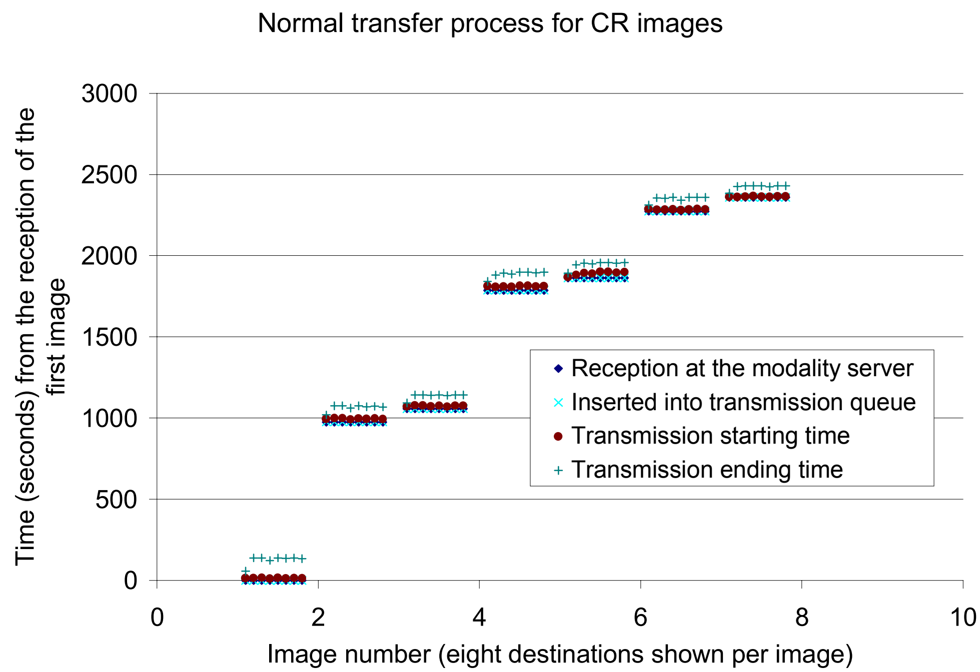

Figure 4.13 Normal transfer processes for CR images.

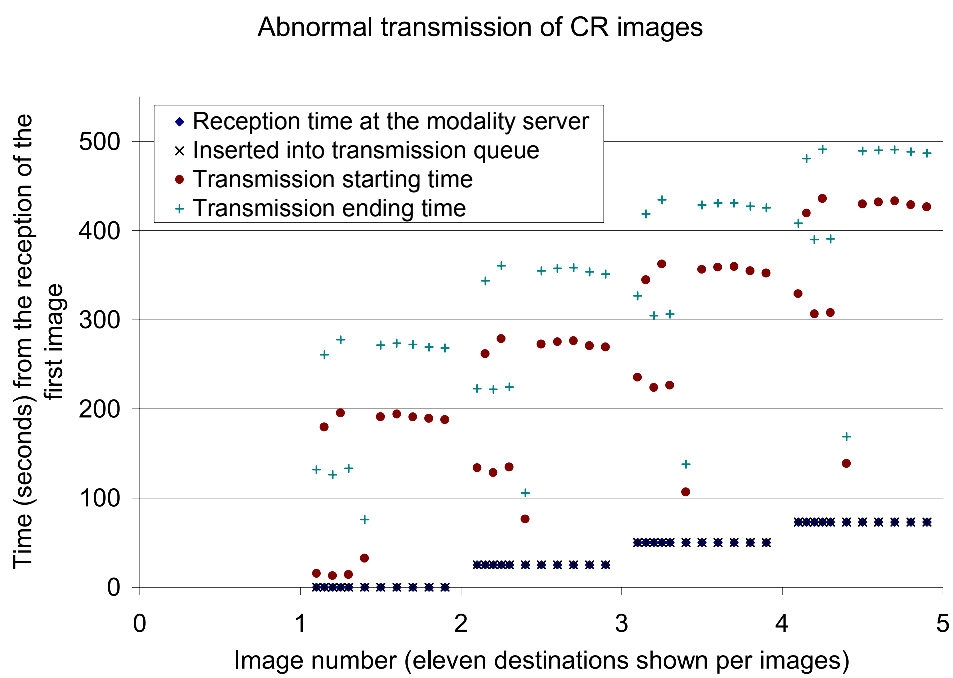

Figure 4.14 Queuing of CR images for transfer to several destinations.

In Figure 4.13, where images are transmitted within 150 seconds, each image is processed individually because the next one is always received once the previous image is completely sent. There is no additional delay due to queuing of the new images while another one is transferred.

In the other example (Figure 4.14), four images were received and inserted in the transmission queue within 100 seconds. The fourth image took 413 seconds before it was transferred to all of its destinations. The transmission took the same time as the previous images, but did not begin until the complete transfer of the previous images. These delays are exceedingly long and can become longer as shown in the graph. The simultaneous sending of images to a large number of destinations implies sharing of the available bandwidth of the server. In these two examples (Figure 4.13 and 4.14) one of the destinations always receives the image faster then all the rest. This increase efficiency in the transmission depends on the workstation. This workstation was identified as one of the archive servers. This destination runs Linux as the operating system and is more efficient for multi-tasking than the other regular review workstations.

