UW PACS Prototype Performance Measurements, Computer Model, and Simulation

R. K. Panwar, C. S. Wang, L. A. DeSoto, H. W. Park, and Y. Kim

Image Computing Systems Laboratory, FT-10
Department of Electrical Engineering
University of Washington
Seattle, WA 98195

ABSTRACT

A PACS prototype has been installed and evaluated at the University of Washington. This paper presents the work done in the performance evaluation of the PACS prototype. The work involved network and workstation performance measurements and development of a simulation model based on the performance measurements. The simulation model was then used to do a parametric study of the PACS prototype to pinpoint the bottlenecks and suggest corrective measures. Results show that there are some local bottlenecks in the PACS prototype and an overall global bottleneck in the Data Management System (DMS) which forms the hub of the PACS prototype.

1. INTRODUCTION

Efforts have been made to develop simulation models of various PACS architectures [1-3]. The simulation models can be used to develop a systematic and objective means to propose and evaluate various PACS architectures, estimate the expected cost and performance, and fine-tune the proposed architecture. They can also be used to pinpoint the bottlenecks in an already existing system. While Bakker *et al.* [1] described a model that can be used for the quantification and evaluation of costs in the implementation of a PACS, the others [2-3] described models that help in the performance evaluation of PACS. At the University of Washington, we have developed a simulation model that can be used to evaluate the performance of the UW PACS prototype. The model was implemented using N.2 (Endot Inc., Cleveland, OH) which is a set of tools with features for hardware description and modeling and some features for stochastic modeling as well [4-7].

The PACS prototype system architecture at the University of Washington Hospital is shown in Figure 1. The PACS equipment installed consists largely of CommView equipment manufactured by AT&T and marketed jointly by Philips Medical Systems and AT&T. The CommView system is an active star-configured network with the Data Management System (DMS) as the central node. It supports image capture from X-ray CT (Computed Tomography), MRI (Magnetic Resonance Imaging), DSA (Digital Subtraction Angiography), PCR (Philips Computed Radiography) and laser digitizers. An Acquisition Module (AM1) digitizes the analog video signals from the CT, MRI and DSA into 512×512×8 images and transmits them to the Data Management System (DMS). The AM0 handles the images for the PCR modality. The DMS provides the central database management and transmission of patient images and demographic information. Up to 11 acquisition and display devices can be connected to the DMS. Images placed in the DMS can be displayed at various workstations featuring either single-screen Results Viewing Stations (RVSes), dual-screen Consultation Workstations (CW's) or quad-screen Display Workstations (DW's). The DMS has a prioritized service protocol and the priorities of the different nodes are listed in Table 1 with the highest priority node at the top and the lowest priority node at the bottom. Also, the University of Washington PACS has telecommunication links to the Harborview Medical Center, VA Hospital and Madigan US Army Hospital via a dedicated T1 line and the Sitka Community Hospital in Alaska via a 4.8 kbps telephone link.

The CommView PACS is based on UNIX with most CommView processes being daemons that remain in system memory permanently, occasionally being activated and then returning to a dormant state. The advantage of this

This work is supported in part by the U.S. Army Medical Research Acquisition Activity of the U.S. Army Medical Research and Development Command, Contract No. DAMD 17-86-C-6145. The views, opinions, and/or findings contained in this paper are those of the authors and should not be construed as official Department of Army policy or decision unless so designated by other documents.

approach is that it reduces process startup overhead and decreases user response times. The major disadvantage of this is that a significant amount of memory may be allocated to idle daemon processes. There is heavy software pipelining in the AM's and DW's. The processes in the pipelines communicate by means of shared memory or buffers. Each buffer has the capability of handling one image packet or SPF (Standard Picture Frame) which is 256 kbytes.

Priority	*Node*
0	DW2
1	CW0
2	CW1
3	CW2
4	CW3
5	AM0
6	DW0
7	DW1
8	RVS
9	AM1

Table 1. Node Priority List

This paper presents the work done in the performance evaluation of the commercially available PACS prototype. We first summarize the results from the network performance measurements and parameter estimates derived thereof. Then we present our simulation model and parametric studies. Finally, we present the conclusions that we were able to draw based on our performance measurements and simulation studies.

2. NETWORK PERFORMANCE MEASUREMENT

The PACS installed at the University of Washington Hospital is a closed system. No software tools were available to measure the various network and workstation performance parameters. Interaction with the system was limited to control over the communications links between the DMS and the various nodes. Estimates for the parameters were based on observance of status lamps on various modules and the measured times required to complete certain testing protocols.

