Proceedings of the 1992 Winter Simulation Conference
ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson

Hierarchical Rapid Modeling of Picture Archiving and Communication Systems Using LANNET II.5 and NETWORK II.5

Emil Wirsz
Fred W. Prior

Siemens Gammasonics Inc.
2501 N. Barrington Rd.
Hoffman Estates, Illinois 60195, U.S.A.

Glenn A. Meredith
Kenneth R. Anderson

Siemens Corporate Research Inc.
755 College Rd. E.
Princeton, New Jersey 08540, U.S.A.

ABSTRACT

Siemens and Loral Western Development Labs have designed a Picture Archiving and Communication System capable of supporting a large, fully digital hospital. Its functions include the management, storage and retrieval of medical images. The system may be modeled as a heterogeneous network of processing elements, transfer devices and storage units.

A series of discrete event simulation models have been developed to investigate different levels of the design. These models include the Workstation Models, focusing on the internal processing in the different types of workstations, the Communication Network Model, focusing on the control communication and host computer processing, and the System Model, focusing on the flow of image traffic throughout the system.

This paper describes some of the issues addressed with the models, the modeling techniques used and the performance results from the simulations. Important parameters of interest include: time to retrieve images from different possible storage locations and the utilization levels of the transfer devices and other key hardware components. The models were an important part of the Loral/Siemens proposal for the Medical Diagnostic Imaging Support System (MDIS), helping to understand system performance under fully loaded conditions while the architecture was still in the design stage. With the first MDIS systems now online, the models continue to be useful in refining the design.

1 INTRODUCTION

The Medical Diagnostic Imaging Support System (MDIS) contract, awarded to Loral Western Development Labs and Siemens Gammasonics in September 1991, involves the implementation of state of the art Picture Archiving and Communication Systems (PACS) in several U.S. Military hospitals. A PACS is a computer based system for the storage, retrieval and viewing of radiological images. In the past, these images have primarily been handled with film, or some combination of film and digital imaging on a small scale. The Loral/Siemens PACS architecture introduces a new level of performance which is capable of supporting a large, fully digital hospital. The MDIS Request For Proposal (RFP), from the U.S. Army Engineering Division (1990), specifies a number of performance requirements to assure that the fully digital hospital will operate efficiently. Using computer-based modeling and simulation techniques, a performance analysis project for the Loral/Siemens PACS was undertaken to prove that the proposed design meets these performance requirements.

In addition to addressing the RFP requirements, the modeling and simulation efforts were an integral part of the design process, helping to explore design tradeoffs and identify key system tuning parameters. Within the basic system architecture, there are many design and operational issues which can benefit from a detailed performance analysis. The system performance is a multi-faceted issue due to the wide variety of hardware components and functional modes of operation involved. Because of the size of the PACS, building prototypes of the entire system to investigate design decisions is impractical. Thus, during the design stage modeling takes on a significant role in investigating the system performance under fully loaded conditions.

In describing their work in PACS performance analysis, Martinez et al. (1990) classify the nodes on a PACS network as image acquisition equipment, viewing workstations and database archive systems. While these components are common to all PACS, from one architecture to another there can be significant differences in the network configurations and the structure of the database archive system. In the Loral/Siemens PACS architecture, the network consists of a number of separate network segments, which can be grouped into two major classifications: a LAN handling the command traffic at 10 Mbps and a fiber optic distribution system handling the image traffic at 100 Mbps.

The Loral/Siemens PACS is based on a shared file system architecture, as described by Glicksman, Wilson, Perry and Prior (1992). Such an architecture has the advantage of giving all workstations equal access to the complete file system. A key challenge in this approach is to minimize the contention delays at the storage system. The Working Storage Unit (WSU) is a RAID 2 disk array that serves as short term and local storage for all image data in the system. The WSU has 28 I/O ports that are connected to a fiber distribution system. The WSU backplane is capable of supporting three simultaneous 100 Mbps transfers through the fiber distribution system. Long term storage is provided by two optical disk jukeboxes connected to the WSU. New image data is fed into the WSU from the image acquisition equipment, involving a number of different modalities, Computed Radiography (CR) being the most significant.

PACS workstations are classified as two basic types: the Optimized workstation (OWS) and the Standardized workstation (SWS). An SWS provides Radiologists with primary diagnosis and reporting capability via two to eight portrait mode displays. These displays may be "A" type and support 1535 pixels x 2048 lines or "B" type at 1024 pixels x 1280 lines. Special dual channel display cards for the Macintosh NuBus permit up to eight such monitors to be driven by a modified Macintosh IIFX computer.

An Optimized workstation provides lower cost softcopy viewing capability for all other clinicians within the Medical Treatment Facility. Up to four displays are provided with a resolution of 1152 pixels by 882 lines. Like the SWS, the OWS is based on an enhanced Macintosh IIFX computer.

