

**Simulation and Modeling of a 10 Gb/s
Metropolitan Area Network for Radiology**

BY

MPHOENTLE OTUKILE

A Thesis submitted to

The Faculty of Graduate Studies

In Partial Fulfillment of the Requirements for the Degree of

MASTER OF SCIENCE

Department of Electrical and Computer Engineering

University of Manitoba

Winnipeg, Manitoba

© Copyright by Mphoentle Otukile, January 2004

THE UNIVERSITY OF MANITOBA
FACULTY OF GRADUATE STUDIES

COPYRIGHT PERMISSION

**Simulation and Modeling of a 10 Gb/s
Metropolitan Area Network for Radiology**

BY

Mphoentle Otukile

**A Thesis/Practicum submitted to the Faculty of Graduate Studies of The University of
Manitoba in partial fulfillment of the requirement of the degree**

Of

MASTER OF SCIENCE

Mphoentle Otukile © 2004

Permission has been granted to the Library of the University of Manitoba to lend or sell copies of this thesis/practicum, to the National Library of Canada to microfilm this thesis and to lend or sell copies of the film, and to University Microfilms Inc. to publish an abstract of this thesis/practicum.

This reproduction or copy of this thesis has been made available by authority of the copyright owner solely for the purpose of private study and research, and may only be reproduced and copied as permitted by copyright laws or with express written authorization from the copyright owner.

ABSTRACT

The tremendous growth in digital imaging applications and the development of filmless radiology has generated the need to integrate the radiology departments into a digital metropolitan area enterprise system. The medical images generated by these applications are typically large, requiring high bandwidth connections. The bandwidth must be sufficient for the metropolitan links to transmit the medical images within a required response time. The bandwidth allocation should also allow scalability for potential future growth.

In this thesis, we characterize the radiology workload for the Winnipeg hospitals. We use this study to define the bandwidth requirements for the Metropolitan Area Network (MAN) that will interconnect the major hospitals in Winnipeg. Data collection was conducted at the hospitals to gather information on the amount of traffic generated by the imaging acquisition devices. Analysis of the current workload and forecasted workload was conducted, and based on this, our initial guess was to start working with 10 Gigabit Ethernet.

Simulations are performed in OPNET to validate this 10 gigabit bandwidth requirement. A traffic generator was implemented as process models in OPTimum NETWORK performance (OPNET). Comparisons are performed between a distributed and a centralized topology. Simulation results indicate that both the distributed and centralized topologies will be able to sustain the Winnipeg radiology workload with 10 Gigabit Ethernet metro links for a period of 5 years.

ACKNOWLEDGEMENTS

I would like to acknowledge my advisor Dr. J. A Rueda for his support, enthusiasm and inspiration, which was always there when I needed it. I am grateful for his patience and guidance throughout this research. My sincere thanks go to my co-supervisors Dr. Mirek Pawlak and Dr. Jeff Diamond who were responsible for the completion of this thesis. I express my gratitude to Dr. Jeff Diamond and Mrs. Alice Rueda for the technical discussions and for proof-reading the thesis. I also appreciate Mr. Adam Zilkie for his help with OPNET.

I would also like to thank St. Boniface Research Center (SBRC) for the great collaboration in this project. Special thanks go to Mr. Sergio Camorlinga who introduced me to the world of health informatics. I also thank Mr. Ron VanDenakker at Health Science Center who made the data collection possible by introducing me to the necessary people. I thank the diagnostic imaging managers at the Winnipeg hospitals for their time and support in providing me with the data that I needed for this thesis.

I appreciate Nortel Networks' contribution in this thesis. Special thanks go to Mr. Kent Felske and his group, especially Zhonghui Yao and Michel Ouellette.

I am thankful for my parents for their moral support and for encouraging me to pursue my dreams. I thank all my friends who made my student life more enjoyable and adventurous.

Finally, I would like to acknowledge TRILabs for providing financial support and research facilities.

TABLE OF CONTENTS

ABSTRACT	i
ACKNOWLEDGEMENTS.....	ii
TABLE OF CONTENTS.....	iii
LIST OF FIGURES	vii
LIST OF TABLES	ix
CONFERENCES.....	xi
ACRONYMS	xii
CHAPTER 1. INTRODUCTION	1
1.1 OBJECTIVES	1
1.2 BACKGROUND.....	2
1.3 MOTIVATION	5
1.4 10 GIGABIT ETHERNET.....	6
1.5 WINNIPEG HOSPITALS – A CASE STUDY	7
1.6 ORGANIZATION OF THE THESIS	9
CHAPTER 2. RELATED WORK.....	10
CHAPTER 3. HOSPITAL WORKFLOW ASSESSMENT.....	14
3.1 MOTIVATION	15
3.2 WORKLOAD COLLECTION AND ANALYSIS	15

3.2.1 METHOD.....	16
3.2.2 ANALYSIS	18
3.3 HOSPITAL WORKLOAD FORECAST.....	22
3.4 BANDWIDTH REQUIREMENT	25
3.5 WORKLOAD MODELING	27
3.6 SUMMARY	28
CHAPTER 4. NETWORK SIMULATION MODELS	29
4.1 INTRODUCTION.....	30
4.2 THE DICOM AND TCP/IP PROTOCOLS.....	31
4.3 DISCRETE-EVENT SIMULATION	33
4.4 NETWORK MODELS	34
4.5 NETWORK MODELS COMPONENTS	36
4.5.1 THE LONG-TERM STORAGE ARCHIVE	39
4.5.2 THE SHORT-TERM STORAGE SERVERS	39
4.5.3 MODALITIES	40
4.6 TRAFFIC MODELING	42
4.7 IMAGE TRAFFIC SOURCE MODELS	51
4.8 PARAMETERS CONFIGURED.....	54
4.8.1 FTP APPLICATION MODELS	54
4.8.2 PROFILES	56
4.8.3 ARCHIVE PARAMETERS.....	57
4.8.4 SERVER PARAMETERS	58

4.8.5 MODALITY PARAMETERS	59
4.8.6 IP PROCESSING AND CPU SETTINGS.....	61
4.9 ASSUMPTIONS	62
4.10 LAYER 2 AND LAYER 3 SWITCHING	63
4.10.1 LAYER 2 SWITCHING	63
4.10.2 ROUTING.....	64
4.10.3 LAYER 3 SWITCHING	64
4.10.4 ROUTING PROTOCOLS: RIP	65
4.10.5 NORTEL PASSPORT 8600 ROUTING SWITCH.....	65
4.11 TCP PERFORMANCE.....	66
4.11.1 BANDWIDTH*DELAY PRODUCT (BDP).....	67
4.12 SUMMARY	70
CHAPTER 5. SIMULATION RESULTS AND ANALYSIS.....	71
5.1 PERFORMANCE MEASURES	71
5.2 RESULTS.....	72
5.3 DISCUSSION	78
CHAPTER 6. CONCLUSIONS.....	80
6.1 CONTRIBUTIONS.....	80
6.2 RECOMMENDATIONS FOR FUTURE WORK.....	82
APPENDIX A. OPNET PROCESS MODELS.....	85

A.1 OPNET SIMULATION TOOL.....	85
A.2 PROCESS MODELS	87
APPENDIX B. HOSPITAL WORKLOAD MODELS.....	91
B.1 IMPLEMENTATION OF TRAFFIC SOURCE GENERATOR.....	91
B.1.1 ON/OFF TRAFFIC SOURCE MULTIPLEXER PROCESS MODEL.....	92
B.1.2 ON/OFF TRAFFIC SOURCE PROCESS MODEL	94
B.1.3 DICOM SERVER DISPATCHER PROCESS MODEL	98
B.1.4 DICOM SERVER APPLICATION PROCESS MODEL.....	100
REFERENCES	103

LIST OF FIGURES

<i>Figure</i>	<i>Page</i>
1.1 PACS workflow	4
1.2 Current hospitals ATM OC3 network (June 2002).....	8
3.1 Hospitals workload forecast.....	25
4.1 Centralized MAN architecture	35
4.2 Distributed MAN architecture.....	36
4.3 Hospital PACS architecture	37
4.4 Image acquisition node.....	42
4.5 Examination retrieval applications.....	55
4.6 Examination retrieval profiles.....	57
4.7 Configuring archives to support applications.....	58
4.8 Configuring servers to support applications.....	59
4.9 Modality traffic parameters	61
4.10 TCP parameters	69
5.1 Response time for retrieving adhoc examinations	73
5.2 Response time for retrieving new and prior examinations.....	74
5.3 Probability distribution for adhoc response times for distributed topology....	76
5.4 Probability distribution for adhoc response times for centralized topology ...	76
5.5 Probability distribution for new and prior response times for centralized topology	77
A.1 The three-tiered OPNET hierarchy	87

A.2 Example of a state transition diagram	88
B.1 State transition diagram of the process model for creating an ON/OFF traffic source multiplexer	92
B.2 Compound attribute for the ON/OFF traffic source multiplexer.....	92
B.3 State transition diagram of the renewal process model for creating an ON/OFF traffic source	94
B.4 Model attributes of the ON/OFF traffic source	95
B.5 State transition diagram of the server dispatcher process model for creating multiple server applications	98
B.6 State transition diagram of the process model for creating a server application	100

LIST OF TABLES

<i>Table</i>	<i>Page</i>
3.1 Grace Hospital workload assessment.....	17
3.2 Total number of examinations/day.....	18
3.3 Aggregated new volume/day.....	19
3.4 Aggregated prior retrievals/day.....	20
3.5 Estimated examinations performed per day forecast	21
3.6 Weighted retrievals inflation rate per year.....	23
3.7 Examination size annual inflation factor rate.....	24
4.1 Estimated volumes for the Mayo Hospitals	44
4.2 Examination size growth rate.....	45
4.3 Image retrieval file sizes	46
4.4 Image retrieval inflation rate.....	47
4.5 Inter-request time in seconds.....	49
4.6 Adhoc inter-request time in seconds.....	51
5.1 Percentiles for the adhoc retrievals	77
5.2 Percentiles for the new and priors retrievals	78
B.1 State descriptions of the ON/OFF traffic source multiplexer.....	93
B.2 Transition condition descriptions of the ON/OFF traffic source multiplexer	94
B.3 State descriptions of the ON/OFF traffic source	95
B.4 Transition condition descriptions of the ON/OFF traffic source.....	97
B.5 State descriptions of the server dispatcher process model	98

B.6 Transition condition descriptions of the server dispatcher process model...	100
B.7 State descriptions of the server application process model	101
B.8 Transition condition descriptions of the server application process model..	102

CONFERENCES

Part of this thesis has been presented at the following conferences:

Mpho Otukile, Alice Rueda, Jeff Diamond, Jose Rueda, Sergio Camorlinga, "*Modeling and Simulation of 10 Gb/s Radiology Metropolitan Area Networks with OPNET*", OPNETWORK Annual Technology Conference, Washington DC, USA, August 25-29 2003.

Mpho Otukile, Sergio Camorlinga, Jose Rueda, "*A Measurement Study of Diagnostic Imaging Modalities and Workgroups to Design a Suitable Enterprise PACS Network*", SCAR, The 20th Symposium for Computer Applications in Radiology, Boston, Massachusetts, USA, June 7-10 2003.

Mpho Otukile, "*Workload Characterization Analysis and Data*", 2012 Medical Network Requirements: Winnipeg Workshop, Winnipeg, Manitoba, Canada, March 12 2003.

Mpho Otukile, Jose Rueda, Sergio Camorlinga, "*Performance Characterization of 10 Gb/s Radiology Networks*", TRILabs Technology Forum, Edmonton, Alberta, Canada, October 22-23 2002.

ACRONYMS

10GE	10 Gigabit Ethernet
APPD	Advanced Peer-to-Peer Networking
ATM	Asynchronous Transfer Mode
BDP	Bandwidth*Delay Product
BISDN	Broadband Integrated Services Digital Network
Cath Lab	Cardio-angiography
CH	Concordia Hospital
CT	Computed Tomography
CT Ped	Computed Tomography Pediatric
CR	Computed Radiography
DICOM	Digital Image Communications in Medicine
Ded Chest Rad	Dedicated Chest Radiography
DLC	Deer Lodge Center
Echo	Cardiac Ultrasonography
Echo Ped	Echocardiography Pediatric
FDDI	Fibre Distributed Data Interface
Fluoro	Digital Fluoroscopy
Fluoro Ped	Fluoroscopy Pediatric
FTP	File Transfer Protocol
FSM	Finite State Machine
Gbps	Gigabits per second or billion bits per second

Acronyms

Gen Angio	General Angiography
GH	Grace Hospital
HIS	Hospital Information System
HILS	Hardware-in-the-loop simulation
HL7	Health Level Seven
HSC	Health Science Center
I/O	Input/Output
IP	Internet Protocol
IPX	Internetwork Packet Exchange
KP	Kernel Procedure
Km	kilometers
LAN	Local Area Network
MAC	Media Access Control
MAN	Metropolitan Area Network
MARN	Metropolitan Area Radiology Network
Mbps	Megabits per second or million bits per second
MHC	Misericordia Health Center
MRI	Magnetic Resonance Imaging
NM	Nuclear Medicine
OPNET	OPtimum NETwork performance
OSI	Open System Interconnection
OSPF	Open Shortest Path First
PACS	Picture Archival and Communication System

Acronyms

pdf	probability density function
PET	Positron Emission Tomography
Rad Ped	Radiography Pediatric
RIP	Routing Information Protocol
RIS	Radiology Information System
RTT	Round Trip Time
RHC	Riverview Health Center
SACK	Selective Acknowledgements
SBGH	St Boniface General Hospital
SOGH	Seven Oaks General Hospital
SONET	Service Optical Networks
STD	State Transition Diagram
US	Ultrasound
US FA	Ultrasound Fetal Assessment
US Ped	Ultrasound Pediatric
VH	Victoria Hospital
WAN	Wide Area Network
X-ray	Digitized pain film radiography

CHAPTER ONE

INTRODUCTION

1.1 OBJECTIVES

The objective of this research is to characterize the radiology workload and to define networking requirements for the metropolitan area radiology network (MARN) that would interconnect major hospitals in the city of Winnipeg. A good understanding of the radiology workload is essential when planning a MARN. The study also provides the guidelines for future network planning. The approach used in this thesis is:

1. To survey the hospitals in Winnipeg to understand the current radiology network architecture and the amount of data generated by the image acquisition devices.
2. To analyze the image data to get the current network traffic and forecast this amount for future traffic for the next 5 years.
3. To estimate current and future bandwidth requirements.
4. To build a network simulation model to validate the bandwidth requirement over 5 years.
5. To evaluate the performance of the MARN that would connect the major hospitals in Winnipeg.

While other researchers have been focusing on ATM networks [RALP99] [PECH98] [WONG95], the approach in this research is to support 10 Gigabit Ethernet initiatives. The focus of this research is to consider the bandwidth requirement for handling radiology images at the metropolitan area such that the MARN can support the needs of a regional health network. We expect that in the next few years the regional health network will be a necessity as more hospitals and clinics become filmless and more powerful digital image acquisition devices capable of acquiring higher volume of data are being developed.

1.2 BACKGROUND

Digital imaging is rapidly becoming the basis for diagnostic imaging in the radiology departments. The conventional method of capturing medical images on films is gradually being replaced by digital, filmless operation. The film-based diagnostic imaging is a cumbersome, inefficient method. The captured films need to be physically delivered to different medical staff for review and diagnosis. The process of delivering the images to the medical staff is complicated and slow. Furthermore, medical films use a lot of physical storage space. With the advancement of computer and communications technologies, new digital acquisition devices have been developed to replace films. Computerized medical images are generated and stored into a Picture Archival and Communication System (PACS). PACS systems are at the core of the digital diagnostic imaging environment and constitute the important components of an integrated healthcare enterprise system. PACS systems consist of acquisition devices to acquire medical

images, archives and servers for storage and processing, communication networks for the transfer of medical images, and diagnostic workstations for displaying and manipulating medical images and related patient data. The digital images stored in PACS are transferred using the DICOM (Digital Imaging and Communications in Medicine) standard developed by the American College of Radiology and the National Electrical Manufacturers Association [DICO03]. This standard allows interoperability between devices from different vendors, allowing them to communicate medical images.

PACS are mostly integrated with the hospital information system (HIS) and radiology information system (RIS) to optimize the hospital workflows. Figure 1.1 illustrates the PACS workflow. The image acquisition devices get the patients' information from the RIS. The RIS sets schedules for the modalities, routes and prefetches coordinated examinations, keeps track of images and produces modality reports. A RIS/modality interface is crucial and should be bi-directional for the modality to receive patient demographics from the RIS and to allow the modality to send information to the RIS for validation. HIS is responsible for billing, keeping track of appointments and managing schedules. The Health Level Seven (HL7) standard is used to interface HIS and PACS.

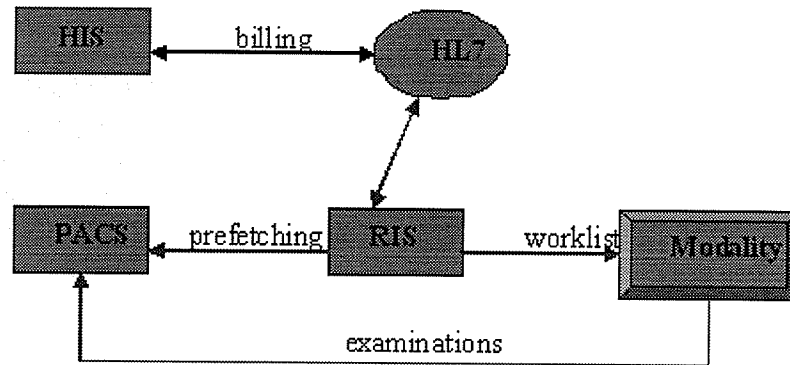


Figure 1.1: PACS workflow

Until recently, PACS have been operating as isolated “islands” within a hospital or healthcare enterprise system [CARM01] [SIEG02]. Many healthcare institutions now realize that integrating these isolated islands is the key to a better quality of health care. Having an integrated healthcare enterprise system can improve accessibility to medical data and reduce healthcare costs.

However, the transition from isolated islands to a filmless integrated healthcare enterprise-wide system has many challenges. While enterprise-wide PACS systems reduce the cost of imaging, the biggest challenges are storing and archiving of medical images, and building the network infrastructure. Medical images are typically large. The PACS system requires large storage archives and high bandwidth connections for communication. The healthcare network infrastructure should also be carefully planned such that it can handle the large volumes of medical data and ensure adequate response time for retrieval of medical images. The challenge is to keep in pace with the emerging

networking technologies and telecommunication infrastructures to provide the required quality of service needed in healthcare environment.

1.3 MOTIVATION

Digital imaging applications are one of the most bandwidth demanding applications. To integrate the hospital network into a single digital enterprise system over the metropolitan area is not a simple task. The high volume of digital images produced has to be distributed in an enterprise-wide PACS. Changes need to be made to the existing radiology networks and storage archival architectures. Current network infrastructures suffer from the deficiency of bandwidth limitations that cannot satisfy future demands. In addition, they suffer from their centralized nature that fails to satisfy the distributed requirement of a regional health enterprise system.

Many high-speed broadband networks have been designed for radiology applications to improve access, retrieval and distribution of medical images. These networks are mostly implemented using Asynchronous Transfer Mode (ATM)-based Broadband Integrated Services Digital Network (BISDN) [RUGG98] [CGGR97] [TSIK96]. But with the volume of gigabytes generated each day by the acquisition devices and archived images accessed by the diagnostic workstations, there is a need for higher bandwidths in the radiology networks.

Technology advancement has made possible complex acquisition devices capable of creating an overwhelming amount of data. For example, the Mayo Hospital [MAYO01] conducted a study that showed that the number of images per examination, and consequently the data set per examination, for computed tomography (CT) and

magnetic resonance imaging (MRI) increases every year. The UCLA Laboratory of Neuro Imaging [SGI01] found that the total stored data from a positron emission tomography (PET) examination increases at a rate of 8 tera bytes per year. In addition to the technology advancement, hospitals [MAYO01] are reporting a proliferation of 24 % in the number of examinations performed each year. As such the integrated healthcare enterprise systems require higher bandwidth links in their metropolitan area networks (MAN) to handle the growing amount of image data generated daily.