2.1. Parameters in the Model

From our studies, we identified several parameters that would be critical in the development of the N.2 simulation model of the UW PACS. Our model uses a packet of 256 kbytes as the basic entity. The typical size of a PCR chest X-ray image is 18 packets while that of a video-captured X-ray CT image is 1 packet.

The deterministic parameters that were most pertinent to our needs were -

1. Image Acquisition Time: This is the time required to acquire and digitize an image packet at the AM. Our measurements showed that the PCR image acquisition time at AM0 for a typical 14"×17" chest X-ray plate was approximately 77 seconds. We assume that the time for image acquisition and digitization is 72 seconds. Given that there are 18 packets in a PCR image, this works out to an average time of 4 seconds for the acquisition of an image packet. For CT images, our measurements showed that the image acquisition time was 1.17 seconds per image. In our model we have assumed that the acquisition time for a CT image is 1.2 seconds.

2. Local Disk Access Time: This is the time required to read/write a packet at the local disks at the AM's or the DW's. Our model assumes that the local disk access time is 1 second/SPF which implies an effective disk transfer rate of 0.25 Mbytes/s.

3. Compression Time: This is the time to compress an image packet at either AM. In our simulation model we assume a compression time of 3 seconds.

4. Transmission Time: This is the time required for the transmission of an image packet between 2 nodes. We assume a transmission time of 0.3 seconds for a packet. Furthermore, we assume that each packet that is transmitted is a compressed packet and that the compression ratio is 3:1. Thus a transmission packet has $256/3 = 85.3$ kbytes of data. Thus a transmission time of 0.3 seconds implies a data throughput of 2.3 Mbps. We have found from our measurements that the application-to-application data throughput on the fiber optic link is about 1.9 Mbps. The T1 links are clocked at 1.344 Mbps though the potential maximum data rate is 1.544 Mbps. In our model, we assume a transmission time of 0.6 seconds for a packet on the T1 links.

5. Decompression Time: This is the time required to decompress an image packet. In our simulation model, we assume a decompression time of 3 seconds for a packet.

6. DMS Service Time: This is an aggregate time in the DMS for a packet. It is supposed to model the DMS task switch, the DMS database operations and the DMS hard disk access time. We assume that it takes the DMS 1 second to service a packet whether it be a packet sent from an AM or a packet requested by a DW.

7. Database Search Time: This is the time for a database search operation at the DMS when a radiologist at a DW requests a case. We found that the database search time is approximately 6 seconds.

8. Gateway to RVS Workstation Time: This time is a macro factor which is supposed to account for the image decompression at the RVS Gateway, the image transmission time from the RVS Gateway to the requesting RVS workstation, and the store/display of the image at the requesting workstation. We assume that it takes the RVS Gateway 5 minutes to service a "macro" workstation with the "macro" workstation modeling all the workstations that are attached to the gateway. This "macro" workstation generates requests at a uniform rate varying from 10 - 15 minutes.

Although we have modeled the above factors as deterministic parameters in our model, in reality they may show some statistical or other variation. For instance, the DMS Service Time is known to depend on the number of images on the disk. When the DMS disk is almost full, the DMS slows down. However, we have found that the DMS service time is approximately constant as long as the DMS disk utilization is below 90% of disk capacity.

In addition to the measurements done to determine the values of the various system parameters, we also applied testing protocols to determine the interaction between different nodes, the service protocol at the DMS and the scalability of the measurements to the examinations of varying sizes. Thus, for example, a testing protocol was applied to DW0 to determine whether a single data set could be scaled to examinations of varying sizes. In this test, CT examinations consisting of a varying number of images were requested at the DW and timed. The results are shown in Figure 2.

2.2. Multinode Measurements

Multinode measurements were conducted to determine the service protocol and performance of the DMS. The communications links were disabled for all nodes except those under test. Four CT examinations, each containing a set of 40 images, were designated as test image sets. An individual examination was assigned to DW0, DW1, CW1, and CW3. For each test protocol involving one or more of these nodes, the assigned image set was requested from the DMS. The testing protocol for the AM's was similar in that both AM links were disabled and images were queued on the AM local disks. If a particular test required one or both of the AM's, the appropriate links were enabled at the start of the test and the time required to flush the local database into the DMS was measured. This was done by measuring the time between enabling the communications link to the DMS and the end of memory/disk activity at the AM. Table 2 shows the measurement results.