In addition to the dual channel display cards, the Macintosh IIFX-based workstations are enhanced with specialized hardware and software. The Siemens' OPUS image processing card provides up to 65 MBytes of high speed buffer storage with specialized image processing hardware. The OPUS moves data on the Macintosh II NuBus at block transfer speeds up to 33 MBytes/second.

The Loral Fiber Optic Interface provides the workstation connection to the Shared File Server over fiber optic links at rates up to 100 Mbits/second. The card provides reversible image compression which effectively doubles both image cache storage capacity and data rate while preserving full image fidelity.

In the larger MDIS facilities there will be over 100 workstations all of which have the ability to retrieve image data from the WSU. A Host computer is also used for the database transactions in the system.

While the MDIS contract involves several specific hospitals, the performance analysis focuses on the Madigan Army Medical Center, as it is one of the largest systems in the contract and the earliest to be implemented. Three phases of the Madigan system are defined to represent the hardware configuration and expected loads at different stages of the implementation. While the modeling project investigated each of these three phases in detail, the results in this paper pertain to the Madigan Phase III configuration.

2 MODELING TECHNIQUES

To more efficiently represent the various aspects of the system, the modeling effort followed a hierarchical approach, involving analytic models and several discrete event simulation models. At the highest level, extensive analytic models were implemented in Microsoft Excel spreadsheets. These provided a rough analysis of the architecture-specific system requirements relative to the general requirements provided in the RFP. These spreadsheets also provided some mean value performance estimates using Markov queuing models.

The discrete event simulation effort involved three distinct modeling areas, each focusing on a different aspect of the PACS performance. The Workstation Models address the performance internal to the various viewing workstation configurations. The Communication Network Model focuses on the control communication and host computer processing. The System Model focuses on the flow of image data throughout the system, with a particular emphasis on the storage systems. The Workstation Models and the System Model were implemented using NETWORK II.5, while the Communication Network Model used LANNET II.5, both of these tools being from the CACI Products Company.

NETWORK II.5 and LANNET II.5 are both discrete event simulation tools aimed at modeling computer, communication and information systems. The two are closely related, with LANNET II.5 being a slightly higher level tool, as it provides certain predefined constructs and protocols which are more specific to modeling LANs. The code produced by LANNET II.5 is downward compatible with NETWORK II.5's simulation engine. Both tools allow a system to be modeled in terms of hardware and software at various levels of abstraction. Most of the following descriptions relate to NETWORK II.5, although the techniques in LANNET II.5 are very similar.

As the three simulation models do overlap in their view of the system, some cross-referencing is done to represent the performance results from one model in another. These relationships are shown in Figure 1. Because the sequences in the actual system have some distinct breaking points between different actions, some of the delays measured in one model can be aggregated in another model without loss of accuracy.

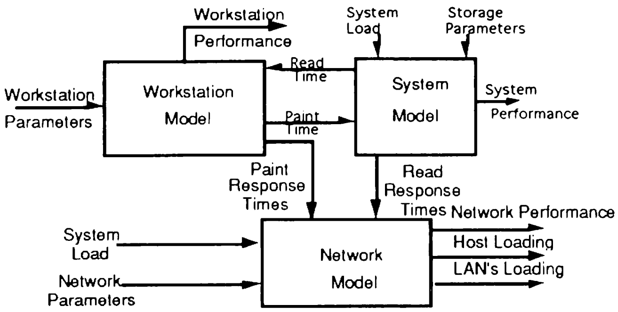

Figure 1: MDIS Simulation Model Relationships

2.1 Hardware Representation

In modeling a PACS with NETWORK II.5, each hardware component is represented as a "processing element", a "storage device" or a "transfer device". Thus, the Ethernet and fiber optic segments are all modeled as transfer devices, and the WSU disks and optical disks are modeled as storage devices. The processing element construct allows a little more abstraction than the others. A node such as a viewing workstation could be modeled as a single processing element, or as a CPU, an internal bus, some memory, etc., depending on the focus of the model and the detail required. The more detailed approach is taken in the Workstation Models, representing the internals of a single workstation by a number of hardware components. The System Model is primarily concerned with image traffic into and out of the WSU, and as such has a relatively detailed representation of the WSU, while most of the other nodes in the system are modeled with single processing elements.

2.2 Software Representation

Within the hardware configuration, NETWORK II.5 or LANNET II.5 allows a detailed software representation. The PACS activities which must be modeled in terms of software processes can be broken down into the following basic tasks:

- send a new image from imaging equipment to WSU,
- retrieve an image from WSU to a viewing workstation,
- send an image from WSU to archive storage,
- retrieve an image from archive storage to WSU.

Within each of the three simulation models, different aspects of these basic tasks are modeled in more detail, depending on the focus of that model. For example, in the Communication Network Model, the task of a viewing workstation retrieving image data from the WSU will have a detailed scenario of database queries and other command messages passed back and forth between the workstation and the host before getting to the actual image transfer. In this case the image transfer is not modeled in as much detail as it is in the System Model, because the focus is on measuring performance of the LAN and the host.

