Proceedings of the 1985 Winter Simulation Conference
D. Gantz, G. Blais, S. Solomon (eds.)

SIMULATION OF A HOSPITAL PICTURE ARCHIVING AND CONTROL SYSTEM (PACS)

G. R. Lawrence
G. A. Marin
S. E. Naron

IBM Federal Systems Division
Network Systems Laboratory
708 Quince Orchard Road
Gaithersburg, Maryland 20878

ABSTRACT

The nation's hospitals have long used distributed data processing as a means of reducing operational costs and providing timely service. Radiology Departments are now also taking advantage of these facilities in order to decrease the cost of producing and archiving radiological images. Today a typical medium scale hospital consumes large quantities of silver oxide film which, along with attendant labor costs, is expensive compared with costs for digital image processing technology now available. Using this technology large image files can be stored and retrieved through local area networks that can also support the transaction traffic essential in a hospital environment. The evolving systems are called Picture Archiving and Control Systems (PACS).

PACS will include radiology imaging equipment, distributed and central image archiving facilities, and significant numbers of user work-stations and graphics display nodes. The devices will be interconnected by high speed local area networks capable of distributing information ranging from simple control messages to large image files of several megabytes in a fashion offering most users a response time of several seconds.

This paper illustrates the PACS system concept, present a queuing model approach to analyzing PACS performance, and discuss results acquired for a variety of parametric samples. The IBM Research Queuing Package (RESQ) has been used for the exercise and will be discussed sufficiently for the reader to appreciate its capability.

RESQ simulation results indicate that system response times will be more dependent on the internal architecture and programs of the workstation than on the speed of the transmission media.

THE PROBLEM

Discrete simulation was used successfully to gain insight into architectural tradeoffs for a complex preprototype distributed information system. The application involved the archiving and distribution of radiological image data (X-Ray, Cat Scan, Magnetic Resonance, Ultrasound, etc.) within hospitals. The work was done by the Network Systems Laboratory of IBM's Federal Systems Division.

The data processing requirements of many hospitals make it cost effective for them to implement distributed data processing systems. Hospitals need to track numerous important transactions on large transient patient populations in order to deliver effective health care, perform inventory and staff accounting, and bill for services. Because of the high cost of health care professionals, hospitals were early to recognize the value of providing distributed access to their computers with digital terminal equipment (DTE) provided in or near most work areas.

Today's diagnostic radiological images are being generated in digital form, as well as on film, and can be distributed electronically throughout the hospital. Perhaps even more important, they can be economically archived (digitally) to avoid the large expense of silver-oxide film and attendant labor costs for film management, archiving, retrieval, and display. In addition to reducing costs, Picture Archiving and Control Systems (PACS) offer far more responsive service.

The typical PACS system will consist of radiology source imaging equipment (called modalities), distributed and central image archiving facilities, and significant numbers of high-resolution graphics display devices. The devices could be interconnected by high-speed local areas networks (LANs) capable of distributing information ranging from simple control messages (<2000 Bytes) to large image files of several megabytes. Typically PACS traffic consists of either: (a) large files being requested infrequently by operators at workstations throughout the hospital or (b) large image files being transferred from modalities located in the radiology department. The requested image files must be delivered within several seconds with high confidence. Performance standards are being defined in detail by the hospitals, the users, and the suppliers. The National Electrical Manufacturers Association (NEMA), and the American College of Radiologists (ACR) have played an active role in developing these standards.

Delivering large image files throughout a LAN environment within a few seconds in response to multiple users is a challenging problem. To contain costs the solution to this problem must include standard 'off-the-shelf' technology. The IBM Federal System Divisions's Network System Laboratory has completed preliminary analysis of this problem using discrete event simulation, and some of the results are presented in this paper.

THE IBM NETWORK SYSTEMS LABORATORY

In support of the PACS application and others like it, the Network Systems Laboratory has developed expertise in distributed system design. The Network Lab has pursued a comprehensive effort to define, configure, and analyze LANs that include products from multiple vendors. In 1983 the Lab began a multi-year internal research and development project that included the following major tasks:

- hybrid network architecture definition,
- laboratory configuration definition,

- o measurement technique development,
- o performance analysis and modeling, and
- o performance measurement and evaluation.

In work accomplished to date, these major tasks have been pursued using investigations of CSMA/CD, token ring, and Hyperbus networks. (Hyperbus is a trademark of Network Systems Corporation.) Based on this work the Lab has developed approaches (a) for defining software/hardware threads (single transaction paths) through LAN architectures, (b) for analyzing the performance of those threads against a varying background load, and (c) for simulating the performance of specific application architectures like those anticipated for PACS.

At the Network Lab a mixture of laboratory measurements, closed-form queueing models and discrete simulation are used. Each of these techniques has strengths and weaknesses that tend to be complementary. The fundamental approach has been to develop analytical methodologies where feasible and simulations where necessary. To some extent both can be benchmarked against measured data; successfully benchmarked models can then be used for extrapolation.

The analytical models developed or used in the Lab to date either address the LAN channel capabilities for homogeneous undetailed stations [1,2,3] or include the device hardware and software details for at most one particular foreground pair of devices running against a homogenous background load. Simulation is being used to look at specific system configurations, requirements for hybrid data terminal equipments, and mixed traffic conditions. Simulation was the tool of choice for PACs.