The raw results were normalized with respect to the performance of each node when it alone was being serviced by the DMS. Figure 3 illustrates the normalized performance for each of the measured nodes as additional nodes are made active. As expected, the addition of various systems causes the expected performance at each node to be decreased. AM0 showed a 150% degradation for the case when there were 6 additional nodes compared to the baseline performance. The degradations of the other nodes varied in the range from 75-100% for the case when there were 6 additional nodes compared to the baseline performances. We also found that the system performance was considerably degraded when the RIS interface was switched on. The importance of RIS and its integration with the PACS has been well understood by PACS developers. Details of the integration and testing results are discussed elsewhere [8].

Event	*Node*	*SPFs*	*Time (s)*
1	CW1	40	215
2	DW0	40	193
	CW1	40	218
3	CW2		
	DW0	40	198
4	CW2		
	CW1	40	218
5	CW2		
	DW0	40	207
	DW1	40	193
6	AM0	63	229
	AM1	40	409
	DW0	40	275
	DW1	40	267
	CW3	40	294
7	AM0	64	309
	AM1	40	443
	DW0	40	279
	DW1	40	280
	CW3	40	295
	RVS		
8	CW2		
	AM0	68	233
	AM1	40	380
	DW0	40	252
	DW1	40	237
9	CW1	40	286
	DW0	40	245
	AM0	69	243
	AM1	40	394
10	CW2		
	AM0	80	322
	AM1	40	438
	DW0	40	292
	DW1	40	267
	CW3	40	300
11	CW2		
	AM0	65	308
	AM1	40	465
	DW0	40	347
	DW1	40	358
	CW1	40	379
	CW3	40	363

Table 2. Multinode Measurements

3. UW PACS SIMULATION

3.1. Other Parameters

Besides the deterministic parameters discussed in the previous section, our model has the following statistical parameters:

1. PCR Exam Schedule: This is the rate at which exams are scheduled at the PCR AM. In our model we have assumed it to be uniformly distributed from 5 - 7 minutes.

2. Number of PCR Images in a Case: We assume that the number of PCR images in a case is uniformly distributed from 1 - 2 images. Since there are 18 packets in a PCR image, we are going to assume that the number of packets in a case are also uniformly distributed. Actually the number of packets in a case has a comb-like distribution. But for for the sake of simplicity, let us assume a uniform distribution. These values associated with PCR images represent roughly a one third of the full load that the PCR and its AM can handle.

3. AM1 Exam Schedule: This is the rate at which exams are scheduled at AM1. This AM in the UW PACS handles 3 modalities - CT, MRI and DSA. For the sake of simplicity, we have only one modality, namely the X-ray CT. The CT exams are scheduled at uniform intervals that vary from 20 - 30 minutes.

4. Number of CT Images in a Case: We assume that the number of CT images in a case are uniformly distributed between 20 - 60 images.

5. Request Rate from CW1 and DW2: We assume that the CW1 and DW2 generate requests for service at uniform intervals from 15 - 25 minutes.

We have chosen the uniform distribution as an approximation to the distributions that arise in reality. The truncated Gaussian distribution could also have been used instead.

3.2. Description of the Model

A model should faithfully represent the system at an adequate level of detail. Levine *et al.* [3] have developed a model of a similar CommView system that has been installed at the Georgetown University Hospital. Their model is a high-level, less-detailed model as compared to ours. For instance, they have modeled each AM by two queues, an input queue and an output queue, while we have a detailed model of the actual transactions in each node. Figures 4 and 5 show some of the process interactions in our model. Thus, Figure 4 shows that our model uses the DW0Rad_Process to model the radiologist at the DW0 console. The radiologist generates a request for a test which is placed in queue for the DMSDBSearch_Process. This process signals the DW0Imagen_Process when it is finished with the database search. The DW0Imagen_Process is used merely to generate the appropriate number of images in a case according to the statistical distributions discussed earlier. The images are then passed onto the DMSDW0_Process which models the service of the DW0 node at the DMS. The service time slot is allocated by the DMS_Process. The DMSDW0_Process also models the image transmission from the DMS to DW0. The image, after it arrives at DW0, is decompressed by the DW0Decompress_Process and then written onto the local disk by the DW0LDWrite_Process. The DW0LDWrite_Process signals the DW0Rad_Process when the last image in a test case has arrived at the DW. The N.2 model uses an ISP'[1] process to model an actual process in the PACS [4-7]. Queues are used to model local storage devices (hard disks at the various nodes) as well as for process synchronization and interprocess communication. All the reads and writes to queues are synchronous (blocking). We shall not delve into the details of how each process is implemented. Suffice to say that there are processes to model the DMS, the AM's, the DW's and CW's at a fairly detailed level and processes to model the RVS Gateway at a description which is not so detailed.