The basic tasks are represented in the models in many different instances, referenced to specific hardware components. In terms of Network II.5, this involves "modules" which will execute on specific processing elements, and "instructions" which are referenced within the module specifications and defined in the processing element specifications. The instructions are of several basic types: message, processing, semaphore, read or write instructions. The modeling of any one specific task involves the execution of a series of modules and their associated instructions.

2.3 Model Loading

The occurrences of these tasks originate with the execution of specific initiating modules, modeled as Poisson arrivals with exponential interarrival times. The exponential interarrival time distributions for each type of task are derived from the estimated number of new images per year, of several different types, and the number of nodes involved in executing each task. The yearly input rates are translated into daily rates in terms of CR equivalent images (10 MB of data per image in uncompressed form), and then, using the 5 Busy Hour Day concept, into rates in CR equivalent images per minute. The 5 Busy Hour Day concept is an approach that derives a worst case scenario for system loading by using the assumption that the peak hour's loading represents 20% of the entire day's loading. The system is continuously operated at this peak rate, compressing the whole day's loading into a "5 Busy Hour Day."

The retrieval rates for each type of viewing workstation are based on these new image input rates. This is done with the assumptions that every new exam will be reviewed an average of one time on an SWS, and an average of three times on an OWS. It is assumed that two historical exams are retrieved with each new exam. Each exam is assumed to consist of 3 CR equivalent images.

Each task, once initiated, involves a long series of steps including communication, queuing of requests, processing delays, data transfer, etc. The "message" instruction type is being used to represent the requests, as well as the data transfers. Thus, the size of these messages varies significantly. One query to the database might only be 12 KB while one CR image in the WSU will be 5 MB when compressed 2:1.

2.4 Objectives

In a PACS with over a hundred nodes, there is a significant potential for contention in several areas. Multiple nodes may simultaneously be attempting to send new image data into the WSU, or to retrieve stored image data from the WSU. Multiple database queries may competing for time on the host. A single Fiber Optic Distributor node may be requested by several of the workstations which share it for access to a WSU port. A major issue in the performance analysis of any PACS is the speed with which images can be retrieved from storage and brought up on the screen of a workstation. There are two sides to such response times. First, there is the peak performance in which the task occurs in the absolute minimum time that is allowed by the actual hardware and software implementation, given that there is no contention. These times can be measured with prototypes. On the other hand, there is the response time which includes some stochastic amount of queuing delay due to contention in the system. Much of the purpose of the simulation-based performance analysis involves investigating the dynamics of the system under these stochastic influences and determining the statistical properties of such responses times.

An accurate contention model requires reasonable representations of both the image transfers and the control logic for processing those transfers. These issues are fully addressed in the simulation models by implementing the same basic control logic that is used in the actual system.

2.5 Output Analysis

The standard output listings provided by Network II.5 include extensive statistics on all hardware elements and software elements defined in the model. Often, we are interested in specific response times which span a long series of module executions. In such cases the semaphore instruction construct is used to measure these response times. By "setting" and "resetting" the semaphores at specific points in the simulation, the response times can be read directly from the "semaphore statistics" section of the output listing, providing average, standard deviation, minimum and maximum values. To get a more detailed picture of these response times, the semaphore "set" and "reset" instructions are added to the trace report. A program has been written in C that parses the trace report, picking out the instructions corresponding to a specific semaphore, calculating the response time and writing each such response time to an output file. In this way, complete data sets are generated for each of the response times of interest. Further analysis of the distributions for each response time is done using these data sets in Microsoft Excel and/or Unifit, from Averill M. Law & Associates.

2.6 Simulation Validation

The validation of the three models, the workstation, the network and communication, and the system model imply comparing the simulation results with measurements taken on the real system. It is easier to perform measurements on the workstations, the networks and the WSU. The high complexity of the database sequences, the variety of configurations in which the database can be configured, and the necessity to accommodate a growing demand of new database requirements, call for additional detailed modeling of the database activities. These modeling activities are well under way at Siemens Gammasonics and will be made available soon in future publications

3 THE WORKSTATION MODELS

A series of Workstation Models were developed, each representing the details of a particular viewing workstation's internal architecture. The performance analysis done with the Workstation Models not only provided data for use in the other models, but was a significant investigation of design tradeoffs in its own right. Different workstation configurations were analyzed, using different monitors (A, B, or C type) and different OPUS memory sizes for storing images retrieved from the Working Storage Unit. All workstation configurations are based on the Macintosh IIfx and use an additional Fiber Optic Interface (SIC) buffer to enhance the system performance when small images have to be displayed. The final workstation configurations, which represent optimum performance and lowest cost solutions to the MDIS requirements, are as follows:

- Standardized Workstation (SWS) - Macintosh IIfx based workstation with up to 8 A and/or B type monitors, 64 MB OPUS memory , 5 MB SIC buffer, and Loral NDDC display cards.
- Optimized Workstation (OWS) - Macintosh IIfx based workstation with up to 2 C monitors, 32 MB OPUS memory, 5 MB SIC buffer, and Radius display cards.
- Software Only Optimized Workstation (OWS-L) - Based on the hardware of the OWS without the OPUS memory, with up to 2 C monitors, 5 MB SIC buffer, and Radius display cards.