SIMULATING BACKGROUND LOAD

RESQ simulation results have been compared with results obtained from analytical models. The results of both approaches are now being compared with measurements. Figure 1 illustrates the results of

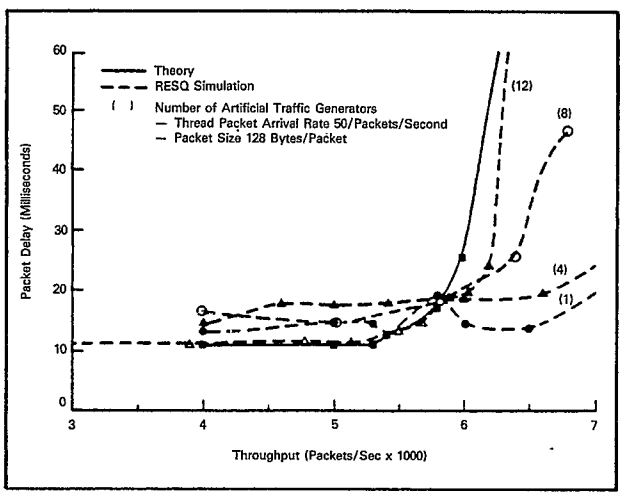

Figure 1: PCXT-PCXT Performance Using Artificial Traffic Generation (50 Prts/Sec into Thread)

one such comparison that reflects performance between two PC's communicating across an Ethernet* LAN using datagrams (no acknowledgement) at OSI level 2.

*A trademark of the Xerox Corporation.

The solid curve (theory) represents results obtained from an IBM PC-based algorithmic model that analyzes in detail a pair of foreground devices running with a homogeneous background load. This background load consists of the usual Poisson traffic uniformly distributed among a large number of background devices. Because we were interested in reproducing such a curve in a laboratory environment, the simulation results were obtained in several runs over which the background load was spread among different numbers of devices. In this way we were able to determine that the typical analytical background load could be achieved, for this particular case, by dividing the load evenly among at least twelve LAN-attached devices. Fewer devices may be required depending on system parameters.

Other comparison runs were made for Ethernet that included large numbers of simulated devices. Laboratory measurements have been made under both load and no-load conditions, and token ring comparisons have also been made. Results are beyond the scope of this paper but are generally supportive of the modeling.

THE TOOL -- RESQ

Simulation of the PACS application has been accomplished using the Research Queueing Package (RESQ) developed by IBM's Research Division [4,5,6]. RESQ is a high-level language designed to facilitate either numerical or simulation analyses of queueing systems. For the rest of the paper we will only be referring to the capabilities of the simulation 'side' of RESQ.

RESQ supports two general classes of queues: active and passive. Active queues are the normal type of service oriented queue that represent the principal system resources. For these queues RESQ provides several service disciplines and allows one or many servers. Passive queues are used to represent the simultaneous possession of resources, such as shared CPU memory. Both types of queues can be configured with one or more waiting lines.

The first step in modeling a communication network using RESQ is to represent the major network resources in terms of the RESQ modeling elements. This is done with a graphics notation language. In addition to the active and passive queues, the principal RESQ modeling elements are shown in Figure 2.

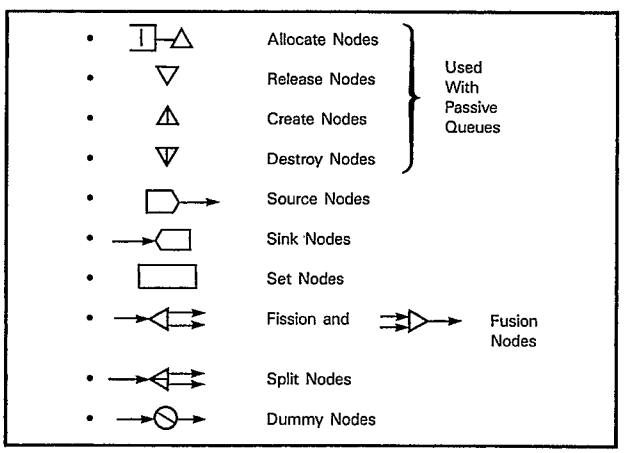

Figure 2: Principal RESQ Modeling Elements

In a communication network model, source nodes are used to introduce traffic into the network. Sink nodes are used to remove traffic from the network. Set nodes are used to assign characteristics to messages and to accomplish computations. The primary use for fission and fusion nodes is to model message packetizing. Messages are subdivided at fission nodes, move through the modeled network and are reassembled automatically by RESQ when the pieces reach a fusion node. A split node is often used to model acknowledgements.

RESQ keeps a complete set of statistics about every queue in the network. It gives the user access to a variety of statistical distributions. It provides for flexible routing and offers other user-friendly features. Perhaps its major advantage is its full support for modeling blocking due to resource constraints with passive queues.

For the past 18 months the Network Lab has pursued analytical and simulation models that can be used to address such application environments as PACS. The work leading to the PACS analysis includes the following:

- o developing and validating a RESQ-based Ethernet sub-model,
- o developing and validating a RESQ-based token ring sub-model, and
- o building a PACS architecture and traffic model "around" the two submodels.

Figures 3 and 4 present a simplified version of the Ethernet simulation model in RESQ graphics notation. Figure 5 presents the token ring model in similar notation. Reference 6 discusses examples related to these figures and is thus a source for additional insight into these RESQ diagrams.