3.3. Parametric Study

The main objective of the parametric studies was to identify possible bottlenecks in the network and suggest remedial measures to alleviate the bottlenecks in the overall system performance. Similar parametric studies have been done, for example, by DeSilva *et al.* [2]. Their study has two shortcomings compared to ours: (1) They have made optimistic estimates for various parameters. For instance, they assume that the DBMS can be modeled as a delay in which it takes 100 ms to locate an image. From our measurements, we found that the time to locate an image in the DMS was at least 6 seconds, and (2) they assumed that the buffers are of infinite sizes. This is a good starting point, but a rather naive and impractical assumption to obtain more reliable simulation results. The main parameters of the network were identified in the previous section. Out of those parameters, we identified the following as the variable parameters of interest:

1. TAM0Acquire: Time to acquire a PCR image at AM0. 4 s/packet.
2. TAM1Acquire: Time to acquire a CT image at AM1. 1.2 s/packet.
3. TAMCompress: Time to compress an image packet at the AM's. 3 s/packet.
4. TAMLDAcc: Time to access (read/write a packet to) the AM local disk. 1 s/packet.
5. TDMSService: Time for the DMS to service 1 packet. 1 s/packet.
6. TDWDecompress: Time for decompressing a packet at the DW. 3 s/packet.
7. TDWWrite: Time for writing a packet at DW local disk. 1 s/packet.
8. TXmit: Time for packet transmission on fiber optic link. 0.3 s/packet.
9. T1Xmit: Time for packet transmission on T1 link. 0.6 s/packet.

[1]ISP' is derived from ISP (Instruction Set Processor). ISP is a hardware description language.

To study the effect of a parameter on the system, we hold all the other parameters at their default values and vary the parameter under investigation from its default value to a value which is about 10 times less (better) than the default. We take 6 samples in the interval. The motivation is that the default value represents the state-of-the-art in PACS technology when the CommView system prototype was built in the mid-1980s. As the technology advances over the next five years or so, it is reasonable to assume that the performance of system components will improve by a factor of 10.

Table 3 shows how the parameters were varied in the different simulation runs. Since there are 8 relevant parameters (TAM0Acquire, TAM1Acquire, TAMCompress, TAMLDAcc, TDMSService, TDWDecompress, TDWWrite, TXmit) for Study 1 (the T1Xmit parameter being irrelevant since all the links that are operational are fiber links, not T1 links), and each parametric variation involves 6 sample points, we have a total of $8 \times 6 = 48$ simulation runs for this study. The output parameters of interest in each simulation run are the total average service time for a case at each node. Thus the service time at the DW would be measured from the time the radiologist made a request for a case till all the images for the case have been displayed on the DW. Similarly, the total service time for an AM is the time it takes to acquire, digitize, compress and store the compressed images on the DMS. The literals N and F in parenthesis beside the nodes CW1 and CW2 indicate whether the CW's are operating under a normal load or a full load. Normal load implies that the CW's request cases at normal (stochastic) rates while a full load implies that the CW's generate requests continuously as soon as the previous request has been satisfied. The DW's are also operated at full load for all the studies. This enables us to know what the maximum throughput of the system will be. The AM's however run at normal loads.

Study	*Enabled Nodes*	*Total Runs*
1	AM0,AM1,DW0	48
2	AM0,AM1,DW0,DW1	48
3	AM0,AM1,DW0,DW1,CW1(N)	48
4	AM0,AM1,DW0,DW1,CW1(F)	48
5	AM0,AM1,DW0,DW1,CW1(N),CW2(N)	54
6	AM0,AM1,DW0,DW1,CW1(F),CW2(F)	54
7	AM0,AM1,DW0,DW1,RVS	54

Table 3. Events of PACS Modeling

3.4. Results

Table 4 shows the mean service times predicted by our model for various nodes. The first row is the baseline times, i.e., the service time for a node when only that particular node is up. The second row gives the times when 3 nodes (AM0, AM1, and DW0) were up in the case of Study 1.