The emergence of promising technologies such as 10 Gigabit Ethernet provides the prospect for deploying MARN at a lower cost than the currently deployed technologies such as ATM and SONET [GUNS00]. Migration to 10 Gigabit bandwidth connections on the MARN provides the infrastructure to meet the bandwidth requirement of the radiology enterprise networks.

1.4 10 GIGABIT ETHERNET

With the official de facto standard released on June 2002, 10 Gigabit Ethernet (10GE) is the latest development of Ethernet. It is mostly deployed in local area networks (LAN) [10GEA]. 10GE is basically the fastest version of Ethernet. It supports data rate at 10 gigabits per second (Gbps). It provides not only the big bandwidth but allows a seamless integration of LANs to form a single shared high capacity MAN in ring or mesh topology since all traffic is in Ethernet format. Therefore there is no need for packet fragmentation, reassembly or address translation. This eliminates the need for conventional routers that are much slower than switches.

10 Gigabit Ethernet extends the distance limitation by using single mode dark fibers in MAN to support up to 40 kilometers (km) links. It can also interoperate with the current technologies without modifying the existing network infrastructures. The key benefits of 10GE are:

- It is easy to use.
- It allows a straightforward migration from the previous Ethernet versions (10/100/1000 Mbps).
- It provides a lower cost for maintenance and support compared to the current high speed technologies such as ATM and SONET.
- Network administrators are familiar with Ethernet so little training is required.
- It supports new applications and new data types.
- It provides flexibility in network design.
- It provides interoperability between different vendors.

As a result of these benefits, 10 Gigabit Ethernet is a logical path towards the deployment of MARN. It will provide radiology applications with lower network latency due to its high speed links. Its high bandwidth can also accommodate the bursty nature of the enterprise traffic [10GE02].

1.5 WINNIPEG HOSPITALS – A CASE STUDY

The Winnipeg Regional Health Authority [WRHA99] is comprised of nine major hospitals within a 25 km radius in its region. These hospitals are: Health Science Center

(HSC), St Boniface General Hospital (SBGH), Grace Hospital (GH), Victoria Hospital (VH), Seven Oaks General Hospital (SOGH), Misericordia Health Center (MHC), Riverview Health Center (RHC), Concordia Hospital (CH) and Deer Lodge Center (DLC). HSC is a 900-bed tertiary care University teaching hospital and the largest health care referral in the province. SBGH is a 631-bed second largest tertiary care University teaching hospital. GH is a 281-bed hospital. SOGH is a 275-bed hospital. MHC is a 418-bed. DLC is a 487-bed long term care and rehabilitation facility.

These hospitals provide a combination of analog and digital diagnostic imaging services such as Ultrasound (US), Magnetic Resonance Imaging (MRI) and Computed Tomography (CT). The radiology departments of the hospitals perform more than 500,000 examinations per year and generate over 20 tera bytes (TB) of new uncompressed imaging data per year. The current IT infrastructure serving the radiology departments consists of a high-speed ATM OC3 (155 Mb/s) metropolitan network (as of June 2002) interconnecting PACS systems at HSC, SBGH and GH (Figure 1.2). These PACS systems are sending images to a central archive at SBGH. Only CT, US and MRI images are stored. The workflows of other modalities are either film or locally stored.




Figure 1.2: Current hospitals ATM OC3 network (June 2002)

The future MARN will connect over 140 modalities distributed in the nine hospitals. The PACS architecture is hybrid, with more than one global storage archive and more than one server for immediate access of new acquired images and prior images (old exams) by radiologists.

1.6 ORGANIZATION OF THE THESIS

Chapter 1 gave a brief introduction to digital imaging in general and provided the objectives statement of the thesis. The rest of the thesis is organized as follows:

- Chapter 2 discusses work that is related to this thesis.
- Chapter 3 provides an overview on the Winnipeg hospitals workload characterization and the analysis of the current and future bandwidth requirements.
- Chapter 4 describes the OPNET network models and assumptions used in the simulations.
- Chapter 5 provides the simulation results of the performance evaluations of the network models.
- Chapter 6 gives the summary and conclusions of the thesis and recommended future work.
- Appendix A provides an introduction to the OPNET simulation tool.
- Appendix B presents the traffic source models based on the hospitals workloads implemented and the model engines in OPNET.

CHAPTER TWO

RELATED WORK

This chapter describes related work aimed at solving problems similar to those outlined in Chapter 1. A number of related projects are compared and contrasted with our research.

In recent years, a growing number of radiology examinations have been generated digitally using imaging modalities such as US, NM, MRI, and CT. This evolution from analog to digital imaging has been challenging the computer networking and medical imaging community for new technologies and techniques to manage the images to acquire, process, store, transmit and receive [MART90]. Computers now have a central role in both the generation of medical images and the management of images. The image transfer also requires a high speed computer network that allows the retrieval and exchange of the digital images in a short period of time.

Many radiology networks are being developed with the same general goal as ours; that is, to distribute medical images and facilitate communication between radiologists and physicians. Many of these networks have addressed solutions for a hospital-wide environment, such as that reported by Georganas et al. [GEOR89], Ritib [RATI93] and Prior et al. [PRIO95]. They did not consider factors such as the rapid health system geographic expansion and the need to access information from any location, that make

the radiology network absolutely critical to getting access to healthcare information. Some researchers such as De et al. [CEDE98] have integrated the hospital sites within their region but their main focus was on the sharing of computerized patients' records such as medical imaging reports and not the medical images. Our goal is to extend beyond these limitations and focus on a metropolitan area radiology network that would allow access to radiology information at all hospital sites within a region. Other health institutions such as the Mayo Hospital [MAYO01] and the US Army Virtual Radiology [MART90] have done similar work as ours by integrating the hospital networks in their region into a single digital enterprise system.

A number of papers have discussed the implementation and/or evaluation of high-speed networks for radiology to link hospital consortia. Hall et al. [HALL92] report the implementation of a teleradiology system that uses a dialup switched wide area Network (WAN) for connecting two military hospitals and an academic diagnostic radiology department. Researchers such as Georganas et al. [GEOR89] have focused on packet switching technologies such as Fiber Distributed Data Interface (FDDI). These technologies have proven to be inadequate to meet the high-bandwidth demand of diagnostic imaging because of their low transmission rates. ATM-based broadband integrated services digital networks (B-ISDN) is the current state-of-the-art high-speed technology and has been implemented in numerous hospitals such as the University of California to connect Mount Zion Hospital and the San Francisco VA Medical Center [RUGG98] [WONG95]. Another ATM project called the U.S. Army Virtual Radiology Environment (USAVRE) was implemented at Brooke Army Medical Center to connect all the major US Army's medical centers and Regional Medical Commands [RALP99].

The USAVRE also discusses issues such as network architecture and traffic modeling. The current ATM networks provide adequate bandwidth and performance. With shortage of bandwidth in transporting high-resolution medical images, due to increased image traffic and future demands, hospitals are now starting the initiatives to deploy 10 Gigabit Ethernet in their backbone networks. The University Health Network (UHN) which links Universities and Teaching Hospitals in Toronto has upgraded its backbone to 10 Gigabit Ethernet to eliminate bottlenecks in their data intensive distributed core and support high-resolution medical image transfers and other bandwidth-intensive applications [FOUNET]. Our research also supports 10 Gigabit Ethernet in the enterprise backbone network. But unlike our work, the UHN community did not analyze the bandwidth requirements and performance evaluation before implementing.

Little research has been done to characterize the radiology workload for bandwidth analysis. Studies such as that by Erickson et al. [MAYO01] have done similar approach as ours by analyzing the workload of their radiology departments, diagnostic imaging devices and network equipment. The only difference with our research is that they have focused on characterizing the workload for storage requirements and not bandwidth requirements. The USAVRE also conducted a similar workload analysis as ours for their US Army ATM backbone network implementation. They developed a detailed analysis of the workload and workflow patterns in their current radiology network. Our approach also is to model and simulate the workload and use it as a baseline for the new proposed metropolitan area network for the Winnipeg hospitals. We analyze the workload for each modality at each hospital site. The workload consists of

the number of examinations per day per modality, the type of images produced, the number of images per examination and the size of the images.

Different techniques have been used to predict and analyze the performance of PACS under various load conditions. One technique is the use of stochastic activity networks (SNAs) reported by Sanders et al. [SAKU92] and Martinez et al. [MART93]. SANs are a stochastic extension to Petri nets. They are used to build models of PACS components such as modalities, viewing stations, network and archives. Other researchers such as Reijns [REIJ95] and William [WILL91] have focused on implementing PACS testbeds to test the network performance. Implementing testbeds is very costly. Researchers such as Lawrence et al. [LAWR85] have used queuing models to analyze PACS performance. The disadvantage of using queuing models is that they need too many simplifications and hence the emerging theoretical model may not be valid. Our method is modeling and simulation which is the most common method nowadays in network performance evaluation. Researchers such as Martinez et al. [MART90] [MART98], Centeno et al. [CENT00] and Lin et al. [SHALIN] have used simulation and modeling to study PACS performance.

CHAPTER THREE

HOSPITAL WORKFLOW ASSESSMENT

The objective of this chapter is to provide, based on the current hospital workload analysis, a five-year workload forecast and to estimate the bandwidth requirement for the metropolitan area radiology network (MARN) interconnecting the nine hospitals in Winnipeg. The bandwidth requirement estimates must be accurate in order to be useful. This is why the analysis has been done on real hospital workloads. The chapter describes the workload analysis of the radiology departments in the Winnipeg hospitals. Planning for a MARN requires a good understanding of the radiology workload characteristics. Section 3.1 discusses the motivation for understanding the hospitals workload characteristics. Section 3.2 describes the workload collection and analysis. Section 3.3 provides an estimated forecast on the workload for the next 5 years. The workload forecast will be used to estimate the future bandwidth requirement, which is discussed in Section 3.4. Section 3.5 presents the hospital workload models for performance analysis. Finally, Section 3.6 provides a summary of this chapter.

3.1 MOTIVATION

Prefetching is used if it is known which images will be required for diagnosis beforehand. Most radiologists want to eliminate prefetching as prefetching algorithms are difficult to implement [MAYO01]. This means that the network should be able to handle on-line traffic as it is requested. The role of the hospital workload characterization is to understand how the radiology communication network will respond to a variation of workload. This enables network capacity planning based on current and projected workloads. The hospital workload characterization has a significant impact on the performance evaluation of radiology networks. An understanding of the nature of the radiology workload can help in defining performance measures and interpreting the simulation results. A lack of understanding of the workload, traffic forecast and bandwidth estimation for the radiology networks could lead to wrong conclusions.

3.2 WORKLOAD COLLECTION AND ANALYSIS

The goal of the hospital workflow assessment was to gather traffic information on the image acquisition devices that generate medical images in the radiology departments. The assessment will be used to predict the amount of traffic that will flow in the MARN. The radiology departments in the 9 hospitals have 17 different types of modalities: the digitized computed tomography (CT), ultrasound (US), magnetic resonance imaging (MRI), nuclear medicine (NM), general angiography (Gen Angio), cardio-angiography (Cath Lab), digitized plain film radiography (X-ray), computed radiography (CR), digital fluoroscopy (Fluoro), ultrasound fetal assessment (US FA), cardiac ultrasonography

(Echo), dedicated chest radiography (Ded Chest Rad), pediatric ultrasound (US Ped), pediatric echocardiography (Echo Ped), pediatric computed tomography (CT Ped), pediatric radiography (Rad Ped) and pediatric fluoroscopy (Fluoro Ped) . The CT and MRI modalities located at two of the hospitals send their images to the central archive at SBGH. The two hospitals each have a commercial Picture Archival Communication System (PACS) to manage, store and distribute digital ultrasound images. Another hospital has a mini-PACS for its local images. The rest of the modalities are still paper-based, printed on films.

3.2.1 METHOD

The method we used in collecting the data involved visiting and interviewing diagnostic imaging managers of the nine hospitals. The managers were also given a questionnaire that surveys both the digital and analog modalities they have. The analog traffic was included in the analysis assuming that in the future all the modalities will be digitized, though with more data. The hospitals have a total number of 144 modalities generating image data. The analysis that we performed on each modality was based on the following:

- The size of the image files generated by the modality
- The number of images per examination
- The number of examinations performed each day
- The number of new examinations that require previous examinations (i.e. priors) for diagnosis comparison

- The number of priors if required
- The projected annual growth of the total number of examinations generated each day by each modality. The projected annual growth is the increase in the number of examinations performed and this does not include the increase due to the number of images per examination or increase in the image size.

Table 3.1 shows an example of the workload analysis for one of the hospitals. For each modality type, we have the image size supported, the number of images per examination, and the number of examinations performed each day. The “Qty” column shows the total number of modalities of the same type from different vendors. The “% priors required” is the percentage of the new examinations that require previous examinations to be sent along for diagnosis. The “No. of priors” is the number of previous examinations that will be sent if required. The percentage growth per year is the projected annual increase in the number of examinations performed by each modality. The same information was collected for the other eight hospitals and the summary is provided in Table 3.2.

Type	Qty	Image size (MB)	Images/ exam	Exams/ day	% priors required	No. of priors	growth/ year (%)
CT	1	1.05	45	20	50	7	10
US	1	0.26	24	10	50	2	10
	2	0.26	24	13	50	3	10
Fluoro	1	2.1	12	10	50	3	10
	2	0.26	6	3	50	1	10
	1	0.26	10	3	50	1	10
	1	0.33	10	1	50	1	10
CR	3	10.5	3	120	50	1	10
NM	1	0.07	9	30	50	1	10
	1	0.13	9	30	50	1	10

Table 3.1: Hospital workload assessment


3.2.2 ANALYSIS

The hospital workflows collected were used to estimate the bandwidth requirement that will meet the needs of a metropolitan area radiology network in Winnipeg. The first step was to estimate the aggregate volume of image data generated each day by the hospitals. Table 3.2 shows the total number of new examinations performed each day by the hospitals. Table 3.3 shows the aggregate volume of new data for each modality type and the total volume of new image data generated each day by the 9 hospitals.

← Hospital →

Type	Hsp 1	Hsp 2	Hsp 3	Hsp 4	Hsp 5	Hsp 6	Hsp 7	Hsp 8	Hsp 9	TOTAL exams/day
US FA	40	62								102
US	1300	80	36	40	27	45	22			1550
Fluoro	72	28	20	8	8	4	25	7		172
US Ped	30									30
NM	84	91	60	11						246
Echo	266	16								282
Echo Ped	25									25
CT Ped	12									12
CT	78	130	20	25	18	14	14			299
Gen Angio	50	10								60
Cath Lab	0.27	40								40.27
X-ray	247	782		55		50		7	15.29	1156.29
Rad Ped	213									213
Ded Chest Rad	110	30								140
Fluoro Ped	4									4
MRI	18	60								78
CR			360		100					460
TOTAL exams/day	2549.27	1329	496	139	153	113	61	14	15.29	4869.56

Table 3.2: Total number of examinations/day




Type	Hsp 1	Hsp 2	Hsp 3	Hsp 4	Hsp 5	Hsp 6	Hsp 7	Hsp 8	Hsp 9	TOTAL MB/day	TOTAL GB/day
US FA	240	400								640	0.63
US	2600	1450	216	450	608	900	66			29689.5	28.99
Fluoro	90	140	260	12	100	20	94	10.5		725.88	0.71
US Ped	600									600	0.59
NM	513	19	51	5						587.81	0.57
Echo	6650	0								6650	6.49
Echo Ped	313									321.5	0.31
CT Ped	528									528	0.52
CT	12100	9360	900	977	1125	210	53			24724.06	24.14
Gen Angio	2188	413								2600	2.54
Cath Lab	0.48	560								560.48	0.55
X-ray	8550	22256		83		400		10.5	151	31449.57	30.71
Rad Ped	2911									2911	2.84
Ded Chest Rad	320	120								440	0.43
Fluoro Ped	60									60	0.06
MRI	432	5100								5532	5.4
CR			10880		3000					13880	13.48
TOTAL	61494	39818	12226	1526	4833	1530	212	21	151	121810.8	
TOTAL	60.05	38.88	11.94	1.49	4.72	1.49	0.21	0.02	0.15		118.96
GB/day											

Table 3.3: Aggregated new volume/day

Tables 3.2 and 3.3 show that the radiology departments perform a total of 4870 examinations per day, generating an equivalent of 119 GB new uncompressed medical data each day. This volume includes both scheduled examinations and emergency examinations and does not take into account any video images. In the fiscal year of 2001-2002, 25% of the performed examinations were emergency examinations [BLAR03].

These examinations require instant access to previous examinations and immediate retrieval for diagnosis.

Radiologists use the viewing stations to access medical images for diagnosis. Previously stored images (called priors) that are relevant to the newly acquired examinations are sent to the desired destinations along with the new images. Table 3.4 shows the volume of priors required for each modality type and the aggregated expected priors' volume for all the hospitals.



Type	Hsp 1	Hsp 2	Hsp 3	Hsp 4	Hsp 5	Hsp 6	Hsp 7	Hsp 8	Hsp 9	TOTAL MB/day	TOTAL GB/day
US FA	120	200								320	0.31
US	104000	725	294	135	365	450	211			106180	103.69
Fluoro	90	70	130	3.6	180	10	141	3.15		627.19	0.61
US Ped	300									300	0.29
NM	379.8	8.34	25.3	1.44						414.89	0.41
Echo	13300	0								13300	12.99
Echo Ped	187.5									187.5	0.18
CT Ped	528									528	0.52
CT	12100	2808	3150	293	675	105	210			19341	18.89
Gen Angio	1968.75	247.5								2216.25	2.16
Cath Lab	0.24	224								224.24	0.22
X-ray	8550	6676.8		24.8		400		3.15	151	15805	15.43
Rad Ped	1455.5									1455.5	1.42
Ded Chest Rad	320	120								440	0.43
Fluoro Ped	30									30	0.03
MRI	324	1530								1854	1.81
CR			5400		3600					9000	8.79
TOTAL	143654	12610	8999	458	4820	965	562	6.3	151	172224	
TOTAL	140.29	12.31	8.79	0.45	4.71	0.94	0.55	0.01	0.15		168.19
GB/day											

Table 3.4: Aggregated prior retrievals/day

Using the percentage estimations of the required priors collected from the hospitals, we found that 66% (3193 studies) of the new examinations require priors. The prior traffic contributes an aggregated volume of 168GB of image traffic per day. At Hsp 3, the estimated average number of automatic retrievals is 175 examinations per day, but only 100 examinations are viewed at the diagnostic viewing stations. We assume that the other retrievals are used by the department for research. Also, based on the hospital usage assessment conducted by [BLAR03], 20% of the scheduled examinations require "ad-hoc" retrievals. Ad-hoc retrievals are examinations acquired in a hospital and accessed by radiologists from a different hospital. This adds an additional of 24GB of image traffic per day to support remote access.

	2002 (Current)	2003 Estimate	2004 Estimate
US FA	102	107	112
US	1550	1562	1574
Fluoro	172	177	184
NM	246	256	266
Echo	282	300	320
CT	299	309	320
Gen Angio	60	81	117
Cath Lab	40	42	44
X-ray	1156	1198	1241
Ded Chest Rad	140	140	140
MRI	78	84	90
CR	460	502	547
Pediatric	284	292	300
TOTAL/day	4870	5049	5256

Table 3.5: Estimated examinations performed per day forecast

There are a total of 144 digital and analog modalities in all the radiology departments combined. These modalities perform 4870 examinations per day. This number is estimated to increase by 4% every year (Table 3.5).