A block diagram of the Standardized Workstation's hardware components, as represented in the Network II.5 animation developed along with the simulation model, is shown in Figure 2. Here the NuBus is the main transfer device internal to the workstation, providing connections between the main Mac IIfx CPU, the OPUS board, the displays and the interface cards. The diagram also shows the workstation's Ethernet connections to the host and the WSU, and the fiber link to the WSU. The other workstation models follow a similar approach.

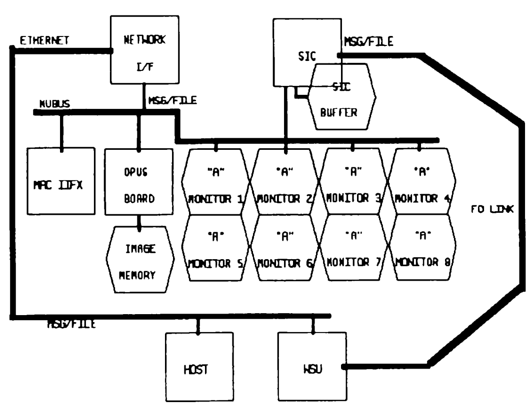

Figure 2: Workstation Model Animation Screen

To come to the final workstation configurations described above, an extensive engineering study was performed. Many parameters were considered: workstation configurations, monitor types, OPUS memory sizes, SIC buffer sizes, image sizes and types, and combinations of these parameters. Thus, much research was done and many simulation models were employed before the final decisions were made.

A major concern was to find the optimum SIC buffer size without sacrificing performance and still keeping costs to a minimum. Although physically present on the Fiber Optic Interface, the SIC buffer is activated by the workstation software only for small images, such as CT, MR, or NM images. This particular software behavior, as well as all the other software processes and activities which determine the functionality of the workstations, are explicit constructs of the simulation models. These software representations consider all the interactions between the workstation hardware components and their environment. The engineering study led to finally considering 5 MB to be the optimal size of the SIC buffer.

The primary performance measures provided by the Workstation Models are the response times measured from the time when an image enters the workstation until the screen is completely painted. These times, referred to as the "paint times", are used in the System Model and in the Communication Network Model, where a workstation is represented as a single processing element with an aggregate "paint time" delay. The average paint times for a CR image, as measured from the simulations, are as follows:

- SWS, A monitors: 803.8 msec
- SWS, B monitors: 366.4 msec
- OWS, C monitors: 176.4 msec
- OWS-L, C monitors: 1019.1 msec

4 THE COMMUNICATION NETWORK MODEL

The Communication Network Model, as is described in greater detail by Anderson et al. (1992), focuses on the control communication between the system nodes and the host computer, and the associated processing delays on the host. Before any transfer of image data into or out of the WSU can occur, the user at an acquisition or viewing workstation will go through a series of steps interacting with the host to set up the transfer. Because there are multiple workstations sharing a particular LAN segment and because all nodes in the system share the same host computer, there are significant possibilities for contention. To fully capture these interactions, the model includes a detailed hardware description including all of the individual LAN segments and bridges. Within this hardware definition the model executes a number of communication and host processing scenarios, in order to represent the demands on the LAN and the host under fully loaded conditions.

Figure 3 is a simplified view of the hardware configuration in the MDIS Communication Network Model. It contains five bridges connecting five backbone LAN segments. Each bridge also connects to one or more 10 Mbps Ethernet LAN tributaries. These tributaries connect to the various acquisition, viewing and storage nodes as described above. In addition there is the Comprehensive Health Care System (CHCS) which represents a hospital information system with which the PACS must interface for such things as patient scheduling.

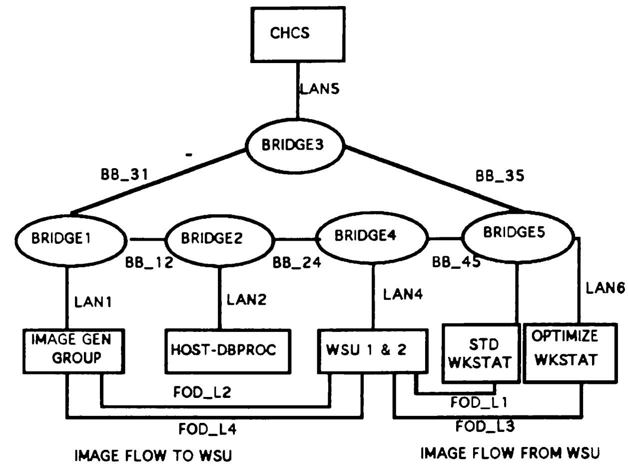

Figure 3: MDIS Communication Network Model

Any particular message going from a node to the host will pass over a LAN to a bridge connected to a backbone segment. The message could possibly have to go through another bridge and another backbone segment before reaching the bridge which connects to the LAN that the host is on. As there are protocols and possible contention involved in each leg, one such transfer of a command message can be an involved process. LANNET II.5 facilitates this type of model by having predefined constructs for the LANs, the bridges and the protocols.