4.3.3 Time needed for manual retrievals

When a user makes a request for image retrieval on the PACS, he expects to get the images quickly. Autorouting transfer times were not as crucial since new images are automatically pushed to workstations prior to viewing. Depending on the age of a study (or the last time it was reviewed), it could be stored either online, on a RAID attached to one of the servers, or offline in a tape library.

4.3.3.1 Retrieval of images available online

The timing probes inserted in the program AutoMove2 provide the timing information about the DICOM association, the image transmission and its decompression (if needed).

The DICOM connection time was found to take an average of 0.36 seconds. But this connection time has a Poisson-like distribution with a median of 0.20 seconds and a 95% percentile of up to 1.0 second.

Several connection times took as long as 40 seconds. This long connection time definitely occurs in irregular situations, and the 40-second wait probably comes from a time-out delay programmed into the function that handles the negotiation of the connection on the receiving side. These irregularities occurred only on the CR server when connecting with three specific destinations. These three destinations frequently had problems first getting connected. Once connected, the transmission of images was often also abnormally slow, taking up to 200 seconds instead of the usual 15 seconds.

The examination of the log files has shown that AutoMove2 opens and closes a DICOM connection for each series sent. This practice is not ideal, especially for CR where a series usually consists of one image. There may be something to be gained if a single connection was used for all images that were requested at the same time and were going to the same destination. Unfortunately, it is a violation of DICOM communication rules if images from more than one modality are transmitted over the same association. This is not a problem when transferring only one series over an association since the modality is always the same for every image in a series.

Apart from the decompression time, the images are sent one after the other without any other loss of time. The time between the end of the transmission of an image and the start of the decompression of the next was observed to be a few thousandths of a second.

In Table 4.4, the transmission and the decompression times are presented for each modality.

Table 4.4 Transmission and decompression time per image.

Image Type	Transmission time (seconds)		Decompression Time (seconds)	
	average	Standard deviation	average	Standard deviation
CR	10.79	16.44	3.16	2.05
CT	1.29	4.16	0.36	0.33
US	2.42	12.25	0.81	0.56
MR	0.58	1.87	0.33	0.38
NM	0.33	1.04	0.23	0.24

In the case where AutoMove2 handles retrieval requests from the automatic prefetching process, the requested images are typically sent to many destinations. The connection to each destination is obtained almost simultaneously and is managed by separate forked processes. Not only is the same image decompressed several times, but this repeated decompression tends to be slower as the number of destinations increases since each decompression process tries concurrently to access the same file. This practice is certainly inefficient, but has little impact on the user since only automatic prefetching can request image transmission to multiple destinations. A single decompression of the series of images, independent of the transmission, would resolve this additional problem and represent a better usage of the CPU resource.

It is important to remember that all the images to be sent to the same destination are put in a queue and are sent one after the other. The number of images requested at the same time will mainly determine the transfer time.

4.3.3.2 Retrieval of archived images

When an image is not available on a RAID attached to one of the Modality Servers, the copy stored on the long-term archive is retrieved. Several steps are needed before the series can be read from tape. These steps where timed and the results are shown in Table 4.5 for the two tape libraries employed in the MGH PACS.

Table 4.5 Summary of the average time (in seconds) needed for various tape library operations.

<table>
<tr><th>Operation</th><th>Sony PetaSite</th><th>ADIC Scalar 458</th></tr>
<tr><td>Robotic move
(tape removal + insertion)</td><td>116[11]</td><td>93[11]</td></tr>
<tr><td>Tape loading</td><td rowspan="3">22.3</td><td>33</td></tr>
<tr><td>Tape positioning</td><td>48</td></tr>
<tr><td>Reading (one series)</td><td>2.8</td></tr>
</table>

[11] This figure includes the time to rewind the tape before it is unloaded.

Mechanical positioning operations of the tape take up the most time when retrieving data. First, a robotic arm brings the tape from the storage slot to the tape drive. Usually a tape is already loaded in the drive and needs to be removed and put back in its storage slot. Robotic motion speed does not vary from time to time. The ADIC tape library needs 26 seconds to remove a tape and 19 seconds to insert a new one. Four different (distinct) robotic motion times were observed for the PetaSite : 37, 58, 73, 96 seconds. Each value is associated with the location of the tape in the basic or drive console. The PetaSite requires more time to move the tape within its larger cabinet, to handle transitions from one module to another, and to manage an additional degree of freedom in its robotic arm. The addition of modules containing more tapes to increase the capacity of this kind of jukebox would certainly increase the time needed to move a tape. On average the ADIC and PetaSite take 45 seconds and 116[11] seconds, respectively, to unload the previous tape and load the selected one.