SCENARIOS AND GOALS

Figure 6 shows a simplified PACS architecture. The devices on the left represent typical radiological sources: Cat Scan (CT), X-Ray, Magnetic Resonance (MR), and film digitizer (DIG). The host computer in the upper right acts as the archiving file server for the radiological pictures. The workstations in the lower right provide access for doctors to request and observe their patients' pictures. The key analytical assumptions include the following:

- o Device speeds and characteristics were educated guesses of generally available products in the 1986 to 1989 time-frame; for example, 4mbps token ring was compared with 10 mbps Ethernet.
- o Typical I/O software delays were included to accomplish protocol at OSI level 4.
- o Acknowledgements and windowing were not modeled.
- o Workstations were of approximately IBM PC AT capability, and the host was of IBM 43xx capability.
- o Data was transfered on the host channel in 64 kbyte units.
- o The transfers were memory-to-memory; no disk delays were included.
- o Messages were packetized in place in memory.
- o No adapter priority was modeled.

Some of these assumptions are optimistic, but do not interfere with our gaining an understanding of the

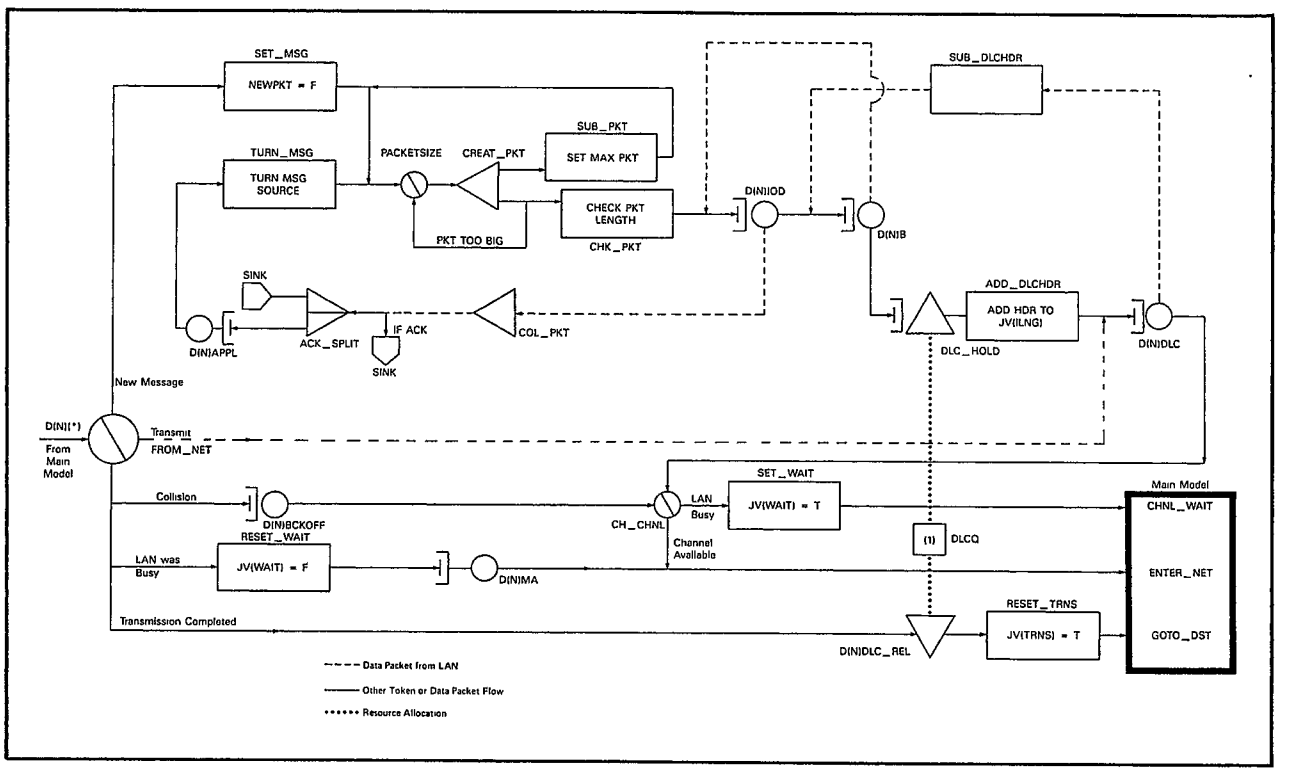

Figure 3: Ethernet Device Submodel

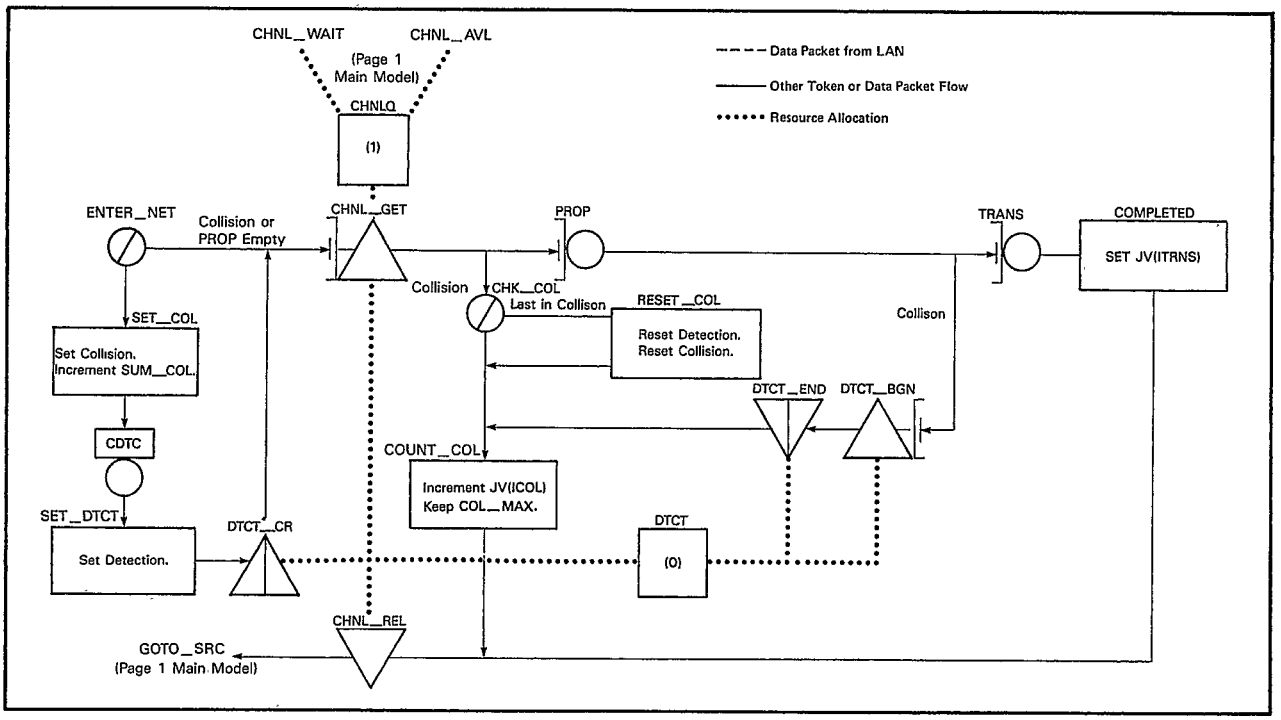

Figure 4: Ethernet Main Model

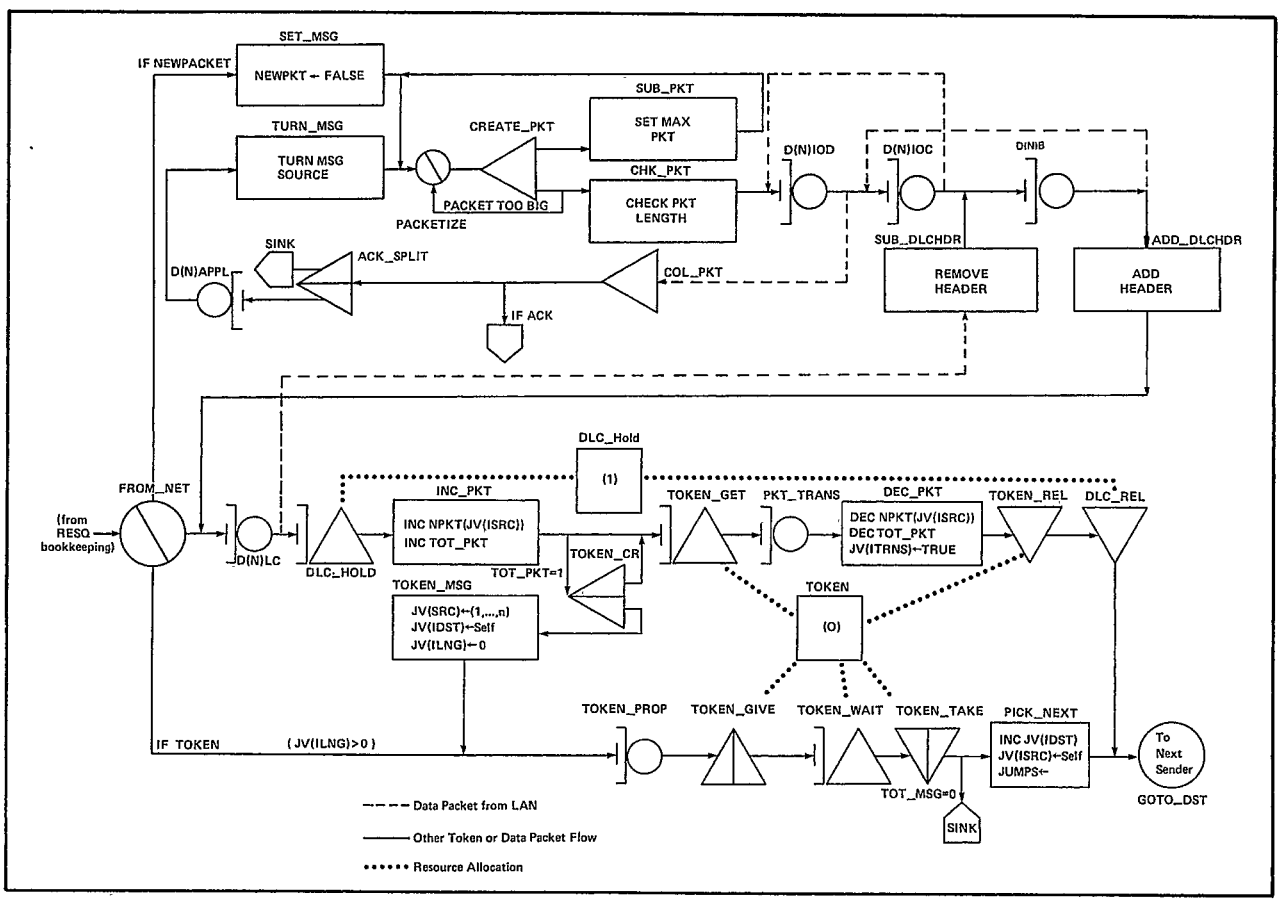

Figure 5: Token Ring Model

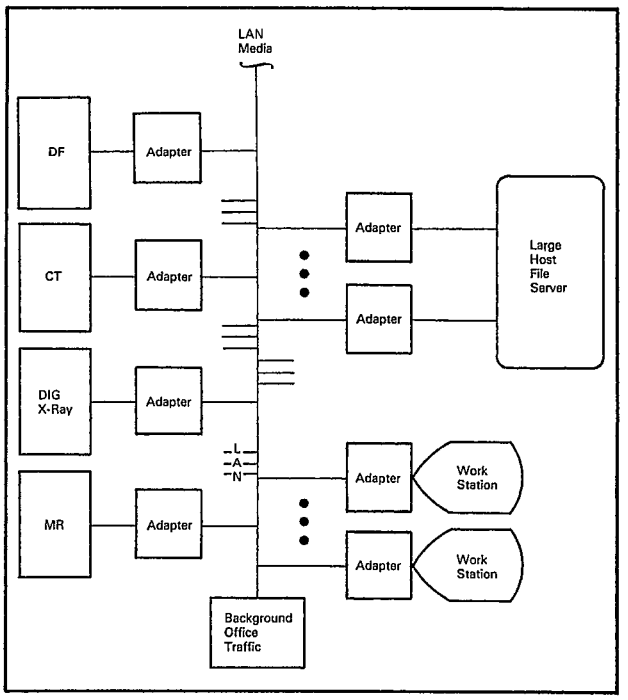

Figure 6: Simplified PACS Architecture