The two major processing tasks of the host computer involve the shared file system and the database (implemented in Sybase). The processing delays which are defined in the model represent aggregate delays that a given database operation could be expected to encounter in a multitasking environment with competing processes. These times were arrived at through experimental measurements with the actual database.

The communication scenarios which model all of the traffic to and from the host, and the host processing delays, correspond to variations of the basic operational tasks described above in section 2. For example, Table 1 shows part of a workstation viewing scenario. Each communication from the workstation to the host is followed by some processing delay at the host and a response sent back to the workstation. In addition to the more deterministic delays included in the scenario, there will possibly be queuing delays associated with accessing the LAN segments and accessing the required host or database processing element.

Table 1: Workstation Viewing Scenario

- values in brackets are message size in bits
- values in parentheses are processing delays in ms

ACTION	SOURCE	DESTINATION
connect [48K]	STD WS	DPROC1 (50)
connect ack [48K]	DPROC1	STD WS
login [96K]	STD WS	DPROC1 (200)
config [288K]	DPROC1	STD WS
wrklist query [96K]	STD WS	DPROC1 (500)
wrklist resp [96K]	DPROC1	STD WS
open query [96K]	STD WS	DPROC1 (2000)
open resp [96K]	DPROC1	STD WS
open file [48K]	STD WS	HOST1 (10)
open ack [144K]	HOST1	STD WS
read file [96K]	STD WS	WSU (10) + (Fetch Delay)
images [40M]	WSU	STD WS
wrklist update	STD WS	DPROC1 (50)
update resp	DPROC1	STD WS
close file	STD WS	HOST1 (2)
close ack	HOST1	STD WS

In initiating scenarios such as that shown in Table 1, the model aggregates the nodes of a certain type into one or two processing elements. The generation rate of the appropriate scenario for a certain node type is then defined so as to simulate the complete load from all of the individual nodes of that type.

The major objectives of the Communication Network Model are to predict the level of congestion at the LANs and the host. Results show that the maximum number of messages in any queue awaiting LAN transfer during a 5 hour simulation is 2, indicating that the LANs and bridges are not congested. This is also shown by the average utilization levels in Table 2. LAN2, which is the LAN segment connecting to the host, shows the highest utilization, but in the 7 to 10% range this is still quite low. The two experiments shown in Table 2 all correspond to the Madigan Phase III configuration. The "4 Busy Hr" experiment has the same total load as in the 5 Busy Hour Day case, but it has been compressed into 4 hours by increasing the image generation and retrieval rates.

Table 2: Network and Host Utilizations (%)

	5 Busy Hr	4 Busy Hr
HOST	29.0	38.5
DBPROC	1.3	1.7
LAN2	7.2	9.7
LAN3	1.2	1.6
FOD_L1	11.6	
FOD_L2	4.7	
FOD_L3	23.0	
FOD_L4	41.0	

5 THE SYSTEM MODEL

The System Model, as is described in greater det Meredith et al. (1992), focuses on the flow of image throughout the system, with a detailed representati the WSU and the archive storage system.

Figure 4 shows the animation screen for the Ma Phase III configuration of the System Model. WSU's 14 I/O cards (C1-C14), each with two port shown connected to the disks (DC1) over the back (BP1). Many of the WSU I/O ports are connect Fiber Optic Distributors (FOD) which allow a s channel to be switched to one of several fibers. B some of the FOD switches there are two or workstations each modeled by individual proce elements connected to the FOD by dedicated fi Thus, at any one fiber there is no contention, but th a possibility for contention at the switch, as mul workstations below the FOD may simultaneously re the same FOD and WSU port.