Based on the above analysis, the current total traffic generated each day that flows into the metro network connecting the hospitals would be 311 GB. This volume consists of traffic flowing between the modalities, storage archives and the short-term servers only. If we assume an 8-hour period of busy day, we would have an average traffic of 38.96 GB per hour each day. This average traffic would require a mean bandwidth of 93 Mb/s for the metro links. Since Ethernet bandwidth technologies come in multiples of 10, the current traffic can be serviced by 100 Mb/s Ethernet metro links.

3.3 HOSPITAL WORKLOAD FORECAST

The basis of any good network planning is an accurate workload forecast. Workload forecasting is the process of predicting how the network workloads will vary in the future. This process answers "what-if" questions such as:

- How will the number of examinations performed daily vary in the next 5 years?
- How will the image retrievals vary over time?
- How will the size of the images change over time?

Answering such questions involves evaluating workload trends, starting with the current workload and identifying the sources of growth on workload.

The purpose of the hospital workload forecast is to predict the radiology workload so as to assess future bandwidth required to carry the workload. Radiology workload

changes with population growth and technology advancement. From the conducted hospital workload assessment discussed in Section 3.2, the managers have estimated the annual projected growth in the number of examinations performed each day by each acquisition device. For the image retrievals, we used the weighted average of the estimated growth reported by each hospital to determine the annual increment in the number of retrievals. Table 3.6 shows the weighted projected growth used in forecasting the retrieval workload. The MRI and Gen Angio are available only at Hsp 1 and Hsp 2. Hsp 1 reported high annual projected growth in the number of examinations of 80% and 250% for Gen Angio and MRI respectively. The calculations we made here for the annual inflation rate in the number of the examination was based on the estimations given by the managers and we have no method of verifying these estimations.

Modality	Annual inflation rate (%)
US FA	5
US	0.77
Fluoro	3
US Ped	2
NM	5.11
Echo	6.42
Echo Ped	10
CT Ped	2
CT	3.28
Gen Angio	37.5
Cath Lab	5
X-ray	3.61
Rad Ped	2
Ded Chest Rad	0
Fluoro Ped	2
MRI	59.23
CR	9.02

Table 3.6: Weighted retrievals inflation rate per year

We also assumed that the size of the examination increases, either due to increase in the image data size or increase in the number of images per examination. Since such data is not available from the local hospitals, we used the estimations from the Mayo Hospital [MAYO01] to determine the examination size inflation rate per year. The estimation is provided in Table 3.7. The table shows the growth factor in the examination size. The growth factor is the rate at which the examination size increases annually.

Modality	Examination size inflation factor rate/year
US, US FA, US Ped, Echo, Echo Ped	1.13
NM	1.39
CT, CT Ped	1.18
Gen Angio, Cath Lab, Fluoro, Fluoro Ped	1
X-ray, Rad Ped, Ded Chest Rad, CR	1.54
MRI	1.28

Table 3.7: Examination size annual inflation factor rate




Figure 3.1: Hospitals workload forecast

Using all the estimations given in this section, we predict that in the next 5 years the estimated radiology workload will be roughly 5 times of the current workload, as shown in Figure 3.1.

3.4 BANDWIDTH REQUIREMENT

This section uses the workload estimation provided in the previous sections to forecast future bandwidth requirement. The estimated bandwidth should be able to handle the workload and support the real-time quality of service. The radiologists at the hospitals will be expecting images to be available at the diagnostic viewing stations within 2

seconds of response time [CART03] [CHIM92]. We have made a few assumptions in determining an initial guess for future bandwidth requirement. The network must accommodate the bursty nature of the enterprise traffic. We assumed the traffic burstiness will increase the bandwidth requirement by a factor of 100. If we start with the current required mean bandwidth of 100 Mb/s from Section 3.2.2, the burst factor will increase the bandwidth to 10 Gb/s. Having the MARN configured in a meshed topology can reduce the bandwidth requirement by 1/3. In other words, reduces the bandwidth requirement to 3.33 Gb/s. Also, taking into consideration that most of the traffic is consumed locally the bandwidth will be further reduced by 1/3 or to 1.11 Gb/s. Locally means the examinations acquired in a hospital are mostly retrieved and viewed within the same hospital. There are also other factors that can increase the radiology workload. New acquisition devices will increase the workload. We assume there is a 10% increase per year due to new acquisition devices. If there is a better network, other networking usage will increase. This can add other additional of 10% traffic per year. If the MARN is in place, we assume there will be an increment of 10% of transit traffic per year. Transit traffic is defined as traffic originating from outside the network and passing through the network. Taking all these factors into account, we can estimate the bandwidth requirement for the next 5 years to be:

$$\text{bandwidth} = 1.11 \text{Gb/s} \times (1.1 \times 1.1 \times 1.1 \times 1.1)^5 = 7.5 \text{Gb/s}$$

Thus our initial guess is that roughly 10 Gigabit Ethernet will be required for the MARN in order to compliant to the bandwidth requirement in the next 5 years.

3.5 WORKLOAD MODELING

It is very difficult to model workloads with large data sets. Therefore, in order to simulate the workload, it is necessary to reduce and summarize the information needed to describe the workload. In other words, a higher level abstraction which captures the most relevant characteristics of the real workload is used to represent the workload model. The choice of characteristics and parameters that describe the workload depends on the purpose of the performance study. Workload models exhibit several advantages over actual workloads or traces [MENA02]. In the workload model, it is easy to change model parameters to reflect changes in the network or in the real workload. The workload model is also useful when predicting the network performance using workload forecasts. For example, if one wants to increase the workload in the model, one has to change just one parameter. This parameter could be either the interarrival time or the off (think) time. Using a workload model, it is possible to repeat simulation experiments under a statistically similar condition that are nevertheless not identical. For example, a simulation can be run several times with a different seed value for each simulation. This is useful in determining confidence intervals for performance measurements.

The Winnipeg hospital radiology workloads analyzed in the previous sections are used to drive discrete event simulations to validate the initial guess of 10 Gigabit connections for the radiology MAN in the next 5 years. The main performance metrics is the response time to retrieve an examination from a storage archive. The parameters of the hospital workload model are the distributions of the examination sizes and the interarrival times between examinations. Appendix A gives an introduction to the OPNET simulation tool used to implement the Winnipeg radiology workload models.

Appendix B describes in detail the implementation of the workload models for the examinations generated by the hospitals.

3.6 SUMMARY

Understanding and characterizing the hospital workload is a key factor in the performance study of radiology metro networks. This chapter has also pointed out why it is necessary to have workload models for performance analysis. The next chapter, Chapter 4, describes simulation network models and the assumptions made in the performance analysis of the metropolitan area radiology network.

CHAPTER FOUR

NETWORK SIMULATION MODELS

One of the thesis objectives is to validate through simulations the bandwidth requirement discussed in Chapter 3. A series of discrete event simulations are performed to investigate the bandwidth requirement of the metropolitan area radiology network (MARN) between the Winnipeg hospitals over a period of 5 years. This chapter describes some of the issues addressed with the simulation models and the modeling techniques. The chapter also outlines the network architecture models for our performance analysis study. It describes the network models and the assumptions made. Section 4.1 provides an introduction to PACS network architectures. Section 4.2 discusses the DICOM protocol and the TCP protocol. Section 4.3 describes the simulation modeling and discrete event simulations. Section 4.4 provides an overview of the two network topologies considered in the performance analysis. The network components are listed in Section 4.5. This section also discusses how the modality models generate their traffic. Section 4.6 provides the workload assumptions and describes how they are translated into traffic model parameters. Section 4.7 discusses the image retrieval request source models for the servers and the archives. Section 4.8 discusses the parameters configured in the source models to generate the required amount of traffic in the network models. Assumptions made in designing the networks are described in Section 4.9. Layer 3 switching is

discussed in Section 4.10. Section 4.11 discusses TCP and its role in the performance analysis of high-speed networks. Finally the chapter is concluded in Section 4.12.

4.1 INTRODUCTION

Most PACS networks are implemented using a centralized architecture where all the medical images and data are stored in a single central storage archive and then distributed to different local workgroups with temporary storage. The main advantage of this architecture is that it reduces the routing of medical images across multiple locations. In addition, this design is made simple and easier to maintain than a distributed architecture by reducing the need to keep track of the location of the medical images across multiple locations in the network. The drawback of this implementation is the potential occurrence of spikes in network traffic due to simultaneous access of images from multiple workstations and modalities. The consequences of these spikes are increased delays for accessing diagnostic imaging data, reduced productivity and user frustration. This architecture also has poor fault tolerance because of the single point of failure. If the archive fails, the whole system goes down. It also has poor scalability as network components are added to the network.

With the advancement of computer technology and communication networks, the trend has moved towards distributed architectures where medical images and data can be stored at different locations on a network and still be accessible from any other part of the network [RATI93]. In a distributed architecture, there are multiple storage archives. Each acquisition device in the network can send their images to any one of the archives. This

architecture reduces the processing and network loads handled by each archive and also improves the network scalability. This architecture also eliminates the single point of failure, since failure of one storage archive will not completely halt the flow of image traffic in the network. Hence it has a better fault tolerance, robustness and availability. This architecture reduces access time and improves user satisfaction. Deploying a distributed architecture can also reduce or smooth out spikes by allowing images to be retrieved from multiple locations instead of a single archive.

The choice between the centralized and distributed architectures depends on the type of PACS network and the flow of medical images that need to be supported. Prefetching algorithms are currently being used in PACS networks to regulate the traffic of images flowing between the storage archives and the workstation servers. The problem with prefetching algorithms is that if they fail to provide the medical images in advance, the radiologists must wait for the images to be retrieved from the archives. Most current radiology networks do not have high bandwidth connections to handle retrieval of large medical image files within a required response time. Therefore, the desirable network infrastructure should be able to support real-time image retrieval such that the radiologists can retrieve the medical images within a required response time of 2 seconds [CART03] [CHIM92].

4.2 THE DICOM AND TCP/IP PROTOCOLS

DICOM [DICO03] is an application layer network protocol that enables reliable and unambiguous communication and transmission of medical images, and other medical

information between various medical sources and users. In other words, it provides communication and image delivery between modalities, workstations, archives, and short-term storage servers. It supports a wide range of medical images across the fields of radiology, cardiology, pathology, and dentistry. It is also designed to integrate with the Health Information System (HIS) and the Radiology Information System (RIS). It features a standard common file format and facilitates standard network send, receive and query operations. A single DICOM file contains a header and image data. The header contains patient demographic information such as patient's name, age, type of scan, image dimensions, etc. The DICOM header is in tens of bytes and is very small compared to the image data, which can be several megabytes. DICOM runs on top of the Transmission Control Protocol/ Internet Protocol (TCP/IP) network communication protocol.

The TCP/IP stack [LEON00] is a set of protocols that allows communication across multiple diverse networks. TCP is a transport layer protocol that provides reliable connection-oriented transfer of stream information over the connectionless IP. It operates on a pair of end hosts across the network. It provides a mechanism for error recovery and flow control on an end-to-end basis to handle problems such as packet loss, delay or misdelivered packets. The flow control mechanism reduces the rate at which packets are transmitted into the network when congestion is detected. The IP protocol handles the transfer of information across multiple networks through the use of gateways and routers. It deals with the routing of packets across different networks. IP packets are routed independently in a connectionless setup, and so they may traverse different paths. The fact that IP is connectionless makes the network robust because if failures occur in the

network, the packets are routed around the points of failures. TCP is responsible for recovery if packets are discarded due to congestion problems. Since the DICOM protocol relies on TCP/IP, the performance of TCP and IP affects the performance of DICOM in transporting medical images.

4.3 DISCRETE-EVENT SIMULATION

Simulation modeling is becoming an increasingly popular method for network performance analysis [CHAN99]. There are two methods of network simulation; analytical modeling and discrete-event simulation [FISH01]. Analytical modeling characterizes a network as a set of mathematical equations. The main disadvantage is its over simplistic view of the network and its inability to simulate the dynamic nature of a network. Thus, study of a complex network requires discrete-event simulation, which allows one to build a more realistic simulation model of a real-life situation.

In a discrete-event simulation the state of the system changes at discrete points in time. Discrete-event simulations can be categorized as *event-driven* or *time-driven*. An event-driven simulation models the network's activities as a series of events that occur asynchronously and at irregular intervals. Time-based simulations synchronize all changes in the network to a single clock.

There are many simulation software tools and packages available today designed specifically for modeling communication networks. Chang [CHAN99] and Law *et al* [MLAW94] describe some of these tools. The simulation tool that is used in this thesis is OPNET, which is described in Appendix A.

4.4 NETWORK MODELS

Two network topologies, centralized and distributed, are modeled in this study to compare their performance for the future MAN interconnecting 9 hospitals in the city of Winnipeg. Figure 4.1 shows a centralized architecture where all the hospitals are connected to a central storage archive located at St. Boniface Hospital and a backup archive at HSC. The reason for placing an archive at St. Boniface hospital is because of its size and location. Figure 4.2 shows a distributed architecture with three storage archives located at Grace Hospital, HSC and St. Boniface Hospital and a backup archive at HSC. These three hospitals are currently generating more than 90% of the total radiology workload. This is the reason that they are selected to have the archives. For both configurations, the interconnection between any two hospitals is a 10 Gigabit Ethernet (10 Gb/s) link. The scale on the diagrams provides a sense of the geographical distance between the hospitals in kilometers (km). The propagation delay of each link is calculated based on the distance between the two hospital connected by the link.

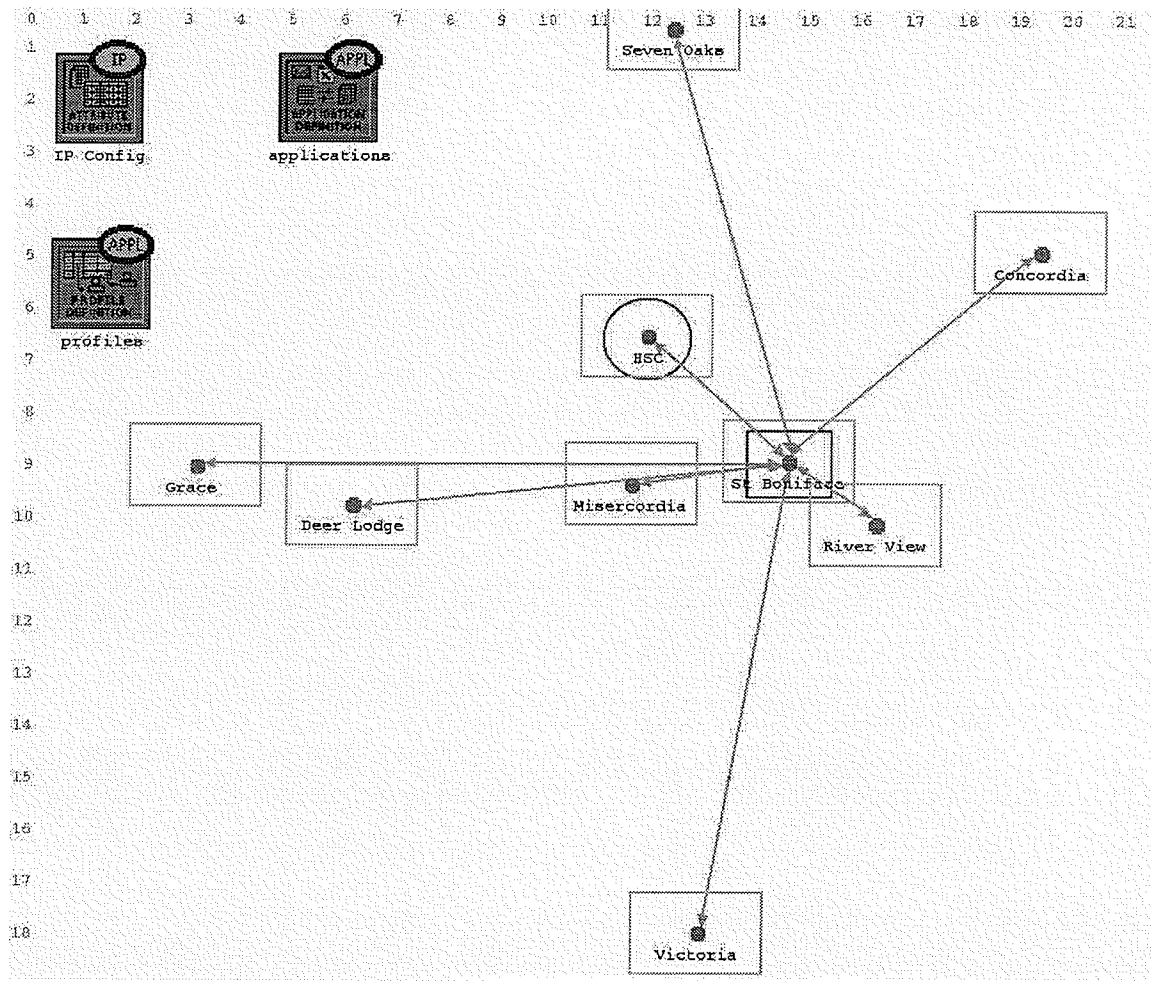


Figure 4.1: Centralized MAN architecture




Figure 4.2: Distributed MAN architecture

4.5 NETWORK MODELS COMPONENTS

Figure 4.3 shows the departmental LANs contained in a hospital. This is a typical configuration that is currently implemented by the radiology departments. Each hospital site has a PACS network which may consist of an archive, servers, and acquisition modules. The archive is only available at the hospitals mentioned in Section 4.4. The archives are for long-term storage of medical images. They are connected to the

acquisition devices which are the source of image creation. The servers belong to different radiology departments and are distributed over the hospital for short-term storage of images. These images are to be accessed by the viewing stations for diagnosis. The number of subnets varies from hospital to hospital and depends on the types of modality each hospital has. For example, Figure 4.3 shows a part of a PACS system with 5 subnets. The 5 subnets represent the 5 radiology departments and their corresponding modality types: ultrasound (US), computed tomography (CT), computed radiography (CR), nuclear medicine (NM), and digital fluoroscopy (Fluoro).

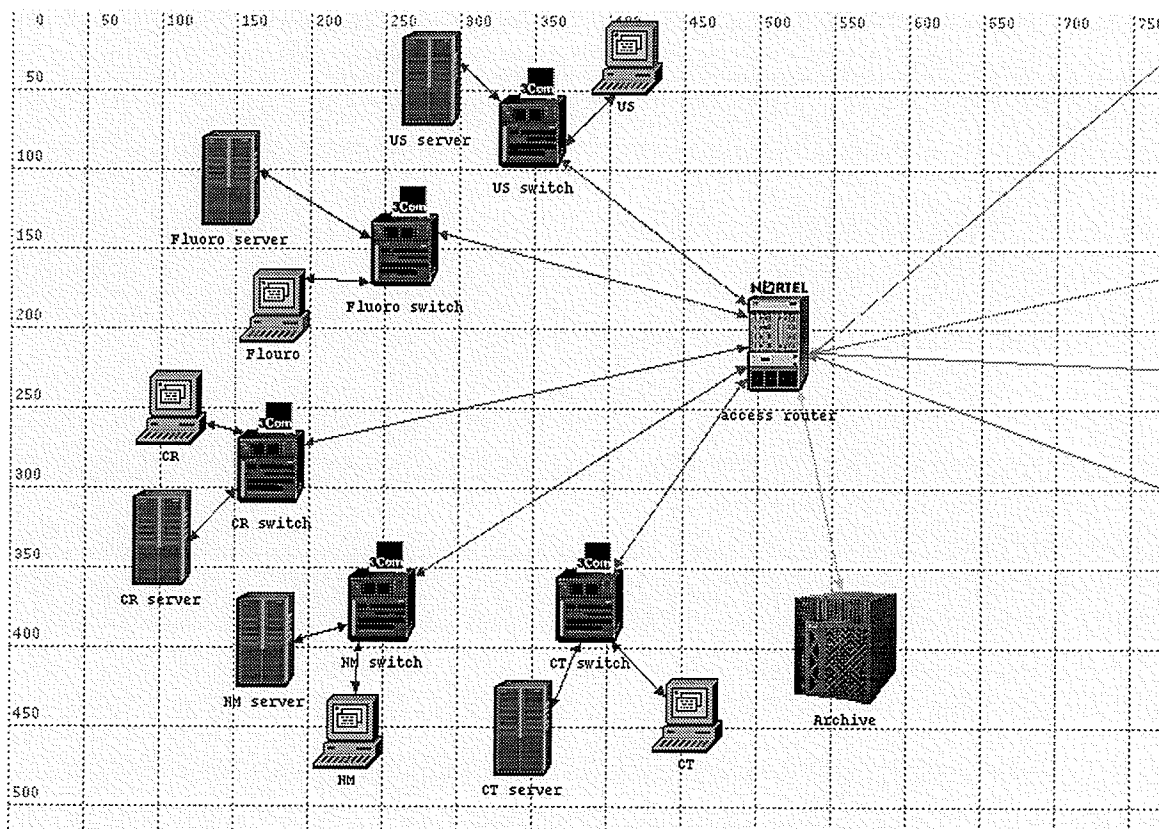


Figure 4.3: Hospital PACS architecture

Each subnet has a server and a traffic source. The traffic source produces the aggregate traffic that is generated by all the modalities of that type. For example, if a hospital has 5 NM devices, the traffic source will generate the equivalent amount of traffic that would have been generated by the 5 NM devices combined. The traffic source node consists of multiple objects, each object representing a single modality, that generate traffic streams simultaneously. Each object is assigned its own traffic parameters and operates independently. The advantage of using aggregated multiple traffic sources instead of using separate traffic sources is that the network model is simpler, and easy to construct and work with. An important challenge of modeling a network with many devices lies in the “capture” of the network model. That is, entering the model of the network into the modeling environment.