architectures and media access methods. In the case of acknowledgement and windowing, an appropriate sliding window protocol would minimize any effect on response. Such a protocol might have the destination device acknowledge every two packets and the source device pause only if four sends have happened without acknowledgment. In a well-balanced system using such a protocol, a file-transfer source may never wait. The considerations of archive access and retrieval times were not modelled since both LAN technologies would incur these delays if the same archive was used in each cases.

Figure 7 shows a simple picture of a single communication thread. A message only spends a part of its

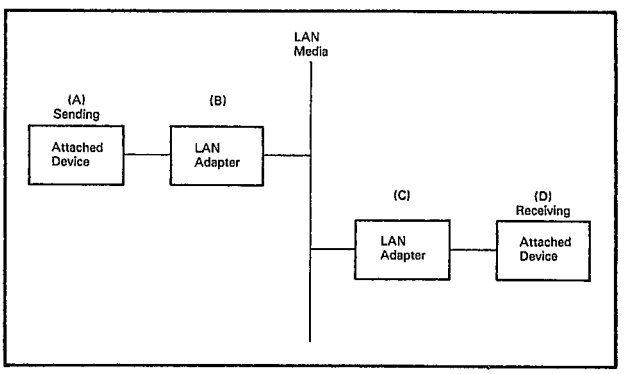

Figure 7: Simple Single Communication Thread

travel time being sent on the LAN's media (B-C). This is true whether this media can transmit 4 or 10 megabits per second. Messages must be built into one or more packets of the proper size and form, typically handled by the source device's processor (A). The packets must be moved from the attached device's memory into LAN adapter for transmission (A-B). Some cost may be paid for acquiring the right to transmit on the LAN (the media access delay). After receipt of the packet(s), the packet(s) must be moved to the attaching device's memory and recomposed into a message (C-D). Often messages have to be moved more than once at each end. Since data movement speeds inside a device (and between attaching devices and the LAN adapters) are typically in the same order of magnitude as the LAN transmission, poor internal device architectures can overshadow efficient LAN methods. We assumed the ideal case of no movement of packets in the attached devices -- that messages were packetized in place.