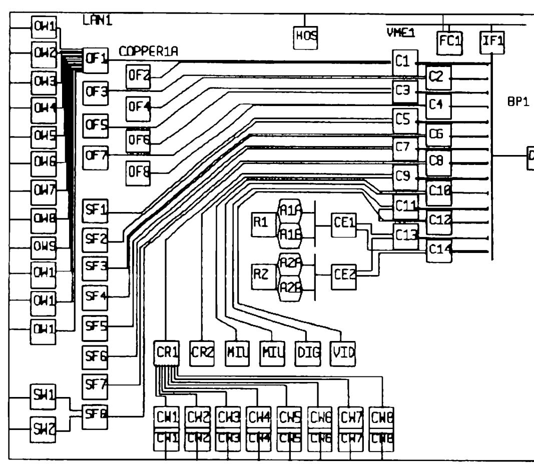

Figure 4: System Model Animation Screen

Since there is a fair amount of repetition in the sy configuration, with multiple FODs which are shared by some number of nodes of a certain type, abstractions are introduced to simplify the model. WSU port is shared by 12 optimized worksta (OW1-12) or 2 standardized workstations (SW1-2) CR acquisition workstations (CW1-8). Among nodes sharing one port, the possibility of conte among them to gain access to the port is an i However, the behavior of any one of these F connecting to a certain node type, can be consider be characteristic of any other FOD which is share the same number of nodes of that type. So, for node type, it is sufficient to model only one channel in complete detail with individual f connected to each individual node, in order to stud FOD contention. The other FODs which connect t type of node are still important individually, since do connect to individual WSU ports which may ea competing for access to the backplane. However they can more simply be modeled as aggregate FOD nodes defined with certain aggregate traffic rates.

In addition to the above described imaging equipment and viewing workstations, the other major hardware components which are connected to the WSU relate to the long term archive storage. The system includes two optical disk jukeboxes (ODJ), each modeled with two disk drives (A1A and A1B, or A2A and A2B) and one robot (R1 or R2) which services those two drives. The two drives in one ODJ are both connected to one Archive Controller (AC, or CE in Figure 4) by a SCSI bus. Each AC is connected to one WSU input port and one WSU output port.

The level of abstraction in the System Model is aimed at accurately representing the image transfers throughout the system. The major delays which make up the overall response times experienced by the user are the image transfer time and the queuing delays waiting for the transfers to be initiated. There are some CPU intensive operations associated with the database queries before a particular image is requested, these operations being modeled in greater detail in the Communications Network Model. However, once a particular transfer is requested, there are no significant CPU processing delays. Thus, the System Model focuses more on the transfer devices and the blocks of image data, and goes into less detail in modeling the actual processors. For the processors other than the host, the important factors are not so much the processing delays, but rather the queuing logic which the actual processors implement.

In general, the requests for a particular hardware element are handled on a first-in-first-out basis. However, the model often involves a series of queues through which requests must pass before some particular action is taken, assuring that competing processes have equal access. Relative to the WSU, exam requests are broken down into image requests and into block requests where the complete image can not be transmitted continuously because of buffering limits. Thus, once an image request gets past the queue at the workstation, each block of image data must essentially queue up for the required WSU port and then for the WSU backplane assignment. The objective of the queuing logic is to equitably balance the delays experienced by any one process with those of any competing processes. This implies maximizing the utilization of the WSU backplane which is the most central transfer device over which all image data transfers must flow.

Another aspect of the model relating to transfers into or out of the WSU is that, although it is a discrete-event simulation model, the continuous pipeline nature of the data transfers is accounted for. When a block of image data is being transferred from the WSU disks to a workstation, different parts of that data will simultaneously be moving over the WSU backplane, through the port buffer, over the copper connection to the FOD and over the fiber optic link to the workstation. The timing aspects of the transfer are specifically modeled using short dummy messages that activate the various legs of the transfer. In this way, the model captures the fact that the backplane will be released significantly before the port buffer has finished sending out the data block.

Regarding the queuing for the Archive Controller units and for the optical disk drives, these are essentially FIFO queues, although priority is given to reads over writes. Another priority scheme distinguishes between the two drives in one ODJ, with the A drive being given priority for reads and the B drive being given priority for writes.

With the basic System Model a number of experiments were performed to investigate different operating scenarios. Three phases of the Madigan PACS were modeled in detail, with each phase having a different hardware configuration and different loading parameters. The results presented in Table 3 are all based on the Madigan Phase III model, under the 5 Busy Hour Day and 4 Busy Hour Day assumptions. The third experiment in Table 3, labeled "diskfail", also using the 5 BHD assumption, simulates the system performance after one of the WSU disks has failed and been replaced with a new disk. To restore the data that was on the failed disk, the WSU runs a restoration process which results in a significant increase in traffic on the backplane.