Once the tape is in the drive, the ADIC takes an average of 33 seconds to initialize and put tension on the DLT before the file can be accessed. On the other hand, the DTF tapes used in the Sony PetaSite have flash memory chips in each tape that provide a tape log and a search map. Once the DTF tape is loaded it is ready to be set to the right file position. Usually many series from the same study are requested from the same tape. When many series wait in queue to be retrieved, the retrieval program looks for other retrievals on the tape currently loaded before it changes to a new tape. Before the series can be read from the tape, the tape must be positioned at the beginning of the file. The ADIC tape library uses DLT drives. For these drives, the mean time observed to position the tape is 48 seconds. After that, the series can be read in an average time of 2.8 seconds. The drive in the PetaSite cannot provide separate information regarding positioning and reading of the tape since the drive caches the request to forward the tape upon reception. Therefore, the forwarding time and reading time are combined. The average sum of the forwarding and reading times on the PetaSite is 22.3 seconds.

Typical usage patterns for both tape libraries are shown in Figures 4.16 and 4.17. Both show similar behavior under clinical load. Normal requests are received throughout the daytime. In Figure 4.16, there are a total of 41 small sharp peaks in the curves that show requests of many

---

series (an average of eight series per peak) at the same time, probably done by a single user. These small requests are usually completely retrieved within 1,000 seconds (17 minutes). When the retrieval delays become unacceptable to the users, the addition of a new tape drive plus a dedicated server is the only strategy that will improve the retrieval performance of the system.

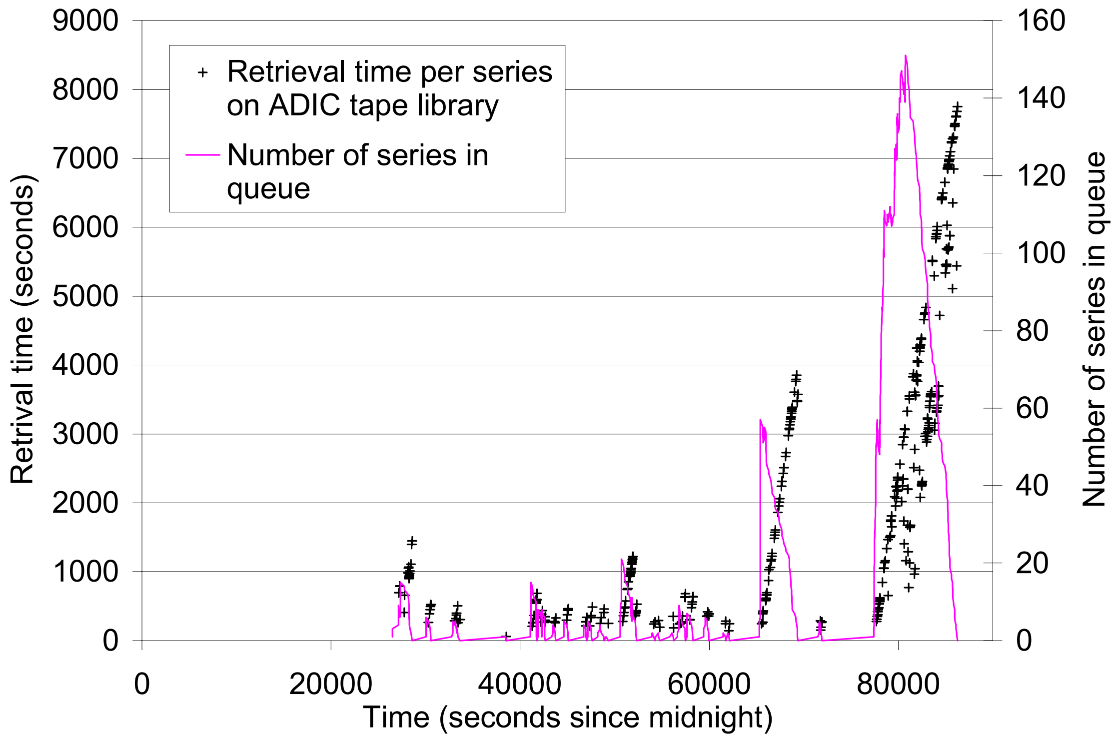

Figure 4.15 Typical usage of the ADIC tape library over one day.