Each hospital has an access router connecting to other hospitals. The subnets are connected to this access router through the departmental switches. The subnet links in our model are 1 Gigabit Ethernet links. The long-term storage archives are connected to the access router through 10 Gigabit Ethernet links.

We used OPNET's existing network models for all the components in the network except for the acquisition stations. We constructed our own image acquisition node models using OPNET's node editor. Every network object (except for the links) has an underlying node model that specifies the internal flow of information in the object. Node models are made up of one or more modules connected by packet streams or statistic wires. Node modules in turn contain process models as described in Appendix A. A process model is represented by a state transition diagram (STD) that describes the behavior of a node module in terms of states and transitions.

4.5.1 THE LONG-TERM STORAGE ARCHIVE

The archives store newly acquired images, old examinations and patient data related to the images. The archives store medical images generated from different acquisition stations. They handle three types of image requests:

1. Adhoc retrievals: These are examinations retrieved at any time by any radiologist from any of the 9 hospitals with the necessary privileges.
2. Archive requests: These are requests made by the archive processes. The archive queues all the new incoming examinations for archiving, and sends a backup copy of the images to the backup archive.
3. Dearchive requests: These are retrievals of the new examinations and priors from the archive to different servers. Priors are the required old examinations that are attached to the new examinations and sent to the servers for diagnosis. Not all new examinations require priors. For example, Table 3.1 shows that 50% of the new examinations generated by all modalities at this hospital require priors. Unlike adhoc retrievals which could occur at any time, priors are tied to the archive requests and occur only when there is a new examination.

4.5.2 THE SHORT-TERM STORAGE SERVERS

The servers provide a temporary local storage for images retrieved from the PACS archive. These images are used by the radiologists and the referring physicians for reviewing. Thus images are stored at several servers to be available at different locations within each hospital. Previous examinations are also sent through the same path together

with the new examinations. A radiologist viewing images at a given viewing workstation can only access images available on the corresponding server. If the radiologist requires an image that is not available on the local server, an adhoc request is posted from the server to the archive to retrieve the image.

4.5.3 MODALITIES

The modalities are image acquisition devices which acquire digital diagnostic images. Figure 4.4 shows the layout of the image acquisition node. The primary component of the acquisition node is the dicom application module, which is responsible for generating traffic. The traffic generator for the dicom application module is implemented as a set of ON/OFF process models. The implementation of these process models is described in detail in Appendix B. Each of the image acquisition nodes in the network models contains the same state machines. The state machine transits between four states; opening a TCP connection, sending, receiving, and closing the connection. Different acquisition stations have different characteristics and features, and hence are configured with different attributes depending on the amount of workload they generate.

Each ON/OFF process model is supplied with parameters such as the size of the examination, the number of examinations to be generated and the IP address of the archive. The ON/OFF process starts by opening a TCP connection to the archive and waits for an acknowledgment that the connection is successful. Once the TCP connection has been established successfully, the ON/OFF process transits to the send state to generate traffic. The period in the send state varies from process to process and depends

on how large the examination to be generated is. It stays in this state until traffic equivalent to the size of the examination has been generated. If there is any traffic waiting to be received the ON/OFF process then transits to the receive state to process the data. Otherwise it transits to the off state and closes the TCP connection. The ON/OFF process keeps track of the number of examinations it has to generate. If the number of examinations parameter has been reached, the ON/OFF process stops generating traffic. Otherwise it stays in the off state until it goes back to the open state to open another TCP connection. The period of the off state is a random variable from a Normal distribution. We assume each process model generates scheduled examinations within 8 hours per day. We also assume that 90% of the examinations can be deviated and be completed within 9 hours per day. The off state parameters are explained in more detail in the next section.

The traffic generated at the dicom application module is sent to the lower layer protocols TCP and IP to be transported to the required destination.




Figure 4.4: Image acquisition node

4.6 TRAFFIC MODELING

Most of the radiology examinations are scheduled examinations. For the radiology assessment conducted in Chapter 3 we made the following assumptions:

1. All the examinations are scheduled within an 8-hour period day.
2. For each modality, the examinations durations are modeled as normally distributed random variables.

3. The mean and variance of the time taken by each examination to be acquired varies from modality to modality.
4. The scheduled examinations can be delayed due to emergency examinations.
5. Taking into consideration the emergency examinations, we assume that 90% of the time all scheduled examinations are performed within a 9-hour period of the day.

In our workload forecast analysis, we also made the following three assumptions:

1. The number of examinations performed each day increases annually.
2. The size of the examinations increases every year either due to the increase in the number of images per examination or increase in the size of images.
3. The rate of image retrieval increases annually.

The above assumptions are translated into the traffic model parameters to be configured in our simulations.

Since the examinations are scheduled and each modality requires the same amount of time to acquire each examination, we assume that the time between the examinations for each modality is normally distributed with mean μ and standard deviation σ where,

$$\mu = \frac{8 \text{ hours}}{N}, N = \text{exams / day} \text{ and } \sigma = \frac{0.7803}{\sqrt{N}} \text{ hours}$$

The standard deviation is derived from the equation:

$$P(X > 9 \text{ hours}) = 0.1, X = \sum_{j=1}^N Y_j \text{ where } Y_j \text{ is the time between the } (j-1)^{\text{th}} \text{ and the } j^{\text{th}}$$

examination (0th examination considered a start of day) assuming 90% of the time, all examinations are completed within 9 hours. The mean and the standard deviation

parameters are used by the ON/OFF traffic source to model traffic generated by the examinations acquired each day by each modality.

We did not have any historical data to analyze the trend in the annual increase in the examination size. As such we used the Mayo Hospital [MAYO01] estimations to calculate the annual growth rate of the examination size. The Mayo Hospital conducted a workload analysis to estimate the archive storage requirements for their radiology departments. They found out that the number of images per examination and the number of images per examination for CT, MRI, and US have increased substantially over the past few years. Table 4.1 shows the overall increase in the examinations performed and the total volume per year from the Mayo Hospital analysis.

	1998 Actual		2003 Estimate	
	Procedures	Data (GB)	Procedures	Data (GB)
Plain Film	407 913	1 400	425 000	12 500
CT	82 242	1 500	108 422	4 500
MRI	32 605	600	49 998	3 200
US	60 724	400	71 515	860
NM	57 902	50	64 023	290
Angio/Fluoro	51 264	0	55 934	1 400
Mammography	46 128	0	54 513	11 000

Table 4.1: Estimated volumes for the Mayo Hospital [MAYO01]

The growth rate in the examination size used in our performance analysis is calculated as follows:

$$\text{Annual Growth Rate} = \left(\frac{\frac{2003 \text{ Data}}{2003 \text{ Procedures}}}{\frac{1998 \text{ Data}}{1998 \text{ Procedures}}} \right)^{\frac{1}{5}} - 1$$

Table 4.2 shows the examination size annual growth rate used for each of the modality type.

Modality	Growth Rate (%)
US FA	13
US	13
Fluoro	1
US Ped	13
NM	39
Echo Ped	13
CT Ped	18
CT	18
Gen Angio	1
Cath Lab	1
X-Ray	54
Rad Ped	54
Ded Chest Rad	54
Fluoro Ped	1
MRI	28
CR	54

Table 4.2: Examination size growth rate

We used the data from Health Science Center (HSC) to represent the current image size for the retrievals. HSC contains almost all the modality types. It represents a typical hospital setup. Table 4.3 shows the file sizes for each year using the growth rates in Table 4.2.

	Current (MB)	Year 1 (MB)	Year 2 (MB)	Year 3 (MB)	Year 4 (MB)	Year 5 (MB)	Year 6 (MB)
US	18.9	21.3	24.1	27.2	30.8	34.8	39.3
US FA	13.1	14.8	16.7	18.9	21.4	24.1	27.3
US Ped	21.0	23.7	26.8	30.3	34.2	38.6	43.7
Echo	26.2	29.6	33.5	37.8	42.7	48.3	54.6
Echo Ped	13.1	14.8	16.7	18.9	21.4	24.1	27.3
NM	6.56	9.12	12.7	17.6	24.5	34.0	47.3
CT	163	192	226	267	315	372	439
CT Ped	46.1	54.4	64.2	75.8	89.5	106	125
Gen Angio	28.7	28.7	28.7	28.7	28.7	28.7	28.7
Cardio Angio	1.84	1.84	1.84	1.84	1.84	1.84	1.84
Cath Lab	14.7	14.7	14.7	14.7	14.7	14.7	14.7
X-ray	36.3	55.9	86.1	133	204	314	484
Rad Ped	14.3	22.1	34.0	52.3	80.6	124	191
Ded Chest Rad	3.05	4.70	7.23	11.1	17.2	26.4	40.7
Fluoro	1.31	1.31	1.31	1.31	1.31	1.31	1.31
Fluoro Ped	15.7	15.7	15.7	15.7	15.7	15.7	15.7
MRI	25.2	32.2	41.2	52.8	67.6	86.5	111
CR	31.5	48.4	74.6	115	177	272	420

Table 4.3: Image retrieval file sizes

The inflation rate for the number of retrievals was calculated by using the weighted average of the annual increase in the number of examinations performed daily by each hospital which is given in Table 3.1. The inflation rate was calculated as follows:

$$Inflation\ Rate = \frac{\sum_{hospital=1}^9 growth\ rate_{hospital} \times total\ exams_{hospital}}{\sum_{hospital=1}^9 total\ exams_{hospital}}$$

Table 4.4 shows the annual inflation rate in the number of examination retrievals for each modality calculated using the equation above.

Modality	Annual Inflation Rate (%)
US FA	5
US	0.77
Fluoro	3.21
US Ped	2
NM	5.11
Echo Ped	10
CT Ped	2
CT	3.28
Gen Angio	37.5
Cath Lab	5
X-Ray	3.6
Rad Ped	2
Ded Chest Rad	0
Fluoro Ped	2
MRI	29.23
CR	9.02

Table 4.4: Image retrieval inflation rate

Assuming an 8-hour period of day, the rate of retrieving new images is calculated as follows:

$$\text{Image Retrieval Rate} = \frac{\text{total examinations/day}}{8 \times 3600 \text{ seconds}}$$

The rate of retrieving old examinations is calculated as follows:


$$\text{Priors Retrieval Rate} = \frac{\text{total examinations/day} \times \% \text{ of required priors} \times \text{number of required priors}}{8 \times 3600 \text{ seconds}}$$

The average inter-request time for retrieving the new and prior examinations is calculated as follows:

$$\text{Inter - request Time} = \frac{1}{\text{Image Retrieval Rate} + \text{Priors Retrieval Rate}}$$

Table 4.5 shows the current year inter-request time in seconds for retrieving new and prior examinations. To forecast the inter-request time for subsequent years, we use the following equation:

$$\text{New Inter - request Time} = \frac{\text{Previous Year Inter - request Time}}{1 + \frac{\text{Inflation Rate}}{100}}$$



	Hsp 1	Hsp 2	Hsp 3	Hsp 4	Hsp 5	Hsp 6	Hsp 7	Hsp 8	Hsp 9
US	128	72	4.43	160	144		261.82	213.33	
US FA		309.68	480						
US Ped			640						
Echo		1200	72.18						
Echo Ped			720						
NM		171.07	185.33	259.46	1415.2				
CT	1028.6	110.77	184.62	720	576		1028.6	800	
CT Ped			1200						
Gen Angio		1515.8	303.16						
Cath Lab		514.29							
X-ray	288	18.41	58.3		261.82	2057.1			942.06
Rad Ped			90.14						
Ded Chest		480	130.91						
Fluoro	4800	685.71	266.67		2400	2742.9	768	2400	19200
Fluoro Ped			4800	960					
MRI		369.23	1230.8						
CR				53.33				192	

Table 4.5: Inter-request time in seconds

We assumed that adhoc retrievals are done over an 8-hour day and the number of retrievals is 60% of the number of new examinations. This is based on the Mayo statistics [MAYO01]. At Mayo, the number of retrievals was 165000 and the number of new examinations was 275000 for 1999. We considered the busy hour, in which we assume the mean activity is 1.5 times the mean activity for the day. This is just a reasonable

guess and we have no method of proving this. The adhoc retrieval rate is calculated as follows:

$$\text{Adhoc Retrieval Rate} = \frac{0.6 \times 1.5 \times \text{total examinations/day}}{8 \times 3600 \text{ seconds}}$$


Similarly, the inter-request time for adhoc retrievals is calculated as follows:

$$\text{Adhoc Inter - request Time} = \frac{1}{\text{Adhoc Retrieval Rate}}$$

Table 4.6 shows the current year inter-request time in seconds for adhoc retrievals. For subsequent years, the adhoc inter-request time is found as follows:

$$\text{New Adhoc Inter - request Time} = \frac{\text{Previous Year Adhoc Inter - request Time}}{1 + \frac{\text{Inflation Rate}}{100}}$$

For the network model with three archives, we assumed that the archives have an equal adhoc retrieval load. Therefore the retrieval rate is split uniformly among the three archives.



	Hsp 1	Hsp 2	Hsp 3	Hsp 4	Hsp 5	Hsp 6	Hsp 7	Hsp 8	Hsp 9
US	711.11	400	24.62	888.89	800		1454.6	1185.2	
US FA		516.13	800						
US Ped			1066.7						
Echo		2000	120.30						
Echo Ped			1280						
NM		351.65	380.95	533.33	2909.1				
CT	2285.7	246.15	410.26	1600	1280		2285.7	1777.8	
CT Ped			2666.7						
Gen Angio		3200	640						
Cath Lab		800							
X-ray	640	40.92	129.55		581.82	4571.4			2093.5
Rad Ped			150.23						
Ded Chest		1066.7	290.91						
Fluoro	8000	1142.9	444.44		4000	4571.4	1280	4000	32000
Fluoro Ped			8000	1600					
MRI		533.33	1777.8						
CR				88.89				320	

Table 4.6: Adhoc inter-request time in seconds

4.7 IMAGE TRAFFIC SOURCE MODELS

The image transfer in the MARN is based on a client-server model. There are three types of image traffic flowing in the network: traffic flowing from the modalities to the archives, traffic flowing from the archive to a backup archive and traffic between the

archive and the servers. New images are generated at the image acquisition stations and sent to the archives. The archives send backup copies of the new examinations to the backup archive. The new images are also sent, along with the attached prior examinations, to the servers. We assume that each archive has a relational database functionality to determine at which of the servers the images must be stored. The servers can also query the archives for adhoc images. The traffic between image acquisition stations and the archives, and traffic between the archives and backup archive is modeled using traffic source models explained in previous sections.

The transfer of images from the archives to the servers is modeled as renewal processes. Renewal processes are used to model independent identically distributed occurrences. Renewal processes are processes whose events occur over time. The time to an occurrence follows a certain distribution. The application-layer protocol used for transferring medical images is the DICOM protocol discussed in Section 4.2. DICOM addresses image file formats and different commands at the application layer to access relational databases and to enable the fetching and storing of images. Since OPNET does not have built-in DICOM protocol models, File Transfer Protocol (FTP) application models are used for convenience. The FTP is a protocol that is used to transfer files over the Internet. Petri et al. [PETI97] performed some experiments comparing transmission times using FTP and DICOM to transfer images between their hospitals. Their results indicated similarities between the two protocols' transmission times, with FTP faster in some cases and DICOM better in other conditions. They also indicated that there was no great advantage in using either one of the protocols in terms of transmission times.

The events corresponding to the beginning of transmissions constitute a renewal process with exponentially distributed inter-arrival times, and the transmission durations are determined by the deterministic examination sizes and the network conditions. Each modality has a different examination size. Hence different FTP applications are defined for different modalities as a way of having different traffic source models. Each application has different parameters. The two important parameters that define the traffic generation for each application are the file size and the examination inter-request time. Adhoc retrievals are configured as "get" FTP applications. The servers initiate the request to retrieve examinations from the archives. New examinations and prior retrievals are configured as "put" FTP applications forwarded by the archives to the servers. The archives initiate the request to send the images to the servers.

Some approximations described in Section 4.6 were made in the modeling to simplify matters. The workload data for HSC was used to represent the average image size for the applications. Table 4.3 shows the file size used for each modality. The inter-request time used for the adhoc, new and prior applications are calculated in Section 4.6.

We could have modeled the examination transfers as ON/OFF models instead of renewal processes. The deterministic ON period is very small such that the ON/OFF model could be compared to the renewal process. It worked well in our case to use both models.

4.8 PARAMETERS CONFIGURED

This section discusses how the parameters are configured in the traffic source models using OPNET to generate traffic.

4.8.1 FTP APPLICATION MODELS

Figure 4.5 shows the applications configured for the retrieval of examinations from the archives. Each hospital in the model has its own applications, and each modality type has its own application. The applications are either retrievals of new and prior examinations or adhoc retrievals. All the examination retrieval applications are configured as FTP applications. A TCP connection is opened for each examination transfer.