Table 3: Average Response Times and Utilization Levels

	5 BHD	4 BHD	diskfail
response times (milliseconds)			
OWS, 1st image WSU fetch	748	746	840
SWS, 1st image WSU fetch	1401	1371	1467
OWS, 1st image archive fetch	24437	29330	22461
SWS, 1st image archive fetch	25055	28154	21842
request for drive in ODJ1	2559	6487	1795
request for AC1 for decompr.	8703	11406	7689
request for AC for compr.	11688	51974	14424
CRAW wait for FOD channel	15	16	40
OWS wait for FOD channel	12	9	6
OWS wait for WSU backplane	53	76	121
hardware utilization (%)			
WSU backplane	17.3	20.2	48.6
FOD channel to 12 OWS	3.9	4.5	3.6
FOD channel to 2 SWS	1.6	2.0	1.6
FOD channel to 8 CRAW	3.1	3.5	3.1
Archive Controller 1	72.6	84.9	71.9
Archive Controller 2	72.7	82.9	69.1
ODJ 1 drive A	37.0	48.9	37.0
ODJ 1 drive B	48.3	62.4	44.6
ODJ 2 drive A	36.0	43.5	32.9
ODJ 2 drive B	46.1	57.3	43.7

While the response times presented in Table 3 are all average values, at times there will be significant deviations from these times due to the stochastic nature of the system loading and the queues that will build up. Thus, it is important to investigate further the statistical distribution of the complete data sets which these averages represent. Figure 5 is a cumulative probability chart for the OWS first image fetch response time. The chart shows that the mean value is not too much higher than the minimum value, although there are some longer response times which occur with very low probability. The peak performance, with no contention in the system, is indicated to be about 0.6 seconds. As the corresponding mean is about 0.7 seconds, we can say that on average the response time for this first image retrieval, directly from the WSU, will include roughly 0.1 second of contention delay. In the case where the response time is 4.3 seconds, occurring with a probability of about 0.05%, the greater part of that time will be contention delay, most likely because of a number of requests arriving at the WSU closely together.

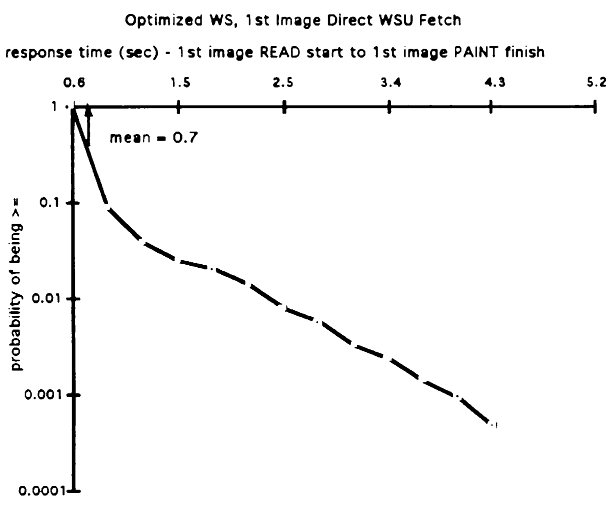

Figure 5: OWS 1st Image WSU Fetch Response Time

In addition to the above experiments, there were other experiments performed with the System Model at earlier stages in the modeling process, which helped to guide some of the decisions in the design of the actual system. An example of such experiments is in the issue of the number of workstations which can efficiently share one WSU port. After a series of experiments considering several alternatives, the conclusion was that 12 OWS, 2 SWS or 8 CRAW could share a single port.

6 CONCLUSIONS

The analytic and simulation models developed for the MDIS project have helped to evaluate the system performance under fully loaded conditions, while still in the design stage. The models indicate that the Loral/Siemens PACS design will meet the MDIS performance requirements. For the complete retrieval of the first image of an exam from short term storage, the requirement is 5 seconds. For the Madigan Phase III simulation, the mean response times of 0.748 seconds and 1.401 seconds, for the Optimized and Standardized Workstations respectively, are well within the requirements. In addition to the means being at low levels, the simulation results show that the random occurrences of significantly higher response times have very low probabilities. The probability of exceeding the 5 second limit for this first image response time in any specific request less than 0.05% for the Optimized Workstations and about 0.2% for the Standardized Workstations.

In addition to these overall response times that are experienced by the user, the simulation models offer detailed insight into every aspect of the system operation which contributes to these overall times. The average host utilization has been shown to be less than 30%. The WSU backplane, which is the central conduit for all image transfers in the system, has been shown to have a very acceptable average utilization of 17%. The LAN segments have been shown to have very low average utilization levels, many below 2%, with the highest being in the 7 to 12% range. The Workstation Models have analyzed the tradeoffs of performance between different internal workstation architectures, providing direct input to the higher level models. The expected queuing delays have been investigated at many different points in the system. Some of these queues have been shown to have fairly minor delays, such as the queues for an individual FOD channel. In other areas, such as the long term archive storage, the utilization levels have been shown to be fairly high and correspondingly there are some longer queuing delays associated with those components. In this case, the model results have led to further refinement of the design since the prototype stage.

The MDIS modeling project has proven to be a valuable part of the design process for a large heterogeneous PACS system. As the models were developed in close collaboration with the system architects and designers, it was assured that the models would accurately reflect the actual system and that immediate feedback could be offered to assist in design decisions. At the final stage of the MDIS proposal process a fully functional prototype, the System Test Bed, was developed for the benchmark test. With the System Test Bed, some initial model validation measurements were obtained. Since the contract award the system design and the models have continued to evolve in an effort to optimize the system performance. The hierarchical decomposition of the PACS into focused submodels has helped assure that the models accurately represent the real system while remaining flexible enough that new issues can be investigated as they arise in the design and implementation process.