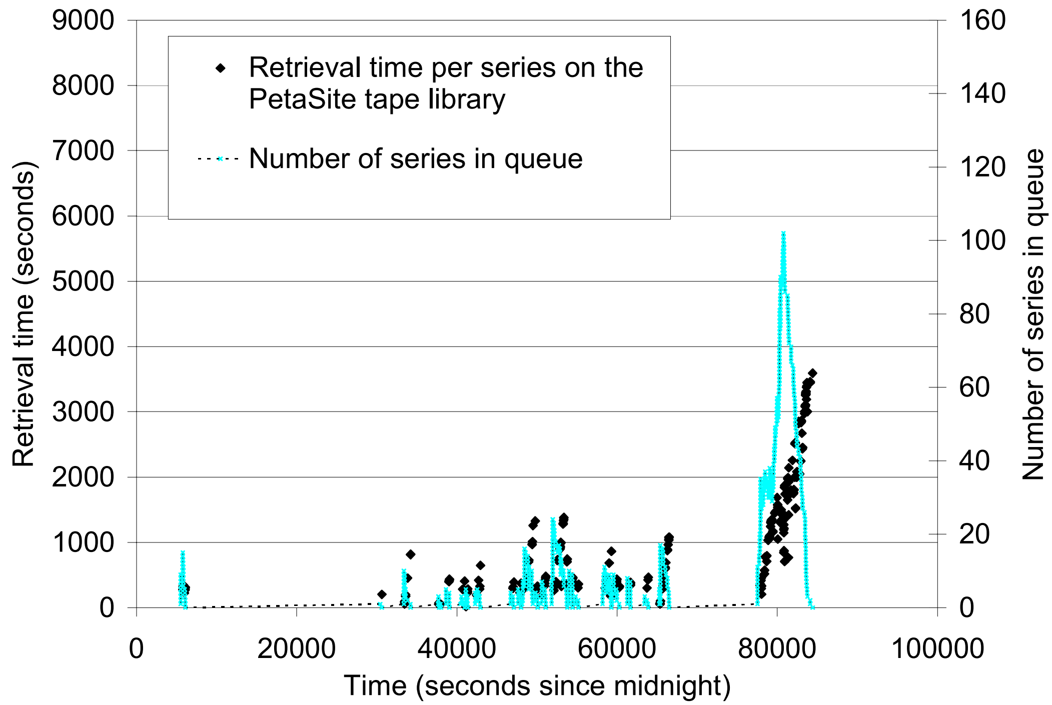

Figure 4.16 Typical usage of the Sony PetaSite tape library over one day.

4.3.4 Discussion

The usage of compression in the MGH PACS is unusual. Images must be transmitted in their uncompressed form since the visualization software installed throughout the hospital does not support compressed image formats conforming to the DICOM standard. For storage purposes on RAID and on tape, images are compressed with a lossless algorithm. When the time comes to transmit an image, it must first be decompressed. AutoMove2 was programmed to decompress each image just before it is sent. This method is probably the simplest to implement, but the decompression time adds to the network transmission time. Ideally, the decompression and the transmission should be performed in parallel. A separate process could be initiated before the connection negotiation to decompress all the images that are to be sent. With this sequence, the network transmission is not delayed by the decompression. As shown in Table 4.4, savings in time would range from 23% (for CR) to 42% (for NM) for the different types of images, assuming that the decompression is started enough in advance to have the first image available when it is needed. The estimate of the savings is done by removing the average decompression time from the sequence formed of the average transmission time plus the average decompression time. Even if the decompression of the first image cannot be initiated soon enough to be ready for the beginning of the transmission, there is still some important saving of time. The first image will remain delayed by the decompression but all the successive images could be sent without this delay since the decompression is generally done faster than the transmission.

The time needed for image transmissions was shown in Table 4.5. These average values hide a Poisson-like distribution with some transmission delays five times bigger than the average value. The cause of such variation is hard to determine. The worst cases were often found to be attributable to particular workstations that were improperly configured in relation to the Ethernet switch to which they were connected. But smaller variations cannot be easily correlated. Different errors in transmission are handled at each layer of the network protocol and retransmission of information occurs at higher levels. Average values of transmission time do not vary enough to be correlated as a function of destination, or the number of concurrent transmissions.

In Figures 4.16 and 4.17, the large peaks at the end of the day come from manually scheduled prefetching retrievals. Hundreds of series are requested at the same time and thus form long queues that take hours to be serviced. This operation is done overnight to minimize interference with user-initiated retrievals. But since there is no priority mechanism to interrupt prefetching retrievals, any series requested by a user during these times will affected. It is a recommendation of this thesis that the retrieval software should be able to prioritize between user requests and prefetching to improve performance from the perspective of the users.

The logic surrounding the processing of retrievals from the tapes might also to be improved. The time needed to read data from the tapes is negligible compared to the whole retrieval process. Removing the tape from the drive when the library is idle would eliminate the time needed to free the tape drive for use and could save time. The read and write throughput of the drive is not critical for retrieval processes but still needs to be efficient to write the data incoming for archiving (over 5,000 MB per day for the MGH) while the system is busy servicing retrieval requests.

4.4 Summary

This chapter has presented the results of the three studies undertaken for the evaluation of the PACS of the MGH. Each study has gathered information on the clinical operation of the system using its databases and other logging facilities. The usage of the PACS equipment is now better known and can be used to plan equipment upgrades or new PACS implementations. The insertion of the timing probes in the transmission processes has allowed us to observe the data-flow and to suggest design improvements and software enhancements.

Chapter 5

5 Conclusion and future work

5.1 Conclusion

PACS were designed and implemented to take over the operations of film libraries in conventional radiology departments in a more effective manner by using digital images, mass storage, networks, and image management software. Three image transfer mechanisms were developed for the MGH PACS in order to distribute the right images to the users when and where they are needed. Autorouting, prefetching, and manual retrieval processes were described and studied in this thesis. It was shown for each process how a continuously running program receives the incoming images or requests and how it uses databases as temporary queues.