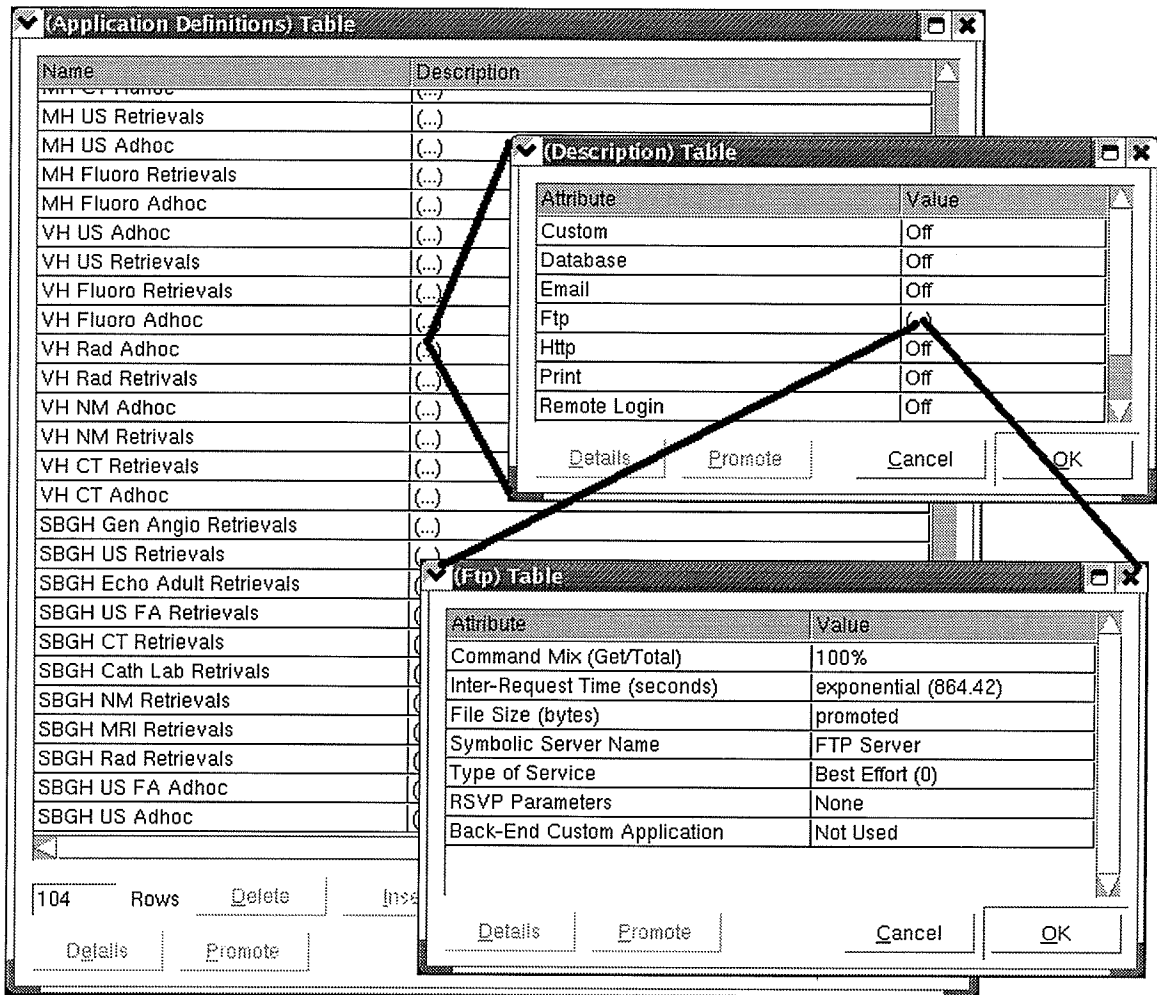


Figure 4.5: Examination retrieval applications

The following attributes are configured for each application:

- **Command get (get/total):** This ratio specifies the ratio of "get" commands to the total number of commands. The examination retrievals are either "get" commands or "put" commands and not both. All adhoc retrievals have this attribute set to 100 % (only get commands) and all prior and new examinations have this attribute set to 0 % (only put commands).

- **Inter-request Time:** This is the model for the time between subsequent examination retrievals.
- **File Size (bytes):** The data size of the examination being transferred. The data includes images from new and prior examinations. The size of each image is deterministic and depends on the modality type.

The inter-request time and the file size attributes can be random variables with distributions for which one can specify parameters. Section 4.6 discusses how the file size and the inter-request time parameters are determined.

4.8.2 PROFILES

The applications defined for different modalities are grouped into profiles that are then specified on the servers to generate application layer traffic. Figure 4.6 shows a list of different profiles configured for different applications. Each profile describe how the servers generate network traffic by specifying the applications to be executed, for how long and how often the applications are used throughout the day. If a profile consists of multiple applications, they are configured to run simultaneously until the end of the simulation. For each application in a profile, usage parameters such as start time, duration of the application and repeatability are specified.

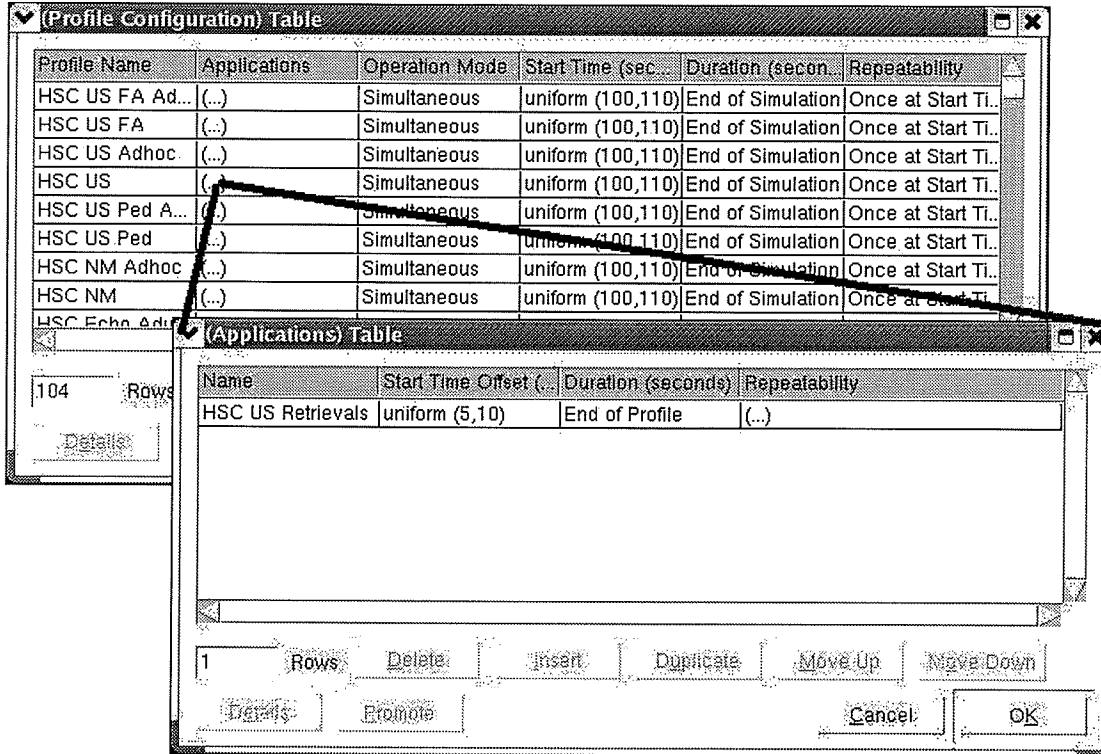


Figure 4.6: Examination retrievals profiles

4.8.3 ARCHIVE PARAMETERS

Figure 4.7 shows how an archive is configured to support the different applications. Each application represents the examination retrieval for each modality in each hospital.

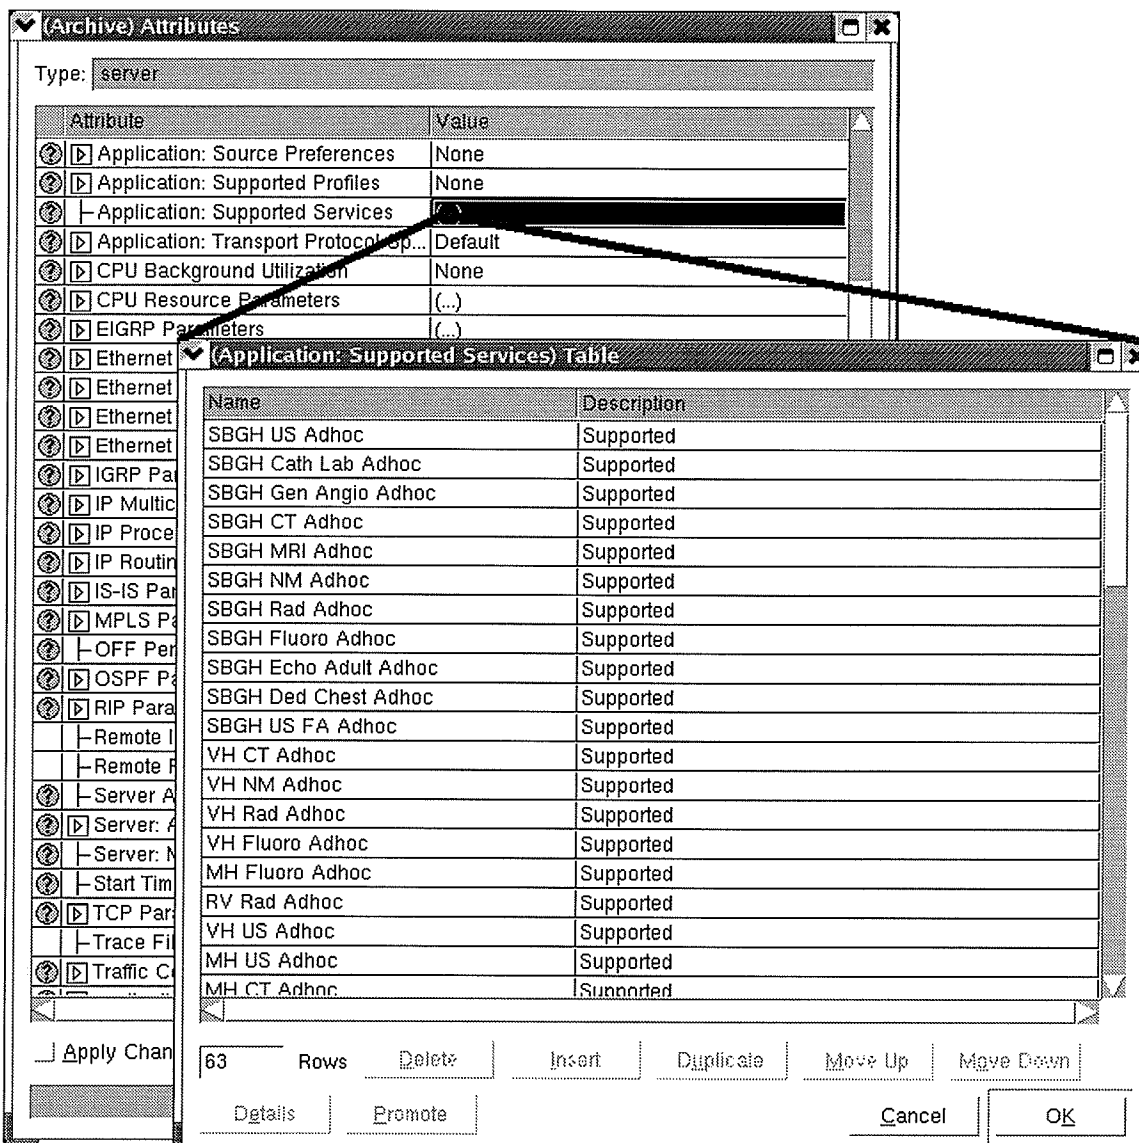


Figure 4.7: Configuring archives to support applications

4.8.4 SERVER PARAMETERS

A server is configured to support profiles that match with its subnet. Each server supports a specific modality type. Figure 4.8 shows the configuration of the US server at GH. From the figure, the server supports only profiles related to US retrieval. The GH US

profile represents retrieval of new and prior ultrasound (US) examinations. The GH US Adhoc profile represents adhoc retrieval of US examinations.

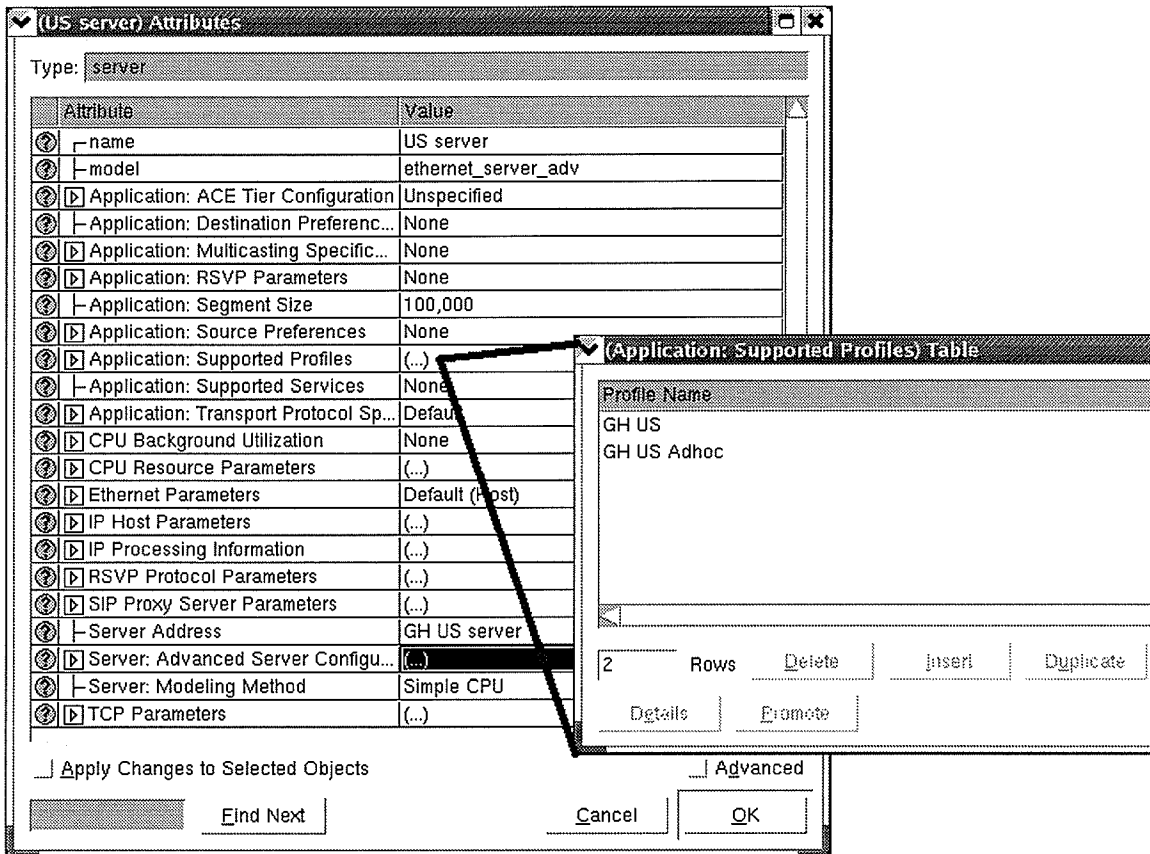


Figure 4.8: Configuring servers to support applications

4.8.5 MODALITY PARAMETERS

Modalities are configured to generate new examinations and send them to the archives for storage. The modality traffic parameters shown in Figure 4.9 include the OFF Period Args. This parameter is being used to calculate the interarrival times between examinations. Other parameters include the remote IP address of the archive, the TCP

remote port and the traffic configurations. The traffic configurations specify the amount of traffic to be generated by the modality.

The OFF Period Args parameter has two values separated by a semicolon. The first value denotes the length of the simulated time period. The second value is used to calculate the standard deviation for the OFF period value. This value is calculated in Section 4.6. The Traffic Configurations parameter contains different traffic attributes for the ON/OFF models to be multiplexed. The number of row represents the number of ON/OFF models to be created. Each row contains parameters for one ON/OFF model. Each ON/OFF model generates traffic of a single machine. The No. of Images attribute and the Image Size attribute are not used in the simulation. The No. of Exams specifies the number of examinations to be generated by the model. The File Size attribute is the size of the examination in bytes that will be generated by the ON/OFF model. Finally, the Local Port attribute is the host port where a TCP connection for each ON/OFF model will be established.

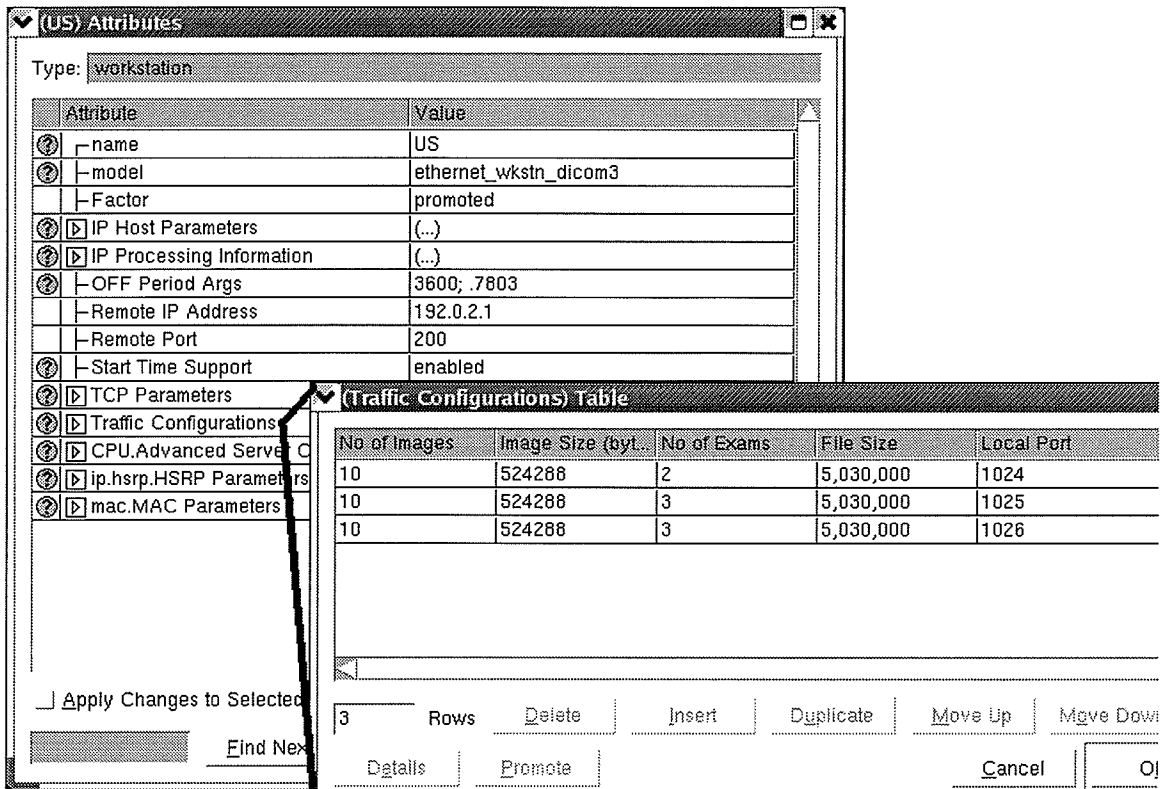


Figure 4.9: Modality traffic parameters

The archives are also configured with the same parameters in order to send the new examinations to the backup archive according to a statistically similar process. Each archive specifies the IP address of the remote backup archive.

4.8.6 IP PROCESSING AND CPU SETTINGS

The modalities, the archives and servers are configured with a datagram forwarding rate of 1 Gbits/second. The datagram forwarding rate is the number of bytes that are processed by the forwarding processor residing in a node in one second. This rate is set to the bottleneck link of each link.

The modalities and the servers each have a single Sun Ultra 10 processor running at a speed of 333 Mhz. The archives are configured to have three processors. The processing speed of each archive or server can be varied by using a multiplication factor parameter. This parameter is set to 100.

4.9 ASSUMPTIONS

In our performance analysis study, we assume a PACS network with a topology shown in Figure 4.3. The network has three network components; image acquisition devices, short-term storage servers for local access by the viewing stations and long-term storage archives. We assume the long-term storage archives are located at Grace, St Boniface and HSC. The connectivity between the hospital PACS networks, or the metro links, is 10 Gigabit Ethernet.

It is assumed that each server has enough disk space to store all the requested examinations. As shown in the Section 4.6, the request rate for each server depends on the number of examinations generated each day by each modality type and the growth estimations derived from the Mayo Hospital study. A connection request from either the server or the archive requires images to be sent from the archive to the server.

The archives are configured to operate as fast as the storage devices. Moreover, we assume that there is enough buffer storage at each archive to handle multiple image retrieval and storage requests without any delay. We assume the archives have the functionality of a relational database and know to which server to forward the examinations. It is also assumed that the requested images or the forwarded images are

always available in the archive buffer and do not need to be fetched from the long-term storage media.