Hospital image generation and acquisition requirements are well documented [7] and realistic values have been used in these models. Source modality image generation rates are reduced to images per hour based on typical study rates and the number of images per study generated by each device. The frequency of image recall from archiving used in this simulation is also consistent with available literature. We picked numbers of terminals and rates to represent more than peak stress load.

To be useful the RESQ model had to address critical PACS issues, which include the following:

- o predicting response time at the workstations,
- o sizing the host (file-server) computer,
- o finding appropriate number of adapters on the host,
- o evaluating different strategies for host/adapter communication,
- o comparing results under different traffic assumptions,
- o comparing token ring and Ethernet.

These issues were all explored in our study. This paper will present some of the results in depth.

TOKEN RING VERSUS ETHERNET

For our assumptions and anticipated PACS of traffic, both the 4 mbps token ring and 10 mbps Ethernet LAN could achieve acceptable response times for both the modality image unloading and workstation image retrieval traffic under conservative assumptions.

In scenario 1 (figure 8) no modalities are included. Requests originate from each of 60 terminals twice a minute with a Poisson distribution. These requests were sent to the host; and responses of 1500 bytes, 1/6 Mbyte, or 1/2 Mbyte were returned with the probability of 0.5, 0.45, 0.05 respectively. This represents office traffic and two different size image files. Response time was defined to be from the start of the request from the terminal processor's memory until the entire message arrived at the terminal processor's memory.