The first monitoring method described, which uses master database queries to retrospectively extract data volume and timing information, showed the daily distribution of incoming data and outgoing image requests for each Modality Server. Manual retrieval is the most critical transfer method in terms of transmission delays experienced by users. But, the autorouting, auto-prefetch, and other tasks performed on arrival of new images were found to be the main source of activity on the PACS, and thus required investigation as well.

When images are prefetched or retrieved, they can either be stored on a RAID of a server or in a long-term library. This distributed hierarchical model was found to be well adapted to the storage usage of the PACS with two different levels of accessibility currently in use. However, some imbalance was observed among servers in the time of availability of the fastest storage level (RAID). A centralized and intermediate level between online and offline storage could balance the load more effectively. Another level with reduced accessibility may be needed to deal with expired digital images. Additional features would be needed to handle the expiration of the images and their migration to the limited access of off-the-shelf storage.

The residency timing method provides important insights on the data-flow within sub-processes. This method allows one to plot each successful event that occurs in the transmission of a series of images. For autorouting, it was shown clearly that the images were sent to the review workstation faster than they were received from the imaging modality. The evaluation of manual retrieval showed irregularities in the time needed for the negotiation of a DICOM connection with specific review workstations. The decompression was also found to be an operation limiting the throughput of the transfer. But, the main source of delay remains the time spent in queue when large numbers of images are transferred. When images are retrieved from long-term storage, large amounts of time were spent on mechanical maneuvering in the tape library. Because of this long delay, when prefetching is done in the evening for a large number of series, large queues are formed and could take a few hours to be resolved. This long delay does not have any consequence for prefetching, but practically suspends the access of the tape library for manual retrieval in the meantime.

In summary, from the results of the evaluation of the MGH PACS, it can be seen that building logging facilities into the PACS has been valuable. In order to have a good performance evaluation of a PACS, a comparison with other commercial PACS would certainly prove to be invaluable.. Commercial vendors would normally not allow access to the database nor to the codes of their PACS. When these limitations are considered, one concludes that the tools developed in this thesis cannot be used on PACS different from the one in place at the MGH. The web-based monitoring tool with SNMP query is an exception and can be exported to monitor the hardware of any PACS, as long as it is possible to install a SNMP agent on the computers used by the PACS. The standardization of certain performance measurements into the DICOM standard would help the purchaser of a new PACS with his choice. The implementation of the monitoring tools would also let the system administrator follow the specifications of the manufacturer are in accordance with the performance warranty including any needs for hardware improvement.

5.2 Future work

Some future developments will be implemented before the monitoring tools are employed in daily practice by the PACS support personnel. First, the residency timing must be streamlined in order to incorporate transmission times and queuing delays into the existing monitoring tools. Second, in order to easily install and configure these monitoring tools on other PACS using the same system, an automated setup "script" should be designed to recognize the number and types of servers in the PACS. Installation of the monitoring tools can then be done with preset configurations for each type of server encountered. Finally the monitoring tools need to be integrated with the management software, called "eMon," which is already in use and issues the alert and warning messages to the support staff. These two management programs complement each other to inform the support staff of the current state of the PACS after an alert is issued. On one side, the monitoring tools show the status of several system parameters of interest showing current activity and its trends while on the other side, eMon gets the attention of the support staff after a specific value has reached a predefined threshold.

The information gathered in this thesis also allowed us to objectively assess the performance while the PACS continues to evolve. The transfer algorithms are currently being revised. Instead of looking periodically into each queue for incoming requests, a new triggering process will take place to start the transfer as new requests are added to the queue.

In the transfer sequence,. starting the decompression before the transmission of a series of images should allow for improved transfer throughput by removing the delay experienced for decompression between each image transmission.

The same program handles prefetching and manual retrieval transfers. Auto-prefetching is being used more with the improvement of its retrieval engine and the link to the calendar of scheduled patients throughout the hospital. Prefetching transfer, particularly for retrieval from the long-term archive in the evening, could slow down manual retrievals. For this reason, an independent manual retrieval program or a prioritization mechanism should be investigated.

Bibliography

[1] F. A. J. Mettler, J. E. Briggs, R. Carchman, K. K. Altobelli, B. L. Hart, and C. A. Kelsey, "Use of radiology in U.S. general short-term hospitals: 1980-1990," *Radiology*, vol. 189, pp. 377-380, 1993.

[2] G. B. Spencer, A. H. Sobel, L. Q. Young, and S. J. D. III, "Processes involved in reading imaging studies: workflow analysis and implications for workstation develoment," *J Digital Imaging*, vol. 10, pp. 40-45, 1997.

[3] M. F. Collen, "A brief historical overview of hospital information system (HIS) evolution in the United States," *Int J Biomed Comput*, vol. 29, pp. 169-189, 1991.

[4] F. Bradley, C. O. Vermillion, and W. Anderson, "Medical records on punch cards," *Modern Hospital*, 1954.

[5] F. J. Moore, "Mechanizing a large register of first order patient data," *Methods Inf. Med.*, vol. 4, pp. 1-10, 1965.