ACKNOWLEDGMENTS

The authors would like to acknowledge Dennis Wilson and Robert Glicksman of Loral Western Development Labs for their insights and support as key members of the MDIS Modeling Project. Particular thanks also go to Paul Sochacki of Siemens Corporate Research for sharing his network expertise, both in simulation techniques and in the experimental validation of model parameters.

REFERENCES

Anderson, K., G. Meredith, E. Wirsz, F. Prior and D. Wilson. 1992. Hierarchical rapid modeling and simulation of high performance picture archive and communication systems. In *Medical Imaging VI: PACS Design and Evaluation, Proc. SPIE* 1654:577-588. The International Society for Optical Engineering, Bellingham Washington.

CACI Products Company. 1991. LANNET II.5 User's Manual. CACI Products Company, La Jolla, CA.

CACI Products Company. 1991. NETWORK II.5 User's Manual. CACI Products Company, La Jolla, CA.

Glicksman, R., D. Wilson, J. Perry and F. Prior. 1992. Architecture of a high performance PACS based on a shared file system. In *Medical Imaging VI: PACS Design and Evaluation, Proc. SPIE* 1654:158-168. The International Society for Optical Engineering, Bellingham Washington.

Law, A.M. and S.G. Vinnt. 1991. *UniFit II User's Guide*. Averill M. Law & Associates, Tucson, AZ.

Martinez, R., W. Sanders, Y. Alsafadi, J. Nam, T. Ozeki and K. Komatsu. 1990. Performance evaluation of a picture archiving and communication system using stochastic activity networks. In: *Medical Imaging IV: PACS System Design and Evaluation, Proc. SPIE*, 1355. The International Society for Optical Engineering, Bellingham Washington.

Meredith, G., K. Anderson, E. Wirsz, F. Prior and D. Wilson. 1992. Modeling and simulation of a high performance PACS based on a shared file system architecture. In *Medical Imaging VI: PACS Design and Evaluation, Proc. SPIE* 1654:169-179. The International Society for Optical Engineering, Bellingham Washington.

Microsoft Corp. 1989. *Microsoft Excel Reference*. Microsoft Corp., Redmond , WA.

U. S. Army Engineering Division. 1990. DACA87-90-R-0058, Multi MDIS System Solicitation. U. S. Army Engineering Division, Huntsville Alabama.

AUTHOR BIOGRAPHIES

EMIL WIRSZ is a Principal System Engineer at Siemens Gammasonics, Inc., Hoffman Estates, IL in the PACS Department. He has a worldwide responsibility for system architecture and modeling for Siemens. His research interests include Object Oriented modeling, and applications of expert systems and neural networks to PACS simulation. He received a MSEE in 1983 from the Polytechnic Institute "Traian Vuia" (Romania) and is currently involved in a program of study for his Ph.D. in Computer Science at the Illinois Institute of Technology.

FRED W. PRIOR is the Manager of Program Planning for PACS at Siemens Gammasonics, Inc., Hoffman Estates, IL. He is one of the primary architects of the Siemens PACS He currently heads a group which has worldwide responsibility for system architecture and new product definition for Siemens. Mr. Prior is also an active participant in the ACR-NEMA standards effort. His research interests include the RIS-PACS interface problem, the application of relational database technology to medical image storage and retrieval and applications of Expert Systems technology to PACS. He received a BA in Chemistry in 1974, an MA in Anthropology in 1976, an MS in Biomedical Engineering in 1984, all from Case Western Reserve University, and a Ph.D. in Computer Science in 1992 from the Illinois Institute of Technology. He is a member of the IEEE Computer Society and EMBS.

GLENN A. MEREDITH is a Research Scientist at Siemens Corporate Research Inc., Princeton, NJ, currently in the Learning Systems Department. His interests include discrete event simulation, Petri net based control, robotics, computer vision and expert systems. He received a BS in Mechanical Engineering in 1980 and an MS in Industrial Engineering in 1990, both from Rutgers University. He is a member of IIE, ORSA and SME.

KENNETH R. ANDERSON is a Group Leader in the System Technology Department at Siemens Corporate Research Inc., Princeton, NJ. There he pioneered the research activities in microelectronics design and test, and computer based rapid modeling of manufacturing and communication systems. Before joining Siemens, he was an engineering manager at RCA's Government Systems Division where he specialized in the reliability and testing of monolithic and hybrid integrated circuits. He received a BSEE from Drexel University in 1968, and an MA in Business Management from Central Michigan University in 1975. He is a Senior Member of the IEEE and is the 1989 President of the IEEE Computer Society. He also holds membership in the ACM, IIE, SCS, ASEE and the NTA. He is an Adjunct Professor of Electrical Engineering at Widener University, Chester, PA