Ethernet allows no more than 1500 bytes of data per packet. Larger messages must be packetized. Therefore, to satisfy a 150,000 byte request from a doctor at least 101 packets must be transfered. (One packet sent out to make the request, and 100 packets returned to satisfy the request). Inside the model, as in the real world, all delays, are for the transmission of packets, but we have to collect statistics on the transaction from the doctors' point of view. Therefore, each outgoing request packet gets a RESQ resource called, in this case, 'LIFE', which is returned with the file transfer packets. RESQ automatically provides detailed

statistics on the LIFE resource. For this application the token ring model also assumes packets of 1500 bytes of data. Actually the token ring protocol allows significantly larger packets. Since many of the delays in a LAN are associated only with the number of packets, the token ring protocol should be able to do somewhat better relative to Ethernet than this model demonstrated.

Scenario 1
Workstation Side of PACS

Mean/Maximum Time for Sending Request and Receipt of Image (seconds)

60 Workstations Each Sending:
- o one 1500 byte request
- o twice per minute (Poisson)
- o for 1500 byte response 50% of time
  1/6 mbyte response 45% of time
  1/2 mbyte response 5% of time

Assumptions:
- o 1500 maximum byte packets
- o no windowing or acknowledgements
- o best estimates delays & architectures

Results:

TOKEN RING	Media Speed	
Traffic	4 mbps	10 mbps
All	.46/2.28	.42/2.50
1500 byte	.24/2.02	.21/2.10
1/6 mbyte	.68/2.28	.53/2.50
1/2 mbyte	1.17/1.27	1.34/2.46
Utilization	41.3%	20.0%

ETHERNET	Adapter Speed	
	Today's	'89
Traffic	10 msec*	3.5msec
All	2.15/10.13	1.20/3.39
1500 byte	1.53/7.78	.92/2.94
1/6 mbyte	2.49/10.13	1.34/3.04
1/2 mbyte	4.80/9.14	2.43/3.39
Utilization	19.4%	19.7%

*three host/adapters were required to keep this case stable

Figure 8: Workstation Load Results

In scenario 2 (figures 9, 10) no workstations are included. The modalities are simply off-loading their image files to the host/file server. As a sensitivity test, the nominal, twice nominal, and five times nominal traffic are run. Comparison of the 10 mpbs Ethernet (with 3.5 ms adapter) and the 4 mbps token ring does not show major difference.

The reason for these results is that Ethernet can not effectively use its 10 mbps bandwidth because it is paced by the slower speed of its adapters. The PACS traffic predominately includes large file transfers. The nominal traffic rates for a PACS would often see the LAN idle when a file must be transfered. Since a file would be split up into many packets to prepare it for transmission and all of these packets would be sent single file through a single sending and a single receiving adapter, the file could be transferred no faster than these an adapter could fill or empty its buffers. This is demonstrated by figure 11 in which no significant delay is added by additional streams of traffic until the total transmission time of a pacekt, times the number of devices, gets in the range of the speed of the adapter. In effect, the adapters have picked time slots through collisions and backoffs, and are paced by adapter speeds. Faster (i.e. >10 mbps) speed token rings, though available, were not modelled. It appears that the LAN media is a less severe constraint than the device or device adapter capabilities (see figure 9). In fact, higher speed token rings (using available adapters) would simply increase the utilization of the media with no gain in throughput within the adapters. Image response times would be reduced by the improved LAN transmission interval.

Scenario 2
Modality Side of PACS with Token Ring

Nominal Traffic Rate: 13.4 Kbytes/sec.
Mean/Maximum Time to Offload to Host (seconds)
Nominal Traffic Mean Rate (Poisson):

	CT	MR	DF	DIG
Rate/Hour (Images):	30	30	8	10
Image Size (KBytes):	525	275	525	2K

Assumptions:
- o 1500 byte packets
- o no windowing or acknowledgements
- o FCFS to adapter; no priority
- o best estimates delays & architectures

Results:

4 mbps Ring	Nominal	Twice Nominal
Traffic	X1	X2
All	1.16/4.12	1.16/4.12
CT	1.16/2.06	1.16/2.11
MR	.57	.57
DF	1.09	1.09
DIG	4.12	4.12
Utilization	2.7%	5.3%

(For Comparison with Ethernet Only:)

10 mbps Ring	Nominal	Twice Nominal
Traffic	X1	X2
All	1.08/3.61	.89
CT	.95	.95
MR	.50	.50
DF	.95	.95
DIG	3.61	3.61
Utilization	1.2%	1.9%

Figure 9: Modalities Only (Token Ring)

Scenario 2
Modality Side of PACS with Ethernet
Nominal Traffice Rate: 13.4 Kbytes/sec.
Mean/Maximum Time to Offload to Host (seconds)

Nominal Traffice Mean Rate (Poisson):

	CT	MR	DF	DIG
Rate/Hour: (Images)	30	30	8	10
Image Size (KBytes):	525	275	525	2K

Assumptions:
- o 1500 byte maximum packet (part of Ethernet standard)
- o FCFS to adapter; no priority
- o no windowing or acknowledgements
- o best estimates delays & architectures

Results:

Current Adapter - 10 msec

	Nominal	Twice Nominal	5 Times Nominal
Traffic	X1	X2	X5
All	4.77/18.0	5.31/21.0	7.46/26.91
CT	4.28/7.04	4.83/9.96	7.42/15.61
MR	2.12/2.56	2.67/6.13	4.00/11.20
DF	4.73/7.04	4.73/7.04	8.88/15.88
DIG	15.99/18.0	16.77/21.0	19.90/26.91
Utilization	1.1%	2.2%	5.4%

'89 Adapter - 3.5 msec

	Nominal	Twice Nominal	5 Times Nominal
Traffic	X1	X2	X5
All	1.92/6.73	1.95/7.11	2.10/7.71
CT	1.70/2.24	1.75/2.80	1.93/3.17
MR	.88/.89	.89/.95	1.02/1.71
DF	1.88/2.49	1.88/2.49	1.84/2.49
DIG	6.42/6.73	6.49/7.10	6.67/7.71
Utilization	1.1%	2.2%	5.6%

Figure 10: Modalities Only (Ethernet)

Transaction Queueing Times (Seconds)
For 1-Way 525 Kbyte Transmissions

Normal Adapter - 10 msec	Mean	Maximum
1 Device	3.95	N/A
2 Devices	3.95	3.95
3 Devices	3.96	3.96
4 Devices	3.96	3.96
5 Devices	5.19	7.04
6 Devices	6.01	7.04

'89 Adapter - 3.5 msec	Mean	Maximum
1 Device	1.67	N/A
2 Devices	1.67	1.67
3 Devices	1.67	1.67
4 Devices	1.74	1.74
5 Devices	2.25	2.48
6 Devices	2.65	2.79

Assumes:
- o No windowing or acknowledgement
- o No application or I/O software delays
- o Single buffering in the adapter
- o Each device sending to distinct device

Figure 11: Simultaneous Modality Starts (Ethernet)

CONSIDERATIONS IN ANALYZING FILE TRANSFER TRAFFIC

In mixed file transfer traffic, sample statistics can be misleading. The LAN and adapters act on packets. A large message or file transfer must be broken into many packets. In these systems when one packet arrives at a server, many more are sure to follow. The last packet in a file transfer must wait until all the earlier packets are served. Thus in an otherwise idle system, a large file transfer will tend to have worse per packet mean queueing times than a small transfer, even if mean packet throughput is the same. In a system with a variety of different sizes of file transfers, sample statistics of average per-packet delay disguise the real phenomenon.

Therefore, in scenario 3 only one type of traffic was tried. Running the model without the modalities and having all the workstations requesting 1/6 mbyte images (scenario 3, figure 12) produces a 2.89 second delay for Ethernet. This is a transaction delay involving one request packet returning in response. The largest component of the delay was in the workstation I/O software (1.00 seconds of the 2.89 seconds). Packets arrived so fast that the workstation processor fell significantly behind in absorbing them. It is noteworthy that although each of the 120 packets in an image file spent about a second (on average) being accepted at the workstation, all 120 packets are absorbed during the 2.89 second end-to-end delay because of pipelining.

RESQ offers several techniques for generating confidence intervals that depend on assumed normal distribution for means. For this application the RESQ-generated response time confidence intervals are not useful. Most of the time LAN utilization is low and takes about the same amount of time for each occurrence on an idle LAN. Therefore, the minimum response time is also the mode of the response-time distribution. The simplistic assumption of normality often captures only 20% of the observed values in a 98% confidence interval. Also, in heavily utilized LANs there is a great deal of autocorrelation between two packets that follow each other. The assumption of independence can only be used if appropriately grouped mean statistics are used.

Scenario 3
Full PACS with Peak Modality Traffic
Nominal Traffic Rate: 13.4 Kbytes/sec.
Mean/Maximum Time to Offload to Host
(seconds)

Traffic includes:
- o 60 workstation traffic (figure 8)
- o Five times nominal modality traffic (figure 9 or 10)

Assumptions same as figures 8 - 10

Results:

TOKEN RING	Media Speed	
Traffic	4 mbps	10 mbps
All	.74/7.43	.40/3.61
CT	1.68/2.19	1.04/1.51
MR	.77/1.15	.55/1.03
DF	3.28/3.28	.99/1.12
DIG	7.43/7.43	3.61
1500 byte	.43/5.61	.17/1.27
1/6 mbyte	.90/4.87	.49/1.49
1/2 mbyte	2.00/5.81	1.08/1.36
Utilization	53.9%	21.6%

ETHERNET	Adapter Speed	
	Today's	'89
Traffic	10 msec*	3.5 msec
All	2.46/15.10	1.56/8.10
CT	4.23/5.02	2.00/2.43
MR	2.41/4.11	1.00/1.39
DF	3.98	1.86/1.86
DIG	15.09/15.10	6.63/6.67
1500 byte	1.63/10.74	1.17/7.40
1/6 mbyte	3.01/11.56	1.94/8.09
1/2 mbyte	5.20/11.32	2.32/3.43
Utilization	20.5%	20.2%

*three host/adapters were required to keep this case stable

Figure 12: Full PACS with Peak Modality Traffic

Trying to mix frequent and infrequent traffic requires long computer runs to get enough observations of the infrequent traffic to understand its effect. The workstations have on the average two transactions as second. The modalities average only slightly more than one event a minute. Further, the workstations are a uniform population, while the modalities have different characteristics. To get a significant number of modality events, one could run the full model for more than an hour of simulated time at a large cost in processor utilization. To get around the problem one can run the uncommon events separately to understand them (figures 8,9, 10) or run them at several times their normal rate with the full model since they are unlikely to occur coincidentally even at higher rates (figure 13). Neither of these techniques is completely satisfactory.