The amount of data traffic, such as DICOM header files, which contain patient demographics information, text and control data is insignificant, approximately 0.001%, compared to the image traffic. As such, traffic of data other than images is not considered in the analysis.

4.10 LAYER 2 AND LAYER 3 SWITCHING

Switching is one way of improving the bandwidth utilization and network performance. Switching makes segmentation on the LAN possible to provide dedicated bandwidth. Consequently, it reduces congestion on the network and can isolate network problems.

4.10.1 LAYER 2 SWITCHING

Layer 2 switching, also called bridging, operates at the data link or Media Access Control (MAC) layer. Layer 2 switches forward packets based on the unique MAC address of each end station. At Layer 2, no modification on the routing information is required when switching between Ethernet and Gigabit Ethernet. However, slight modifications on the routing information may occur when bridging between different interfaces.

The departmental switch shown in Figure 4.3 is the 3Com Superstack II 3300 switch which provides the layer 2 functionality in our intra-hospital network models. The 3300 switch delivers the performance and flexibility of cost-effective Gigabit Ethernet switching to boost performance to desktop users and critical server links [3COM00].

4.10.2 ROUTING

Routers are used to connect two networks at the network layer. Layer 3 of the Open System Interconnection (OSI) model is the network layer which provides switching and routing technologies for transmitting data across different networks. Routers generally route different network layer protocols such as IP, IPX, APPN, etc. Routers offer various LAN, MAN or WAN interfaces such as Ethernet, Token Ring, etc. The major drawback with traditional routers is that most of the routing decisions are processed in software, which adds higher latency and more complex administration. This can cause the LAN routers to become the performance bottleneck [TELCO1].

4.10.3 LAYER 3 SWITCHING

Layer 3 switching refers to a class of high performance routers optimized for LAN. Layer 3 switching blends the attributes of high-speed Layer 2 switching with intelligent network services of Layer 3 routing [ROBE00]. Layer 3 switches work at the network layer (Layer 3 of the OSI reference model) and incorporate the same technology and intelligence features as that of traditional routers. They have advantages over traditional routers because they have optimized hardware that allows fast packet processing. It allows Layer 3 switches to pass data as fast as Layer 2 switches. Hence the performance bottlenecks that can occur when using traditional software-based routers can be eliminated when using faster hardware-based Layer 3 switches.

Providing the boost in performance using Layer 3 switched enables switched networks to efficiently utilize bandwidth, become more stable and reliable. Layer 3

switches act as gateways, exchanging Routing Information Protocols (i.e., RIP, OSPF, etc.) with other routers at wire speed. Layer 3 switches are often called IP switches because they are primarily used to route IP.

4.10.4 ROUTING PROTOCOLS: RIP

The routing protocol used in our network models is Routing Information Protocol (RIP). RIP is based on the Bellman-Ford distance vector algorithm. Each router periodically transmits its routing table to its neighbors. Upon receiving the information, the neighboring routers recalculate their routing tables and then retransmit their own routing tables. RIP is widely available and easy to use. RIP's benefits are:

- RIP is easy to configure and manage.
- RIP has a very low overhead in terms of configuration and management, relative to more sophisticated protocols.
- RIP has a low bandwidth overhead in small networks.

4.10.5 NORTEL PASSPORT 8600 ROUTING SWITCH

The Nortel Passport 8600 router, shown in Figure 4.3 as the access router, is used in our inter-hospital network models to provide the Layer 3 routing capabilities using RIP protocol. The Passport 8600 router is a high-end switch designed for backbone enterprise networks [NORTEL]. It supports manual route configuration, as well as dynamic calculation through RIPv1, RIPv2 or OSPF. The key benefits of the Passport 8600 router include:

- Seamless integration between LAN, MAN and WAN.
- High availability and scalability.
- Layer 3 wire speed routing and filtering.

4.11 TCP PERFORMANCE

TCP has a significant impact on the performance of high speed networks. To take full advantage of the high bandwidths and obtain a good TCP throughput across a metropolitan area network requires tuning some TCP parameters. Each node in OPNET with TCP as one of its protocols provides an attribute called TCP Parameters to specify TCP configurations. The following techniques from [TCPTG] summarize how to get the maximum TCP MAN throughput:

1. Maximize the congestion window size. TCP uses a congestion window to determine how many packets can be sent at one time. The TCP "slow-start" and "congestion avoidance" algorithms determine the size of the congestion window. The two algorithms make TCP adapt to a variety of network speeds automatically. The maximum congestion window is related to the amount of the buffer space that the kernel allocates for each socket. For each socket, there is a default value for the buffer size, which can be changed. The buffer size can be adjusted for both the send and receive ends of the socket. To get the maximal throughput it is critical to use optimal TCP send and receive socket buffer sizes for the link used. If the buffers are too small, the TCP congestion window will never fully open up. If the receive buffers are too large, TCP flow breaks and the sender can overrun the receiver, which will cause the

TCP window to shut down. This is likely to happen if the sending host is faster than the receiving host. An overly large window on the sending side is not a big problem as long as there is excess memory.

2. Use Selective Acknowledgements (SACK). SACK allows a receiver to acknowledge non-consecutive data. SACK enables the fast recovery algorithm and the fast retransmit algorithm to handle multiple packets dropped per window [RFC1323].
3. Use an optimal value for the maximum segment size (MSS). This value is set to the maximum size that the MAC layer can handle, minus the TCP and IP headers. If the MSS is small, the network may be under utilized. Conversely, a large MSS could lead to many IP segments that cannot be acknowledged or retransmitted independently. This can degrade the network performance.

4.11.1 BANDWIDTH*DELAY PRODUCT (BDP)

When transmitting data, there is a delay between sending a segment and receiving it. This is due to the propagation delay and some processing overhead on both ends [ALEX96]. The amount of traffic that can be in transit in the network is called the Bandwidth-Delay-Product (BDP), because it is the product of the bottleneck link bandwidth and the Round Trip Time (RTT). The importance of BDP is that it is the amount of buffering that is required in the end hosts. The optimal buffer size is twice the BDP of the link.

4.11.2 TCP PARAMETERS

Figure 4.10 shows the TCP parameters configured for all the network nodes. The values are set as follows:

1. The maximum segment size (MSS) is set to the default value. Setting the MSS to the default value automatically guarantees that it is configured optimally for Ethernet.
2. The buffer size is set to 1 GB, the maximum allowed buffer size. To get optimal performance, this value should be at least the BDP.
3. The Slow-start initial count is set to "As defined in RFC 2414". This value specifies the number of MSS-sized TCP segments that will be sent upon slow start. RFC 2414 proposes the initial window size to increase from one segment to roughly two to four segments. With an initial window of at least two segments, the receiver will generate an ACK after the second segment arrives. This eliminates the wait on the timeout.
4. The window scaling, fast retransmit, fast recovery and SACK mechanisms are enabled. For more information on these algorithms refer to RFC 2001. TCP uses the fast retransmit algorithm to detect and repair losses based on duplicate ACKs. The fast recovery algorithm governs that transmission of new data until a non-duplicate ACK arrives. The fast recovery algorithm is based on TCP Reno.

Attribute	Value
Maximum Segment Size (bytes)	Auto-Assigned
Receive Buffer (bytes)	1000000000
Receive Buffer Adjustment	None
Receive Buffer Usage Threshold (of ...)	0.0
Delayed ACK Mechanism	Clock Based
Maximum ACK Delay (sec)	0.001
Slow-Start Initial Count (MSS)	As defined in RFC-2414
Fast Retransmit	Enabled
Duplicate ACK Threshold	2
Fast Recovery	Reno
Window Scaling	Enabled
Selective ACK (SACK)	Enabled
ECN Capability	Disabled
Segment Send Threshold	Byte Boundary
Active Connection Threshold	Unlimited
Nagle Algorithm	Disabled
Karn's Algorithm	Enabled
Timestamp	Disabled
Initial Sequence Number	Auto Compute
Retransmission Thresholds	Time Based
Initial RTO (sec)	3.0
Minimum RTO (sec)	1.0
Maximum RTO (sec)	64
RTT Gain	0.125
Deviation Gain	0.25
RTT Deviation Coefficient	4.0
Timer Granularity (sec)	0.5
Persistence Timeout (sec)	1.0
Connection Information	Do Not Print

Details Promote Cancel OK

Figure 4.10: TCP parameters

4.12 SUMMARY

This chapter has described in detail the network models used in our performance analysis. The chapter gave an introduction to the distributed and centralized PACS network topologies and gave the advantages and disadvantages of the two topologies. The traffic models used were also described in details. The ON/OFF process models are used by the modalities to generated new examinations. The renewal process application models are used by the archive and server for the retrieval of adhoc examinations, and new and old examinations. The assumptions made were also explained. The parameters configured in the network components were explained in detail. The next chapter gives the simulation results.

CHAPTER FIVE

SIMULATION RESULTS AND ANALYSIS

In this chapter, performance evaluation of 10 Gigabit Ethernet is conducted to determine if 10 Gigabit bandwidth meets the requirement of an average response time of 2 seconds in retrieving medical images in the next five years. The comparative results for the distributed MAN topology and centralized MAN topology are presented. This chapter is organized as follows. Section 5.1 describes the response time performance measure. In Section 5.2, results of the distributed MAN topology and centralized MAN topology are presented. Section 5.3 discusses the results.

5.1 PERFORMANCE MEASURES

In high speed metropolitan area radiology networks, congestion problems may arise, mainly because of the digital medical image transfers. To address such problems, the performance measures to evaluate in the simulations are the 90th percentile and the 99th percentile of the delay, the average image retrieval response time and the maximum image retrieval response time. The p^{th} percentile is a value so that the probability of the delay being smaller than that value is p . The image retrieval response time is the time to retrieve examinations from different possible storage locations in the network. The

performance measure for the image retrievals is characterized by the rate at which image requests are served and the response time for image requests. Two types of image retrieval response times are considered. The first response time is the adhoc examinations retrievals, which is the time elapsed between an examination request and receiving the examination. It is measured from the time a server sends a request to an archive to the time the server receives an examination from the archive. The second response time is the new examinations and priors retrievals response time. It is measured from the time an archive sends a new examination together with the required priors to a server to the time the archive receives an acknowledgment from the server.

5.2 RESULTS

A series of discrete-event simulations were run to determine the variation of the image retrieval response time over a period of 5 years. All the simulations were run for 3600 seconds of simulated runtime. The seed value for each simulation was changed in order to determine the standard deviation of the response time. Ten simulations were run for each year.

Figure 5.1 and Figure 5.2 show how the average response times increase over time due to the increased traffic volume from one year to the next when retrieving images in the centralized topology and the distributed topology. Figure 5.1 shows the response time in retrieving adhoc examinations. The figure shows no significant difference in the average response time between the two topologies in the next 4 years. Both the topologies exceed the required average response time of 2 seconds in year 6. In the

distributed topology the average adhoc response time is 1.48 seconds in year 5 and 2.08 seconds in year 6. In the centralized topology the average adhoc response time is 1.51 seconds in year 5 and 2.27 seconds in year 6.




Figure 5.1: Response time for retrieving adhoc examinations

Figure 5.2 shows the average response time in retrieving new and prior examinations. Both the topologies meet the required response time of 2 seconds in year 6. In the distributed topology the average retrieval response time is 1.18 seconds in year 5 and

1.79 seconds in year 6. In the centralized topology the average adhoc response time is 1.21 seconds in year 5 and 1.85 seconds in year 6.




Figure 5.2: Response time for retrieving new and prior examinations

The observed utilization in the metro links for the distributed network configuration was 80% in year 5 and 94% in year 6. The utilization is the percentage of the consumption to date of an available link bandwidth. In the centralized network configuration, the utilization in the metro links was 93% in year 5 and 99% in year 6. These numbers show

that the centralized network start to experience congestion at year 5 and the distributed network at year 6.

Figures 5.3 and 5.4 show the complementary cumulative distribution function for the adhoc retrievals both for the distributed and centralized topologies. Figure 5.5 shows the complementary cumulative distribution function for the new and priors retrievals for the centralized topology. The shapes of these curves show the right tail that decreases at a non-exponential rate. Table 5.1 shows the 90th, 99th, 99.99th percentiles for the adhoc response times and the maximum adhoc response times. The table shows that 90% of the time the distributed topology meets the average response time of 2 seconds and 90% of the time the centralized topology response times are less than 2.399 seconds, which is close to the required average response time. In year 5, 99.99% of the time the response time is less than 3.226 seconds in the distributed topology and less than 4.423 seconds in the centralized topology. The maximum delays are less than 5 seconds for both topologies. These delays are tolerable when dealing with emergency cases. Table 5.2 shows the 90th, 99th, 99.99th percentiles for the new and prior examinations response times and their maximum response times. The table shows that 90% of the time both the topologies meet the required average response time. The maximum delays are 4.852 seconds and 5.604 seconds for the distributed and centralized topologies respectively.




Figure 5.3: Probability distribution for adhoc response times for distributed topology




Figure 5.4: Probability distribution for adhoc response times for centralized topology




Figure 5.5: Probability distribution for new and prior response times for centralized topology

Year	Distributed				Centralized			
	90 th	99 th	99.99 th	Max	90 th	99 th	99.99 th	Max
1	0.349	1.543	1.543	2.089	0.347	1.542	1.544	2.353
3	0.818	2.136	2.616	2.727	0.952	2.136	2.137	2.455
5	1.928	2.975	3.226	4.896	2.399	3.808	4.423	4.653

Table 5.1: Percentiles for adhoc retrievals

Year	Distributed				Centralized			
	90 th	99 th	99.99 th	Max	90 th	99 th	99.99 th	Max
1	0.346	1.542	2.533	2.613	0.347	1.542	1.817	2.393
3	0.816	2.136	2.189	2.424	0.816	2.136	3.257	3.453
5	1.928	3.377	4.692	4.852	1.930	3.433	5.600	5.604

Table 5.2: Percentiles for new and priors' retrievals

5.3 DISCUSSION

The results presented in Section 5.2 show that 10 Gigabit bandwidth is sufficient to sustain the radiology workload for a period of 5 years in the Winnipeg metropolitan area radiology network. Though we used the Mayo statistics [MAYO01] to estimate the trend in the technology advancements, we are confident that the Winnipeg radiology departments will follow similar trends by buying more advanced equipment. The results further indicate that either the distributed or the centralized topology can handle the workload to meet the requirement of 2 seconds mean response time in retrieving medical images for diagnosis.

Although high speed bandwidth technologies such as 10 Gigabit Ethernet give alternatives of implementing metropolitan area radiology networks as distributed or centralized, the choice between these two topologies depends on the size of the network and whether there is potential for future expansion. There are other factors that must be

considered which are not covered in this thesis. These factors are reliability and scalability. A centralized architecture is not easily expandable to large scale enterprise PACS. Distributed architectures ensure the capability for future infrastructure growth, connectivity and reliability.

CHAPTER SIX

CONCLUSIONS

In this chapter, we summarize the contributions of this thesis and discuss possibilities for future work. This thesis has characterized the radiology workload at the Winnipeg hospitals for the purpose of identifying the bandwidth requirement needed in the metropolitan area radiology network. The results showed that 10 Gigabit Ethernet will meet the performance required to handle the Winnipeg radiology workload in the next 5 years.

6.1 CONTRIBUTIONS

The contributions of this thesis are summarized as follows:

1. The workload characterization of the radiology departments in Winnipeg hospitals has been conducted. Interviews and surveys were conducted at the hospitals to gather traffic information on the acquisition devices that generate medical images in the radiology departments. The total aggregated volume of image data currently generated each day is 311 GB.
2. The workload was further projected to predict the variation of the network workloads in the future. The purpose of the hospital workload forecast was to assess future

bandwidth that is required to sustain the workload within the required performance. Assumptions and estimations used in the workload forecast were explained. The analysis showed that the estimated workload would be roughly five times the current workload in five years.

3. The current workload and future estimated workload were used to forecast the bandwidth requirement. The forecast bandwidth is expected to deliver medical images at the viewing stations within a response time of 2 seconds. To determine the future bandwidth requirements, some assumptions were made, which either increased or decreased the bandwidth requirement. Using these assumptions, according to some rough calculations the estimated bandwidth for the MAN links was determined to be 10 Gigabit connections in the next 5 years.
4. The future bandwidth requirements were validated through simulations. Discrete event simulations in OPNET were developed to investigate the performance of the metropolitan area radiology network connecting the Winnipeg hospitals. The performance analysis was performed for a forecast period of 5 years.
5. The workload models have been implemented as traffic generators to be used in the simulation. The modality models were implemented as ON/OFF process models in OPNET. The ON/OFF models have alternating ON and OFF periods. Traffic corresponding to a single examination is generated during the ON period. No traffic is generated while the traffic source is on the OFF period. The traffic source models run on top of TCP/IP. The retrieval of examinations at the archives and the servers was modeled as renewal processes at the application layer.

6. A comparative literature review on the distributed and centralized PACS topologies has been carried out. The advantages and disadvantages of each topology were discussed. The network models used in our performance analysis study have been explained. The assumptions made were also described. The parameters configured in the network models were described in detail.
7. The performance of 10 Gigabit Ethernet on the radiology MAN connecting the Winnipeg hospitals has been evaluated. The results showed that a bandwidth of 10 Gigabit will be able to sustain the radiology workload in the next 5 years, using either the distributed topology or the centralized topology. The choice between the centralized and distributed topology depends on the size of the MAN and whether future extension in the network will be required. Knowing the amount of the radiology workload that would flow in the MAN helps to determine the kind of MAN to implement.

6.2 RECOMMENDATIONS FOR FUTURE WORK

Radiology metropolitan area networks are the future basis for better medical diagnosis.

The following issues can be studied in the future:

1. Workload analysis is very important in forecasting bandwidth requirements. Workload analysis should include the use of historical data to estimate future workload volumes. In this thesis, we did not have any historical data except for the projected increase in the current workload. To get more accurate results in forecasting the radiology workloads, historical data should be used to estimate future workloads based on real

- past trends. However this cannot account for changes due to technology advancements.
2. This thesis can be extended to focus not only on radiology only but the entire hospital. Integrating radiology with other non-radiology departments allows clinicians to benefit from a more efficient access to the images. Also, the integration of the modalities and PACS with the RIS and HIS help improve the workflow in the health enterprise network. Although radiology data is large it accounts for most of the traffic.
 3. The radiology LAN in each hospital can be extended to use wireless LAN technology with handheld mobile computers. Wireless mobile networks allow healthcare specialists to access medical applications from virtually anywhere.
 4. Only MAN data network is supported in this thesis. This can be extended to a converged MAN network capable of transporting data, video and voice traffic. This includes any multimedia applications that can streamline delivery of healthcare, such as remote diagnostics and real-time remote consultation on medical images.
 5. The radiology MAN can be extended to include other cities in the province.
 6. The analysis conducted in this thesis can be extended to include remote clinical access using Virtual Private Network (VPN). VPN access provides a secure and encrypted connection between major hospitals and remote clinics. The VPN terminals can be placed at the hospitals with the storage archives.
 7. As the volume of medical data increases over time, storage on dedicated disks directly attached to servers can prove to be inefficient. The Winnipeg radiology workload forecast in Chapter 2 projected that the medical data volume would increase at a rate corresponding to approximately five times the current workload in the next five years.

To handle this traffic, the hospitals can implement Storage Area Networks (SAN).

SAN provides faster data access and simplified image storage management.

8. This thesis considered DICOM traffic which is user-driven traffic and did not consider any block-level input/output (I/O) traffic. Assuming there is a SAN (for block storage) on the radiology MAN network; I/O traffic characteristics should be considered. I/O traffic in the SAN is due to caching, mirroring of virtual volumes from one storage server to another and replication for remote disaster recovery and backup. I/O traffic can be a burden on the network and cause bandwidth limitations.
9. Simulation and modeling plays a very important role in the performance analysis of high speed networks. Hardware-in-the-loop techniques can be applied to simulation and modeling of radiology Metropolitan Area Networks. Hardware-in-the-loop simulation (HILS) refers to a technology where some of the components of a pure simulation are replaced with actual hardware. For example, the process models developed using OPNET for the image acquisition devices and the application models for the short-term storage servers can be replaced by the actual hardware. OPNET can then be used to create process models to interface between the virtual network models and the real hardware. The advantage of HILS is that the real behavior of medical image manipulation is simulated without assumptions or approximations.