[6] J. L. Lehr and F. L. Steinberg, "The radiology information system: its evolution and current status," *Crit Rev Med Inform*, vol. 1, pp. 259-309, 1987.

[7] S. Webb, "In the beginning," in *The Physics of Medical Imaging*, S. Webb, Ed. Bristol: Adam Hilger, 1988, pp. 1-19.

[8] R. P. Woods, "Registration," in *Handbook of Medical Imaging*, I. N. Bankman, Ed. London: Academic Press, 2000, pp. 421-653.

[9] C. Marceau, "What is Picture Archiving and Communication System?," *Proceedings of Spie the International Society for Optical Engineering*, vol. 318, pp. 24-29, 1982.

[10] R. A. Bauman, G. Gell, and S. J. Dwyer, "Large picture archiving and communication systems of the world--Part 1," *J Digital Imaging*, vol. 9, pp. 99-103, 1996.

[11] R. A. Bauman, G. Gell, and S. J. Dwyer, "Large picture archiving and communication systems of the world--Part 2," *J Digital Imaging*, vol. 9, pp. 172-177, 1996.

[12] J. A. Seibert, "Film Digitizers and Laser Printers," in *Practical Digital Imaging and PACS*, vol. 25, *Medical Physics Monograph*, J. A. Seibert, L. J. Filipow, and K. P. Andriole, Eds. Rohnert Park, CA: Medical Physics Publishing, 1999, pp. 107-134.

[13] M. Sonoda, M. Takano, J. Miyahara, and H. Kato, "Computed radiography utilizing scanning laser stimulated luminescence," *Radiology*, vol. 148, pp. 833-838, 1983.

[14] K. P. Andriole, "Computed Radiography Overview," in *Practical Digital Imaging and PACS*, vol. 25, *Medical Physics Monograph*, J. A. Seibert, L. J. Filipow, and K. P. Andriole, Eds. Rohnert Park, CA: Medical Physics Publishing, 1999, pp. 135-156.

[15] M. J. Yaffe, "Digital Mammography," in *Practical Digital Imaging and PACS*, vol. 25, *Medical Physics Monograph*, J. A. Seibert, L. J. Filipow, and K. P. Andriole, Eds. Rohnert Park, CA: Medical Physics Publishing, 1999, pp. 177-224.

[16] T. Chunn and M. Hutchings, "Medical image digital archive: a comparison of storage technologies," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 92-104, 1998.

[17] C. J. Henri, R. K. Rubin, R. D. Cox, and P. M. Bret, "Acquisition and analysis of usage and data-flow statistics for a DICOM-compliant PACS," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 175-177, 1998.

[18] A. Chacko, R. Griffin, J. F. Cook, R. Martinez, B. Lollar, and G. Perez, "Vision and Benefits of a Virtual Radiology Environment for the U.S. Army," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 200-205, 1998.

[19] D. T. Chen, B. K. T. Ho, W. Chao, R. K. Taira, and H. Kangarloo, "PACS War Room for Global Teleradiology," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2435, pp. 178-187, 1995.

[20] R. D. Cox, C. J. Henri, R. K. Rubin, and P. M. Bret, "DICOM-compliant PACS with CD-based image archival," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 135-142, 1998.

[21] R. D. Cox, C. J. Henri, and P. M. Bret, "CD-based image archival and management on a hybrid radiology intranet," *J Digital Imaging*, vol. 10, pp. 168-170, 1997.

[22] R. D. Cox, C. J. Henri, and R. K. Rubin, "Transparent image access in a distributed picture archiving and communications system: the Master Database broker," *J Digital Imaging*, vol. 12, pp. 175-177, 1999.

[23] C. J. Henri, R. D. Cox, and P. M. Bret, "Implementation of a filmless mini picture archiving and communication system in ultrasonography: experience after one year of use," *J Digital Imaging*, vol. 10, pp. 80-82, 1997.

[24] C. J. Henri, R. K. Rubin, R. D. Cox, and P. M. Bret, "Design and implementation of World Wide Web-based tools for image management in computed tomography, magnetic resonance imaging, and ultrasonography," *J Digital Imaging*, vol. 10, pp. 77-79, 1997.

[25] C. J. Henri, R. D. Cox, R. Rubin, P. M. Bret, and G. G. Couch, "Evolution of a filmless digital imaging and communications in medicine-conformant picture archiving and communications system: design issues and lessons learned over the last 3 years," *J Digital Imaging*, vol. 12, pp. 178-180, 1999.

[26] C. J. Henri, R. K. Rubin, R. D. Cox, P. M. Bret, and S. Guberman, "Medical Imaging and the Information Highway-Moving Toward a "Filmless" Radiology Department," *Journal of Optoelectronics and Advanced Materials*, vol. 1, pp. 49-57, 1999.

[27] R. K. Rubin, C. J. Henri, R. D. Cox, and P. M. Bret, "Web-based access to teaching files in a filmless radiology environment," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 520-528, 1998.