Workstations With Ethernet and Simplified Traffic

60 Work Stations Each Sending:
- 1500 byte request
- Twice per minute
- For 180,000 byte picture

10 Adapters Available At Host(s)

Results

Round Trip Time (W/O Application Delay): 2.89 Sec.

Includes (Per Packet):
- 27 msec at Host I/O Channel
- 37 msec at Host I/O Driver Software
- 50 msec at DLC
- 995 msec at Workstation I/O Driver

Figure 13: Workstation Traffic Only

Utilization of resources in a complex communication model needs close watching. In many parts of such a system there can be no queueing delay. For example, in the real world packets arriving at an adapter must be read immediately into an available buffer or be lost. Data loss does happen in the real world, thus knowing how many packets are dropped is useful. Assuming an infinite number of available buffers and calculating the distribution of buffer in use is a valuable way to help system architects pick the appropriate number of buffers for a system that is only in the planning stage.

CONCLUSION

Simulation showed itself to be a useful tool in analyzing the PACS architecture and traffic. We were able to show that given our assumptions either Ethernet or 4 mbps token ring could handle the short transaction traffic and large image data files characteristic of the asynchronous PACS environment. The LAN media (i.e. OSI layers 1-4) was shown to be far less critical to the message propagation delay than the device high level software, especially when accompanied by a disk access requirement. Modelling of the 10 mbps token ring illustrated the constraints being imposed by the device technology. Workstation architectures, internal buses, and software efficiencies became more significant issues in this environment. Wise PACS designers would be well advised to remove disk accesses and complex workstation processing functions from the critical timing path of the operator response times.

REFERENCES

1. Lam, S.S., "A Carrier Sense Multiple Access Protocol for Local Networks," Computer Networks, Vol. 4, pp. 21-32, 1980.

2. Bux, W., "Local-area Subnetworks: A Performance Comparison", IEEE Trans. Com. Vol. COM-29, No. 10, October 1981.

3. Coyle, E.J., B. Liu, "A Matrix Representation of CSMA/CD Networks", IEEE Transactions on Communication, Vol. COM-33, No. 1, January 1985.

4. Sauer, C.H., MacNair, E.A., Kurose, J.F., The Research Queueing Package Version 2, Introduction and Examples, Research Report RA138, IBM T.J. Watson Research Center, Yorktown Heights, New York 10598. April. 1982, 120 Pages.

5. Sauer, C.H., MacNair, E.A., Kurose, J.F., The Research Queueing Package Version 2, CMS User's Guide, (Research Report RA139), IBM T.J. Watson Research Center, Yorktown Heights, New York 10598. April 1982, 180 Pages.

6. Sauer, C.H., MacNair, E.A., Simulation of Computer Communication Systems, Prentice-Hall Inc., Englewood Cliffs, N.J., 1983.

7. Templeton, A.W., Dwyer, S.J. 3rd, Johnson, J.A., "An On-line Digital Image Management System," Radiology, August 1984, pp. 321-325.

AUTHORS

G. ROBERT LAWRENCE

G.R. Lawrence, is employed as an Advisory Engineer in the IBM Federal Systems Division in Gaithersburg, Maryland 20878. Mr. Lawrence has an educational background in Electrical Engineering (BS), Computer Science (MS), and Bio-Medical Engineering/Computer Science (Professional Degree of ENGINEER). He currently lectures graduate engineering courses (Comparative Computer Architectures) at George Washington University as an Associate Professorial Lecturer. His experience has been in large computer systems development, including multimodality PACS medical applications. Mr. Lawrence is a member of the IEEE Computer Society. He has published and presented papers in the United States and Europe, sat on the program committee of national expositions (e.g. CAEDEX '85) and participated on International Standards Committees (e.g. ACR/NEMA).

IBM Federal Systems Division
708 Quince Orchard Road
Gaithersburg, Maryland 20878
301-921-5702

GERALD A. MARIN

G.A. Marin is the Manager of the Network Systems Analysis Department in the IBM Federal Systems Division, Gaithersburg, Maryland. Dr. Marin's educational background includes a Ph.D. in Mathematics (Probability Theory), and both Master's and Bachelor's degrees in Mathematics. He has several years' experience in computer simulations, and his most recent work includes analyzing and simulating the performance of both packet switched radio networks and local area networks for IBM.

Prior to his current assignment Dr. Marin supervised a number of studies addressing U.S. Navy policy. He specialized in analyzing the performance of U.S. Navy ASW systems. Dr. Marin received the U.S. Navy's Distinguished Public Service Award in October 1978.

IBM Federal Systems Division
708 Quince Orchard Road
Gaithersburg, Maryland 20878
(301)921-5983

STEVEN E. NARON

S.E. Naron, is employed as a Systems Analyst in the IBM Federal System Division in Gaithersburg, Maryland U.S.A. 20878. Mr. Naron has an educational background in Operations Research (BES), Computer Science (MS), and Operations Analysis (MBA). He has wide experience in both evaluating and developing complex distributed information systems for research, governmental, and business applications. Mr. Naron is a member of IEEE, ACM, and SCS.

IBM Federal Systems Division
708 Quince Orchard Road
Gaithersburg, Maryland 20878
301-921-5642