APPENDIX A

OPNET PROCESS MODELS

This appendix provides an introduction to OPNET simulation tool and how process models are implemented using OPNET. Section A.1 describes OPNET and Section A.2 explains the implementation of process models.

A.1 OPNET SIMULATION TOOL

OPNET, acronym for **OP**timum **NET**work performance, is a commercial network simulation tool [OPNET3] capable of simulating large communication networks with detailed protocol modeling and performance analysis. OPNET allows the definition of a network topology, nodes and links that make up a network. It features a graphical specification of models, a dynamic event-driven simulation kernel, integrated data analysis tools, and hierarchical object-oriented modeling. The simulation kernel provides a procedure library that simplifies various problems related to the subject of data communication such as manipulation of data packets. The various data analysis tools allow collection of output data, data analysis and the graphical representation of the data.

OPNET uses a hierarchical structure of distinct models to represent the different components of a communication network. Figure A.1 shows the three layers in the

hierarchy of an OPNET simulation. Each model has an associated modeling domain and editor, allowing the structure of models to closely resemble that of the real-world networks. The modeling domains are:

1. **Network layer:** This is a high level specification of the topology of a communication network in terms of subnets, nodes and links between them. Every object in this layer has a set of attributes such as name, position, etc.
2. **Node layer:** This specifies the basic building blocks (modules) in terms of applications, processors, queues and communication interfaces. Nodes are created in the Node Editor by interconnecting a number of process modules using packet streams or statistic wires which, when combined together, model the functions of a node. The node model specifies objects in the network domain.
3. **Process layer:** This specifies the behavior of processes such as algorithms and decision making that operates within the nodes in the network. The states in the process contain user-defined code that details this behavior. The process model specifies objects in the node domain.

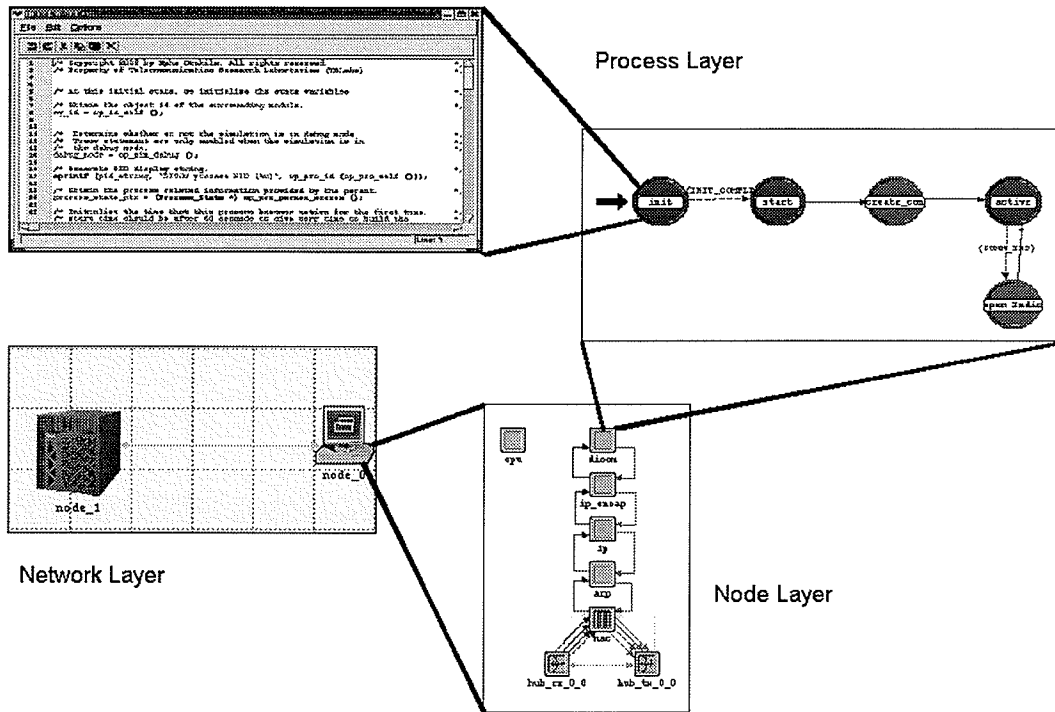


Figure A.1: The three-tiered OPNET hierarchy

A.2 PROCESS MODELS

Process models are user-programmable objects implemented to realize behaviors that represent a real world protocol or algorithm. OPNET expresses process models in a language called Proto-C based on a combination of finite state machines (FSMs) described by state transition diagrams (STDs), a library of high-level procedures known as Kernel Procedures (KP) and the C programming language. A process model's STD defines a set of primary nodes or *states* that the process can enter and, for each state, the conditions that would cause the process to move to another state. The condition needed for a particular change in state to occur and the associated destination state are called a

transition. States are mutually exclusive and complementary, meaning that a process is always in exactly one state; more than one state may never be occupied at a time. Actions may be associated with each state and they are called *executives*. There are two types of executives; enter executives and exit executives. A state's enter executives are executed when a process enters the state. Its exit executives are executed when the process transition to another state. The states can be of two types - *forced* states and *unforced* states. In a forced state, the process retains control of the simulation until it has completed all of its tasks. An unforced state can be interrupted by another event at any time. Unforced states allow a pause between the enter executives and exit executives, and thus can model true states of a system. Forced states are displayed in green color and unforced states in red color. Figure A.2 has one forced state, the *init* state, indicated by a line and two unforced states, the *crt_child* state and the *end* state.

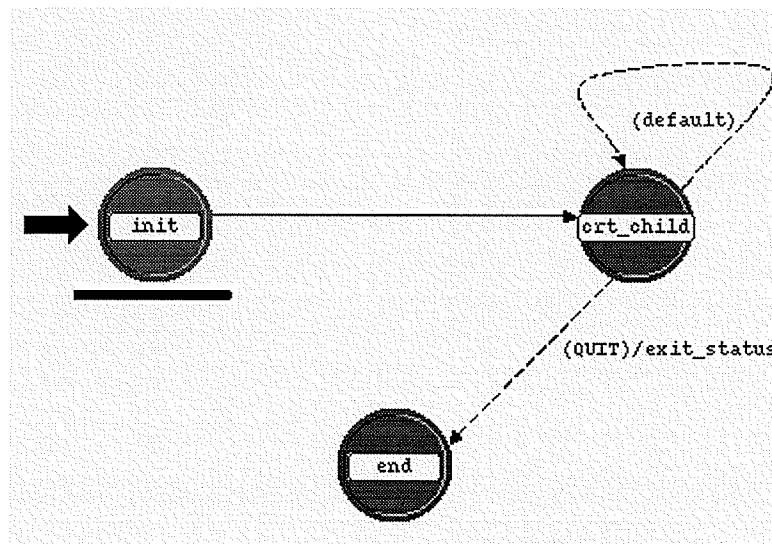


Figure A.2: Example of a state transition diagram

OPNET supports the concept of *process groups*. A process group consists of multiple processes that execute within the same processor or queue. When a simulation begins, each module has only one process, called the *root* process. This process can later create new processes and/or destroy them. When a process creates another one, it is called the *parent* process of the new process. The new process is called the *child* of the process that created it. Process groups are useful to manage tasks that are generated on a dynamic basis, i.e. their exact timing is not known. Such processes are referred to as *dynamic processes*. This thesis uses the technique of process groups to multiplex traffic from a certain number of ON/OFF traffic sources to represent traffic generated from a certain type of modality. For example if a hospital has 10 CT modalities, 10 dynamic processes are created and their traffic will be multiplexed. This feature allows us to minimize the number of nodes in the network model, thus making the network model simple and easy to work with.

One special state, called the *initial* state, must be designated in each process model. Graphically, an initial state is identified by a large arrow located at its immediate left side as shown by the *init* state in Figure A.2. The initial state is the point at which execution begins on the first invocation of the process. Since this state contains executive statements performing initializations that occur only once, most process models do not include transitions that lead back to it.

Transitions describe the possible movement of a process from one state to another and the conditions under which such state change may take place. A transition specification has four components: a source state, a destination state, a condition expression, and an executive expression. Each state may have any number of incoming

and outgoing transitions, which are depicted as directed arcs with an arrow pointing towards the destination state. There are two types of transitions, a conditional transition represented by a dashed line (see Figure A.2) and an unconditional transition represented by a solid line. A conditional transition has a condition, which is evaluated as a Boolean expression to decide whether or not the process should enter the transition's destination state. The condition and the executive appear in a combined label next to the arc, separated by a forward slash ('/'). An unconditional transition is a transition whose condition always evaluates true. Hence a state with an outgoing unconditional state has only one departing transition. A special condition called *default* is provided to represent the complement of the conditions associated with transitions leaving a state. This ensures that the process will not encounter an error at the state of interest due to a missing TRUE transition. A default condition succeeds only if all other transition conditions are false.

APPENDIX B

HOSPITAL WORKLOAD MODELS

This appendix describes in detail the implementation of the workload models of traffic generated by the hospital modalities. Section B1 presents the implementation of traffic source generators for the hospital workflow that was discussed in Chapter 2. The hospital workloads are implemented as ON/OFF process models in OPNET. The traffic generators are implemented by multiplexing ON/OFF traffic source models with alternating ON and OFF periods running on top of the TCP/IP protocol. Server application process model implementation is also described in Section B.2.

B.1 IMPLEMENTATION OF TRAFFIC SOURCE GENERATOR

The traffic generated by modalities is modeled as ON/OFF models shown in Figure B.3, which alternate between active and idle states. In the active state, a TCP connection is opened and an examination is generated. The size of the examination depends on the type of the modality modeled. Once the examination has been generated and transmitted, the TCP connection is closed. The period of the ON state is deterministic and the period of the OFF state is derived from the Normal distribution as discussed in Chapter 4.

B.1.1 ON/OFF TRAFFIC SOURCE MULTIPLEXER PROCESS MODEL

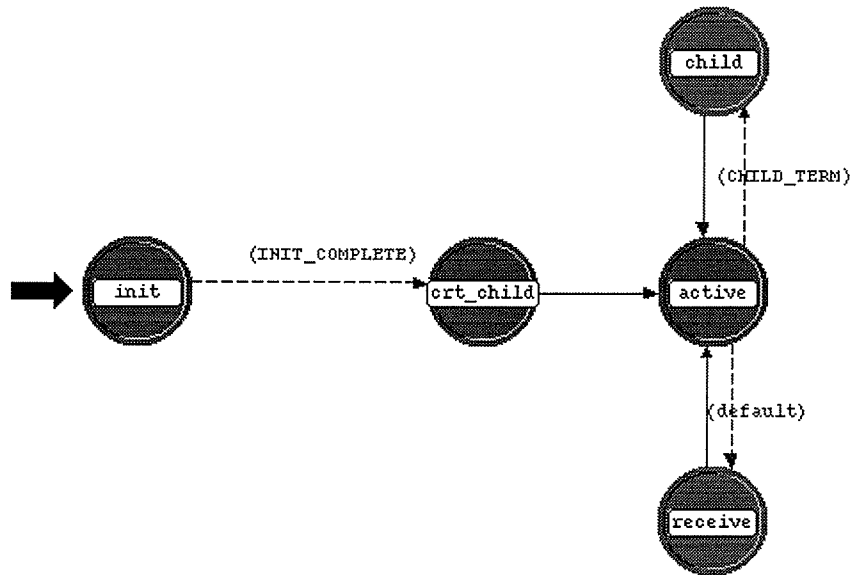


Figure B.1: State transition diagram of the process model for creating an ON/OFF traffic source multiplexer

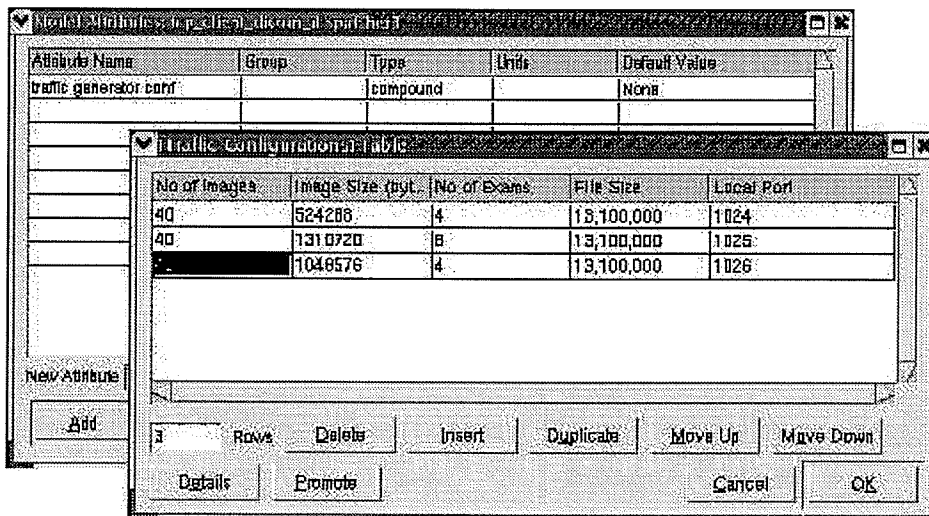


Figure B.2: Compound attribute for the ON/OFF traffic source multiplexer

State Descriptions

State	Description
init	The root process starts at this state and registers itself with the TCP module by using the function <i>tcp_app_register ()</i> . This returns a handle that the root process and its children use to communicate with the TCP layer. Hence the child processes created by this process do not need to re-register. A self interrupt is scheduled to let the process switch to the crt_child state.
crt_child	This state is entered once the registration with TCP is complete. The process reads the user-supplied model compound attribute shown in Figure B.2 and stores the values in the corresponding state variables. It then determines the number of ON/OFF traffic sources to be multiplexed by reading the number of rows in the expanded compound attribute. Each row contains information about each ON/OFF process to be created. The information contains the number of studies to be generated, the size of each study and the local port where the TCP connection will be opened. At this point the root process creates the ON/OFF processes and passes each process's information through the shared memory. The root process also passes the same TCP handle to the ON/OFF processes so the processes can interact with the TCP module directly. It then puts the created child processes into a list for easy tracking and manageability.
active	After creating its child processes, the multiplexer enters the active state and waits for interrupts, either from the ON/OFF processes or the TCP module below. The TCP module interacts with the ON/OFF processes through the root process only. The processes use the same TCP handle and so TCP would not know which process to forward the interrupts. The root process accesses the ON/OFF processes' connection identification variables to determine their TCP connection identifications. If the interrupt received is from the TCP it enters the receive state. If the interrupt is from any of the ON/OFF processes it enters the child state.
receive	In this state the process determines which ON/OFF process the interrupt belongs to by comparing the connection identification in the interrupt and each process's connection identification variable. It then forwards the interrupt to the appropriate ON/OFF process.
child	The ON/OFF processes send interrupts to the multiplexer when they have finished generating studies and terminated. The multiplexer determines the terminated process from the interrupt and removes it from the list and frees the memory the process used.

Table B.1: State descriptions of the ON/OFF traffic source multiplexer

Transition Condition Descriptions

Transition Condition	Description
INIT_COMPLETE	Initialization is complete. Switch from the init state to the crf_child state.
default	If there is no true transition always go to the receive state.
CHILD_TERM	A child process has terminated. Free its memory and remove it from the list.

Table B.2: Transition condition descriptions of the ON/OFF traffic source multiplexer

B.1.2 ON/OFF TRAFFIC SOURCE PROCESS MODEL

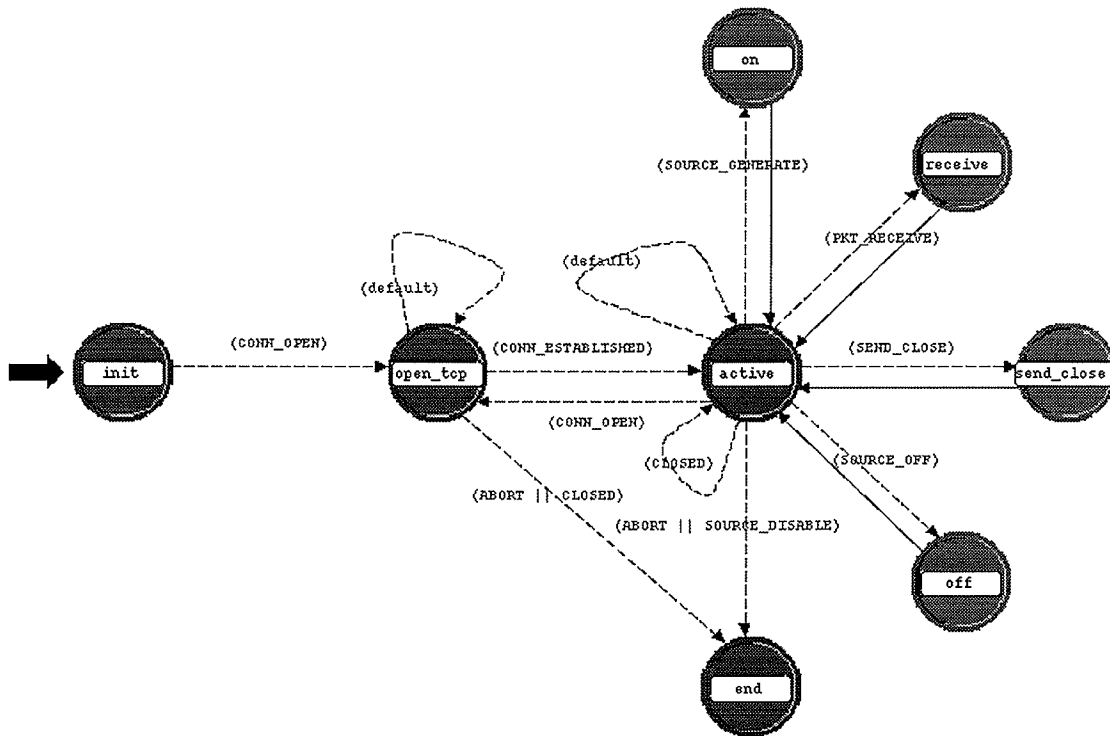


Figure B.3: State transition diagram of the renewal process model for creating an ON/OFF traffic source