[28] R. K. Rubin, C. J. Henri, and R. D. Cox, "Bridging the gap: linking a legacy hospital information system with a filmless radiology picture archiving and

communications system within a nonhomogeneous environment," *J Digit Imaging*, vol. 12, pp. 96-98, 1999.
[29] S. M. Moore, S. A. Hoffman, and D. E. Beecher, "DICOM shareware: A public implementation of the DICOM standard," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 772-778, 1994.
[30] H. Kobayashi, *Modeling and Analysis - An Introduction to System Performance Evaluation Methodology*: Addison-Wesley, 1978.
[31] A. M. Law and W. D. Kelton, *Simulation Modeling and analysis*: McGraw-Hill, 1991.
[32] U. Chandrasekaran and S. Sheppard, "Discrete Event Distributed Simulation," presented at Proceeding Conference on Methodology and Validation, Orlando, FL, 1987.
[33] M. Gerla and T. S. Tai, "LAN/WAN Interconnection to ATM: a Simulation Study," presented at Proc. IEEE INFOCOM, 1992.
[34] A. A. Nilsson and H. Khanmoradi, "A queuing model of picture archiving and communication system (PACS) with a hierarchy of storage," *Proceedings of Third Annual IEEE Symposium on Computer Based Medical Systems*, vol. 6, pp. 1-8, 1990.
[35] H. Kimura and T. Akatsuka, "Modeling and performance analysis of image transfer in PACS," *Medinfo*, vol. 9, pp. 1080-1084, 1998.
[36] T. Masuzawa, A. Miyauchi, H. Kimura, and M. Osada, "Development and evaluation of PACS system simulator," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 642-649, 1994.
[37] Y. Kim and L. Wolfe, "A PACS design and configuration simulator," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 650-658, 1994.
[38] M. Alam, K. Madhav, and S. Kak, "Simulation and analysis of a backbone network for global PACS," *Proceedings of the IASTED International Conference*, pp. 349-352, 1995.
[39] O. R. L. Sheng, H.-M. C. Garcia, C.-P. Wei, T. Ozeki, and P. McNulty, "Distributed Database Design and Modeling for PACS," *Proceedings of Spie the International Society for Optical engineering*, vol. 1234, pp. 256-269, 1990.
[40] R. Martinez, W. H. Sanders, Y. Alsafadi, J. Nam, T. Ozeki, and K. Komatsu, "Performance evaluation of a Picture Archiving and Communication System using stochastic Activity Networks," *Proceedings of Spie the International Society for Optical Engineering*, vol. 1234, pp. 167-178, 1990.
[41] B. Gibaud, Y. Gandon, H. Garfagni, B. Hervault, E. Lussot, B. Barès, and B. Klipfel, "Assessment of the MIMOSA model: experience learned from development of a demonstrator," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2435, pp. 355-367, 1995.
[42] H. Blume, F. Prior, M. C. di Pierro, J. Goble, J. Lodgberg, R. S. Kenney, and F. Goeringer, "Scalability, performance, and fault-tolerance of PACS architectures," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 112-126, 1998.
[43] E. T. Saulnier, R. J. Mitchell, A. A. Abdel-Malek, and K. E. Dudding, "Experience with a Proposed Teleradiology system for Digital Mammography,"

Proceedings of Spie the International Society for Optical Engineering, vol. 2435, pp. 200-209, 1995.

[44] D. T. Chen, "Network Monitoring and Analysis in a Large Scale PACS and Teleradiology Environment," in *Biomedical Physics*. Los Angeles: University of California, 1996, pp. 118.

[45] S. J. Dwyer, B. L. Elias, C. D. Phillips, W. S. Cail, P. T. Huynh, S. Troutman, and M. Meyer, "Evaluation of a PACS for CT and MR: Film System compared to PACS," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2435, pp. 268-275, 1995.

[46] S. E. Hauser, L. E. Berman, and G. R. Thoma, "Performance of RAID as a storage system for Internet image delivery," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2916, pp. 14-22, 1996.

[47] W. Tao, R. H. Zhang, L. J. Huang, and D. J. Valentino, "Optimization of the image cache throughput for a large-scale PACS archive," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3339, pp. 365-374, 1998.

[48] F. W. Prior, K. Niggenaber, S. Khosravani, and G. Meredith, "Impact of Node Architecture on Network Performance," *Proceedings of Spie the International Society for Optical Engineering*, vol. 1234, pp. 334-342, 1990.

[49] S. M. Moore, F. Prior, G. J. Blaine, D. S. Channin, and J. Kilday, "Measured Performance of DICOM-TCP and DICOM-ISO Implementations," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 359-367, 1994.

[50] N. R. Gohel and R. A. Whitman, "Evaluation of multi-megabit networks for medical information delivery," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2711, pp. 560-568, 1996.

[51] H. Kondoh, T. Washiashi, M. Sasagaki, J. Arisawa, H. Nakamura, and K. Inamura, "Development and evaluation of PC-based HIS-RIS-modality-PACS coupling: the results of evaluation of initial stage with personal computer application," *Computer Methods & Programs in Biomedicine*, vol. 57, pp. 63-68, 1998.

[52] S. B. Gay, A. H. Sobel, L. Q. Young, and S. J. D. III, "Processes Involved in Reading Imaging Studies: Workflow Analysis and Implications for Workstation Development," *J Digital Imaging*, vol. 10, pp. 40-45, 1997.

[53] E. L. Siegel and B. I. Reiner, "Recommendations for image prefetch or film digitization strategy based on an analysis of an historic radiology image database," *J Digital Imaging*, vol. 11, pp. 94-99, 1998.

[54] D. V. Beard, B. M. Hemminger, and J. R. Perry, "Interpretation of CT studies: Single-screen workstation versus film alternator," *Radiology*, vol. 187, pp. 565-569, 1993.

[55] W. D. Foley, D. R. Jacobson, and A. J. Taylor, "Display of CT studies on a two-screen electronic workstation versus a film panel alternator," *Radiology*, vol. 174, pp. 769-773, 1990.

[56] L. Benmohamed, A. Mourad, and Y. T. Wang, "Architecture and Perfomance Analysis of A Large Scale Archival and retrival System," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2606, pp. 36-47, 1995.

[57] D. Wilson and R. Glicksman, "Sizing a PACS," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 675-685, 1994.
[58] J. C. Honeyman, M. M. Frost, and E. V. Staab, "PACS component testing: beta and acceptance testing," *Proceedings of Spie the International Society for Optical Engineering*, vol. 3035, pp. 405-412, 1997.
[59] G. Reed and D. H. Reed, "Experience measuring performance improvement in multiphase picture archiving and communications systems implementations," *J Digital Imaging*, vol. 12, pp. 141-143, 1999.
[60] P. D. Shelton, D. K. Lyche, G. S. Norton, J. Romlein, D. P. Lawrence, M. A. Cawthon, J. A. Thomas, and N. Richardson, "Benchmark testing the Digital Imaging Network-Picture Archiving and Communications System proposal of the Department of Defense," *J Digital Imaging*, vol. 12, pp. 94-98, 1999.
[61] N. E. Richardson, J. A. Thomas, D. K. Lyche, J. Romlein, G. S. Norton, and Q. E. Dolecek, "The philosophy of benchmark testing a standards-based picture archiving and communications system," *J Digital Imaging*, vol. 12, pp. 87-93, 1999.
[62] T. E. Lewis, M. C. Horton, T. V. Kinsey, and P. D. Shelton, "Acceptance testing of integrated picture archiving and communications systems," *J Digital Imaging*, vol. 12, pp. 163-165, 1999.
[63] E. T. Saulnier and B. G. Barnett, "High level analysis of medical imaging traffic," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 754-762, 1993.
[64] D. T. Chen, B. K. T. Ho, W. Chao, and R. K. Taira, "Image flow mangement in a PACS network," *Proceedings of Spie the International Society for Optical Engineering*, vol. 2165, pp. 233-240, 1994.
[65] R. L. Kennedy, M. H. Oskin, J. A. Seibert, and S. T. Hecht, "Introduction to DICOM," in *Pratcical Digital Imaging and PACS*, vol. 25, *Medical Physics Monograph*, J. A. Seibert, L. J. Filipow, and K. P. Andriole, Eds. Rohnert Park CA: Medical Physics Publishing, 1999, pp. 287-316.
[66] L. J. Filipow and D. Phil, "Picture Archiving and Communications Systems - An Overview," in *Practical Digital Imaging and PACS*, vol. 25, *Medical Physics Monograph*, J. A. Seibert, L. J. Filipow, and K. P. Andriole, Eds. Rohnert Park: Medical Physics Publishing, 1999, pp. 1-24.
[67] S. Banerjee, "Multimedia traffic analysis and control in a high-speed medical communications environment," *Southcon*, vol. 94, pp. 416-420, 1994.
[68] G. D. Frey and R. Olinger, "PACS Design," in *The Expanding Role of Medical Physics in Diagnostic Imaging*, G. D. Frey and P. Sprawls, Eds. Madison: Advenced Medical Publishing, 1997, pp. 313-337.
[69] B. K. Stewart, "Networks, Pipes, and Connectivity," in *Practical digital imaging and PACS*, vol. 25, *Medical Physics Monograph*, J. A. Seibert, L. J. Filipow, and K. P. Andriole, Eds. Rohnert Park: Medical Physics Publishing, 1999, pp. 259-286.
[70] T. Oetiker and D. Rand, "Multi Router Traffic Grapher," v. 2.8.9, Zurich, 1994, www.mrtg.org.
[71] D. Guerrero, "Network management & monitoring with Linux," vol. 1999: Linux journal, 1997, www.develbnet.es/~david/papers/snmp/.

[72] J. Schönwälder, E. Schönfelder, G. Waters, and S. Waldbusser, "CMU-SNMP for Linux," v. 3.7, Germany, 1999, http://www.gaertner.de/snmp/welcome-3.7last.html.
[73] S. C. Horii, "Networks and Archiving," in *The Expanding Role of Medical Physics in Diagnostic Imaging*, G. D. Frey and P. Sprawls, Eds. Madison: Advanced Medical Publishing, 1997, pp. 339-359.