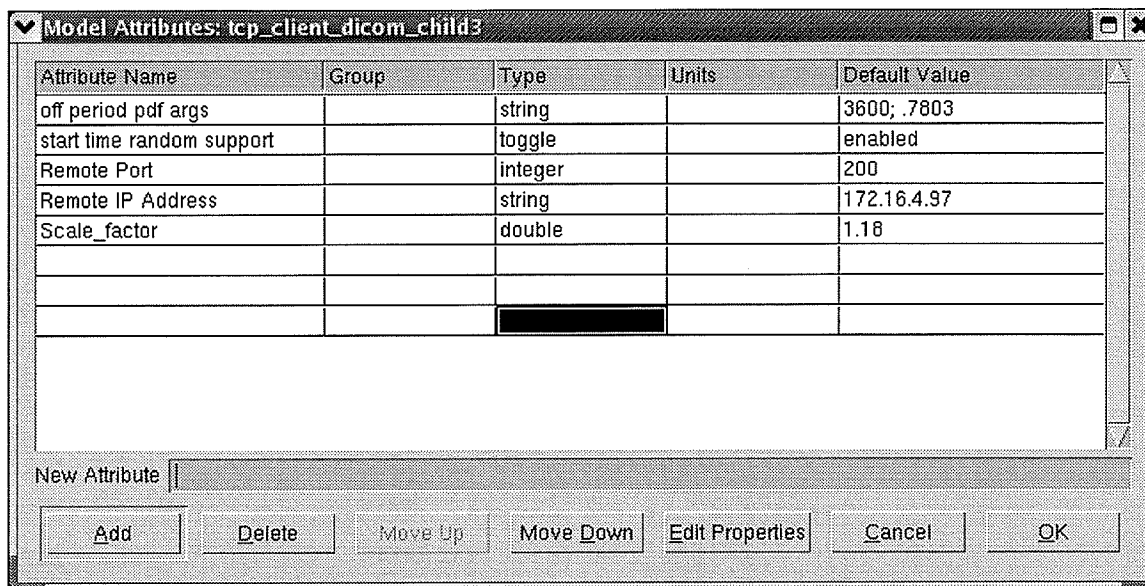


Figure B.4: Model attributes of the ON/OFF traffic source

State Descriptions

State	Description
init	The process starts at this state, reads and saves the user-supplied model attribute values shown in Figure B.4 into the corresponding state variables. The process also accesses the information passed by the root process through the shared memory and saves the values into state variables for quicker access. A self interrupt is scheduled at the <i>start time</i> to activate the process for the very first time. Since this process generates traffic that will be sent from one hospital to another in a different subnet, the simulation start time is configured to be greater than 60 seconds to give routing protocols, OSPF and RIP, time to build the IP routing tables. If the " <i>random start time support</i> " is disabled, the start time is set to 100 seconds. If the " <i>random start time support</i> " is enabled, the start time is set to 100 seconds plus a random outcome from the Uniform distribution between 0 and 200 seconds. This feature allows the created ON/OFF traffic sources to start at different times for the sake of multiplexing. A counter is initialized to keep track of the number of studies generated by this process. The statistics that this process will collect are registered and maintained under a single data structure. This process also loads the distribution that is used to calculate the off period and stores the value into the <i>off_dist_ptr</i> attribute. The distribution is a Normal distribution with its parameters calculated

State	Description
	from the " <i>off period pdf_args</i> " model attribute.
open_tcp	Once the process becomes active, it enters this state. The state opens a TCP connection by issuing an OPEN command using the function <i>tcp_connection_open ()</i> and waits for the response from the TCP. The process specifies the IP address of the remote server, the remote port and the local port of the connection. Upon receiving a connection successful signal, the process switches to the active state. Otherwise the process switches to the end state to terminate.
active	In this state, the process first checks the type of the TCP interrupt that caused the transition to this state. If the remote interrupt is a connection established interrupt, that means a TCP connection was established and the process is now ready to generate traffic. It schedules a self interrupt to go to the on state immediately to generate traffic. It also issues a RECEIVE command using the function <i>tcp_receive_command_send ()</i> to the TCP to receive any packets destined for this process. If the remote interrupt is a closed connection interrupt, it means a TCP connection has successfully closed. The process schedules a self interrupt to go to the off state. The active state is also the default state meaning that if the process is not doing anything, it will switch to this state. While in this state if the process receives an abort connection remote interrupt, that means that a TCP connection has closed prematurely. The process switches to the end state to terminate.
on	In this state the process generates a single study traffic. The size of study depends on each modality's image size and the number of images in each study. An unformatted packet is created with a constant packet size specified by the " <i>Scale_factor</i> " and " <i>File Size</i> " attributes. The scale factor models the increase in the study size per year. The file size is the initial size of study or study size of the current year. After creating the packet, it stamps the packet with the creation time in order to calculate the end-to-end delay statistic later on. It then sends the packet to an output stream going to the TCP layer by issuing a SEND command using the function <i>tcp_send_data ()</i> . It also updates the statistics for this process. It schedules to switch to the send_close state to close the TCP connection.
receive	This state is entered whenever there is a packet from the TCP layer. The process retrieves the packet from the incoming stream and destroys the packet.
send_close	This state is entered when a TCP connection is to be closed. The process creates a 1-byte close indication packet that tells the remote application to close the connection. It sends the packet to the remote application and then informs the TCP layer below that it does not have any more data to send. It issues a CLOSE command to the TCP using the function <i>tcp_connection_close ()</i> .
off	This state is entered when a TCP connection has closed successfully. In this state the process is idle and no packets are generated. The process

State	Description
	increments the total number of studies generated so far. If this number is equal to the "No of Exams" attribute then the process has finished generating the required studies. The process schedules a self interrupt to switch to the end state to terminate. If this number is less than the "No of Exams" attribute the process schedules a self interrupt in a random period of time to switch back to the open_tcp state. The length of this random period of time is an outcome of the Normal distribution specified by the <i>off_dist_ptr</i> attribute.
end	In this state the process cleans up the used memory and then destroys itself.

Table B.3: State descriptions of the ON/OFF traffic source

Transition Condition Descriptions

Transition Condition	Description
CONN_OPEN	It is time to switch to the open_tcp state and open a TCP connection.
Default	If there is no true transition it will always go to the active state. If it is in the open_tcp state, it will stay in the same state.
CONN_ESTABLISHED	A TCP connection has been established successfully. It is time to generate a study at the on state.
SOURCE_GENERATE	Switch to the on state and generate a packet.
PKT_RECEIVE	Switch to the receive state and receive an incoming packet.
SEND_CLOSE	Switch to the send_close state and send a CLOSE command to close a TCP connection.
SOURCE_OFF	The process is idle, switch to the off state.
CLOSED	A TCP connection has been closed successfully. Schedule an interrupt to switch to the off state. If the connection was closed during the setup, switch to the end state and terminate the process.
ABORT	A TCP connection has been aborted. Switch to the end state to terminate the process.
SOURCE_DISABLE	The process has finished generating packets. Switch to the end state and terminate the process.

Table B.4: Transition condition descriptions of the ON/OFF traffic source

B.1.3 DICOM SERVER DISPATCHER PROCESS MODEL




Figure B.5: State transition diagram of the server dispatcher process model for creating multiple server applications

State Descriptions

State	Description
init	The process starts at this state initializing the Application Programming Interface (API) package that is used to interface with the TCP module. A self interrupt is scheduled to let the process move on to performing its own initialization in the start state.
start	After receiving a handle to interact with TCP, the process enters this state and creates a list to hold its application process records. A self interrupt is scheduled to start receiving incoming TCP connection

State	Description
	requests. This also allocates time to all modules, particularly the TCP manager, to process their begin simulation interrupts. Before exiting the state the process opens a passive port to listen to incoming TCP connection requests and creates an application process, shown in Figure 3.8 below to this port. The application record is then inserted into the created list.
active	In this state the process waits for any incoming interrupts. If the interrupt is an interrupt requesting a TCP connection, the process switches to the open_indication state. If the interrupt is a process interrupt from any of the active application processes, the process switches to the child state. If the interrupt is a self interrupt, the process switches to the passive state. If the interrupt is a stream interrupt, the process switches to the dispatch state.
open_indication	In this state the process handles an indication from the TCP layer about a remote client that is attempting to establish a connection with the server process. On receiving an OPEN indication from a remote client process, the process verifies if it should accept the request or deny it. Once the process has indicated that the call can be accepted by TCP, the server process schedules a self interrupt to open another passive port.
passive	In this state the process starts another listening application process on the passive port. This feature allows the server process to have an unlimited number of concurrent TCP connections to different client processes.
dispatch	This is the default state. Whenever the server process receives a packet, it enters this state to determine the TCP connection to which the packet should be dispatched. The process compares the connection identification in the received packet with that of the application process records in the list. If there is a match, the process forwards the packet to the appropriate application process. Otherwise the packet is destroyed since there is no valid application process destination.
child	An interrupt from an application process is to indicate to the server process that it has terminated. Upon receiving this interrupt, the server frees the application memory, cleans up the application record and removes it from the child list.

Table B.5: State descriptions of the server dispatcher process model

Transition Condition Descriptions

Transition Condition	Description
INIT_COMPLETE	Initialization of lower layer has completed. Go to the start state to start initialization.
CONN_ID	TCP connect request has been received. Switch to the open_indication to decide what to do with the request.
NEW_PASSIVE	Open a new TCP listening port
default	Always go to the dispatch state if there is no interrupt. Switch to this state also if there is an incoming packet.
CHILD_TERM	An application has terminated. Clear its record from the list.

Table B.6: Transition condition descriptions of the server dispatcher process model

B.1.4 DICOM SERVER APPLICATION PROCESS MODEL




Figure B.6: State transition diagram of the process model for creating a server application

State Descriptions

State	Description
init	The process starts at this state, accesses the parent shared memory and gets the TCP handle passed by the parent. The end-to-end statistic maintained by this process is also registered.
idle	After initialization, the process switches to this default state and waits for interrupts. Remote interrupts indicate a change in the connection status. If the remote interrupt received is a <i>TCPC_IND_ESTAB</i> , it means that a TCP connection has become established with the remote client. The server application issues a RECEIVE command to the TCP layer using the function <i>tcp_receive_command_send ()</i> specifying that the server is ready to receive packets. If the remote interrupt is a <i>TCPC_IND_ABORTED</i> , it means that the connection has been terminated. The process switches to the done state. If the interrupt is a stream interrupt the process switches to the arrival state.
arrival	In this state the application receives the packet from the incoming stream. If the packet is a close request from the remote client, the application acknowledges with a close indication to the remote client. The process also issues a CLOSE command to the TCP layer using the function <i>tcp_connection_close ()</i> to inform the TCP that it does not have any packets to send or receive. If the packet is not a close request then it is an arrived study. The process accesses its creation time to calculate the end-to-end delay and updates its statistics.
done	A TCP connection has closed successfully or prematurely. The process terminated and a process interrupt is sent to inform the server dispatcher process.

Table B.7: State descriptions of the server application process model

Transition Condition Descriptions

Transition Condition	Description
default	Always stay in the idle state if doing nothing
ARRIVAL	A packet has arrived. Switch to the arrival state to receive the packet.
CONN_ABORTED	The TCP connection has been aborted. Switch to the done state to destroy the process.
exit_requested/CLOSE_CONN	A close connection request has been received. Execute the CLOSE_CONN procedure to close the connection and switch to the done state.

Table B.8: Transition condition descriptions of the server application process model

REFERENCES

- [10GEA] 10 Gigabit Ethernet Alliance, <http://www.10gea.com>.
- [10GE02] 10 GEA White Paper, "10 Gigabit Ethernet Technology Overview White Paper", May 2002, http://www.10gea.org/10GEAWhitePaper_0502.pdf.
- [3COM00] 3300 Superstack II Switch, <http://www.3com.com>.
- [ALEX96] Alex Habermann, "TCP in high speed networks", online paper, http://www.tkn.tu-berlin.de/curricula/ss96/bla/tcp_hs.html.
- [BLAR03] Blair Henderson, "Hospital Usage of Network", Internal Paper, St Boniface General Hospital Research Center Technical Report, Winnipeg, Manitoba Canada, 2003.
- [CARM01] Camorlinga S, Barker K, "Health Distributed Archive System (HDAS)", Internal Paper, St. Boniface General Hospital Research Center Technical Report Number MITR-200105A, Winnipeg, Manitoba Canada, May 2001.
- [CART03] Blake McClarty, "Diagnostic Imaging- Storage/Network Requirements", 2012 Medical Network Requirements Workshop, Winnipeg, Manitoba, Canada, March 13 2003.

References

- [CEDE98] De P., Cerratt T. W., "*An Interorganizational Information System in the Healthcare Arena: Insights Gained from a Hierarchical Analysis*", Proceedings of the 1998 ACM SIGCPR Conference on Computer Personal Research, March 1998, pp. 214-223.
- [CENT00] Centeno M., Albacete C., Terzano D., Carrillo M., Ogazon T., "A *Simulation Study of the Radiology Department at JMH*", Proceedings of the 2000 Winter Simulation Conference, December 2000, pp.1978–1984.
- [CGGR97] Cinotti M., Garroppo R. G., Giordano S., Russo F., "*Analysis and Efficient Modeling of the Traffic Generated by Broadband Teleradiology Applications*", Global Telecommunications Conference, GlobeCom' 97, IEEE, Vol. 2, November 1997, pp. 753-757.
- [CHAN99] Xinjie Chang, "*Network Simulations with OPNET*", Proceedings of the 1999 Winter Simulation Conference, Vol. 1, December 1999, pp 307–314.
- [CHIM92] Chimiak W., "*The Digital Radiology Environment*", IEEE Journal on Selected Areas in Communications, Vol. 10, No. 7, September 1992, pp. 1133–1144.
- [DICO03] The DICOM Standard, <http://medical.nema.org>.

References

- [FISH01] George S. Fishman, "*Discrete-Event Simulation: Modeling, Programming and Analysis*", Springer Series in Operations Research, Springer-Verlag, 2001.
- [FOUNET] Foundrynet Case Study, "*University Health Network*" online paper
<http://www.foundrynet.com/leadership/custBase/casestudies/PDFs/UHN.pdf>
- [GEOR89] Kositpaiboon R., Tsingotjidis P., Barbosa L, Georganas N., "*Packetized Radiographic Image Transfers over Local Area Networks for Diagnosis and Conferencing*", IEEE Journal on Selected Area in Communication, Vol. 7, No. 5, June 1989, pp. 842-856.
- [GUNS00] Trace Gunsch, "*A Case for Gigabit Ethernet for U.S. Army Network Backbones*", Technical Report TR00/009, Technology Integration Center, March 2000.
- [HALL92] Hall Charles, et al., "*Teleradiology Using Switched Dialup Networks*", IEEE Journal on Selected Areas in Communications, Vol. 10, No. 7, September 1992, pp. 1161-1172.

References

- [KLAU00] Klaus Krauter, "*PACS: A systems Approach*", presented at the Winnipeg Health Science Center, December 2000
<http://www.umanitoba.ca/faculties/medicine/units/radiology/presentations>.
- [LAWR85] Lawrence G., Marini G., Navon S., "*Simulation of a Hospital Picture Archiving and Control System (PACS)*", Proceedings of the 1985 Winter Simulation Conference, December 2003, pp. 543–550.
- [LEON00] Alberto Leon-Garcia, Indra Widjaja, "*Communication Networks Fundamental Concepts and Key Architectures*", McGraw-Hill, 2000.
- [MART98] Martinez R., Baxter C., Rivera A., "*United States Army South (USARSO) Network Simulation and Analysis Using OPNET*", 1998 OPNETWORK Conference, online paper, <http://www.opnet.com>.
- [MART93] Martinez R., Sanders W., Alsafadi Y., Nam J., "*Performance Evaluation of Picture Archiving and Communication Network using Stochastic Activity Networks*", IEEE Transactions on Medical Imaging, Vol. 12, No. 1, March 1993, pp. 12-29.
- [MART90] Nemat M., Martinez R., Osada M., Tawara K., Komatsu K., "*A high speed integrated computer network for picture archiving and communication system (PACS)*", Computer-Based Medical Systems,

References

- 1990., Proceedings of Third Annual IEEE Symposium on , June 1990, pp. 14-23.
- [MAYO01] B. J. Erickson, K. R. Persons, N. J. Hangiandeou, E. M. James, C. J. Hanna, D. G. Gehring, "*Requirements for an Enterprise Digital Image Archive*", Journal of Digital Imaging, Vol. 14, No. 2, 2001, pp. 72-82.
- [MENA02] Daniel A. Menasce, Virgilio A. F. Almeida, "*Capacity Planning for Web Services: Metrics, Models, and Methods*", Prentice Hall PTR, 2002.
- [MLAW94] Averill M. Law, Michael G. McComas, "*Simulation Software for Communications Networks: The State of the Art*", IEEE Communications Magazine, Vol. 32, No. 3, March 1994, pp 44-50.
- [NORTEL] Nortel Networks, <http://www.nortel.com>.
- [OPNET1] Modeling Concepts Reference Manual, Modeling overview, OPNET Manual, <http://www.opnet.com>.
- [OPNET2] Modeling Concepts Reference Manual, Process Domain, OPNET Manual, <http://www.opnet.com>.
- [OPNET3] OPNET Technologies Inc., <http://www.opnet.com>.

References

- [OPNET4] OPNET Model Descriptions Reference Manual, Applications Model Description, <http://www.opnet.com>.
- [PECH98] Timothy M. Persons, William J. Chimiak, "*Distributed Radiology: WFUBMC PACS Implementation*", 1998.
<http://citeseer.nj.nec.com/persons98distributed.html>.
- [PEGI97] Petri Chaudio, Giachetti Andrea, "*Medical Images transimmission on heterogeneous networks*", Proceedings of the 15th EuroPACS Annual Meeting, September 1997, pp. 117-120.
- [PRIO95] Prior F. W., Glicksman R. A., Wilson D. L., Keller S., "*Image Management in a Multi-Hospital Environment*", proceedings of the Fourth International Conference, August 1995, pp. 173-181.
- [RALP99] R. Martinez, D. Bradford, J. Hatch, J. Sochan, W. Chimiak, "*Modeling and Simulation of the USAVRE Network and Radiology Operations*", Virtual Radiology Environment Symposium, May 1999, online paper, <http://www.ece.arizon.edu/~cerl/symposium99>.
- [RATI93] Ratib O., "*The Geneva PACS Project*", Proceedings of the Third International Conference on Image Management and Communication in Patient Care, IMAC93, June 1993, pp. 37-44.

References

- [REIJ95] Reijns Gerard, "A *Distributed Picture Archiving and Communications System for Hospitals using Image Pre-fetching*", Proceedings of the 28th Annual Hawaii International Conference on System Sciences, 1995, pp. 470–475.
- [RFC1323] Jacobson V., Braden R., Borman D. "RFC 1323, *TCP Extensions for High Performance*", The RFC-Editor, May 1992, <http://www.rfc-editor.org>.
- [RFC2414] Allanman M., Floyd S., Patridge C., "RFC 2414, *Increasing TCP's Initial Window*", The RFC-Editor, September 1998, <http://www.rfc-editor.org>.
- [ROBE00] Ed Roberson, "The Bandwidth Battle", Healthcare Informatics, January 2000, <http://www.healthcareinformatics.com>.
- [RUGG98] Carmelina Ruggiero, "A *Teleradiology Primer*", Journal of Telemedicine and Telecare, Vol. 4, No.1 1998, pp.25–35.
- [SAKU92] Kudrimoti A.S., Sanders W.H., "A *modular method for evaluating the performance of picture archiving and communication systems*", Computer-Based Medical Systems, 1992. Proceedings, Fifth Annual IEEE Symposium on, June 1992, pp. 44-53.

References

- [SHALIN] Lin C., Shahnasser H., “*A Networked Picture Archiving and Communication System*”, presented at the IEEE International Conference on Telecommunications, June 2001.
- [SIEG02] Siegel Elliot, “*PACS and the Rest of the Healthcare Enterprise System*”, Imaging Technology News, February 2002, <http://www.itnonline.net>.
- [SIG01] SGI and UCLA LONI, “*Taking Brain Research to the next Level*”, UCLA Laboratory Neuro Imaging, <http://www.sgi.com/pdfs/3155.pdf>.
- [TCPPT] TCP Tuning Guide for Distributed Application on Wide Area Networks, <http://www.didc.lbl.gov/tcp-wan.html>.
- [TELCO1] Telco Systems, A BATM Company, “*Layer 3 Switching White Paper*”, online paper, www.telco.com.
- [TSIK96] Stelios C. Orphanoudakis, Eleni Kaldoudi, Manolis Tsiknalis, “*Technological Advances in Teleradiology*”, European Journal of Radiology, Vol. 2, 1996, pp. 205–217.
- [WILL91] Dallas W. J., “*A Hi-Speed Network for PACS Research*”, The Second International Conference on Image Management and Communication

References

(IMAC), in Patient Care: New Technologies for Better Patient Care, April 1991, pp. 32–35.

[WONG95] Huang H., Wong W, et. al., “*Asynchronous transfer mode technology for radiologic image communication*”, American Journal of Roentgenology, Vol. 164, No. 6, 1995, pp.1533–1536.

[WRHA99] Winnipeg Regional Health Authority Website, <http://www.wrha.mb.ca>.