AM0 Service Time	*AM1 Service Time*	*DW0 Service Time*
259.6 s	168.0 s	171.0 s
259.8 s	184.6 s	171.0 s

Table 4

Note that the degradations (increase in service times) are different from the degradations found in the measurements. This is because in the case of the measurements, our procedure for multinode system studies involved acquiring all the images at the AM's and then giving the "GO" signal to all the nodes. Thus the AM images were already stored at the AM local disks prior to the study. This is not true in a real system where images are acquired on the fly. The simulations show that there is no appreciable degradation in either the DW service time or the AM0 service time. This indicates that the DMS is not a bottleneck for either of the DW's or the AM0. However the DMS is a bottleneck for AM1 since there is some degradation for AM1. What this effectively means is that the AM1 hardware throughput is greater than what the DMS can handle when the DMS is servicing 3 nodes.

Figure 6 shows some of the simulation results. The results show the dependence of the service times on a specific parameter. The unit on the Y axis (i.e., for the service time) are seconds in real world time while the units on the X axis (the variable parameter of interest) are simulation ticks with 10 ticks representing 1 second. Thus, Study 1 involving the parameter TAMLDAcc, shows that the service times for DW0 and AM0 are independent of the AM local disk access time. However, the AM1 service time has a slight dependence on TAMLDAcc. The service time for AM1 shows a lot of variation (has a large standard deviation) because AM1 has the lowest priority and the major bottleneck for AM1 is the DMS. The DMS services the stochastically generated requests from various nodes in a prioritized order, and hence the standard deviation of the service time (in the DMS) for the node of lowest priority will be the largest. Study 1 involving the parameter TDWDecompress, shows that the DW0 service time has a linear dependence on the decompression time in the DW, but the service times for the AM's are unaffected by this factor. Study 2 involving TDWWrite, shows that the DW service times are only slightly affected by this factor. Study 6 involving TDMSService, shows that the service times for all nodes except AM0 are dependent on the DMS service time. The results from our studies are summarized below:

The AM0 service time had a linear dependence on the AM0 acquisition time. The AM1 service time was not at all dependent on the AM1 acquisition time. This implies that in the AM's, the acquisition time for CT images was fast enough for it not to be a bottleneck, but the PCR image acquisition time was a bottleneck. None of the service times showed any dependence on the transmission time TXmit. This is similar to the result obtained by DeSilva *et al.* [2] for the parametric study of the compression ratio factor. They found that the compression ratio was not a significant factor if the bus speed was high enough. It is intuitive that the compression ratio will be important if the transmission time is a bottleneck. Improving the compression time in the AM's by a factor of 6 led to an improvement of the service times in the AM's by 20% and 35% in AM0 and AM1 respectively. The dependence of the AM service time on the compression time was not as strong in AM0 because the acquisition time was the primary bottleneck. Improving the DMS service time by a factor of 10 for the case of 3 nodes running simultaneously did not have any noticeable effect on the service times. However, with 4 nodes running simultaneously, improving the DMS service time by the same factor led to an improvement of the service times by about 4% in the DW's but no noticeable improvement in the AM's. With 6 nodes running simultaneously, improving the DMS service time by the same factor (as shown in Study 6, Figure 6) led to an improvement of the service times by 6%, 8%, and 33% for DW0, DW1, and AM1 respectively. This implies that the DMS becomes a bottleneck if 4 or more nodes need service simultaneously. With 3 nodes running simultaneously, improving the DW decompression time by a factor of 10, led to an improvement of the DW service time by approximately 60% (as shown in Study 1, Figure 6). Hence the decompression was a bottleneck in the DW's. However, with 6 nodes running simultaneously, improving the DW decompression time by the same factor led to only a 6% improvement in the DW service time. This shows that with more nodes, the local bottlenecks have a lesser effect because of the overall global bottleneck – the DMS itself. The DW service time had a minor dependence on the DW disk access time with less than 4 nodes running. Adding the RVS Gateway to a DMS with the two DW's and the two AM's did not affect the system bottlenecks in any way.

4. CONCLUSIONS

We have completed the evaluation of the commercially available PACS prototype at the University of Washington. Based on our measurements we developed a simulation model which we implemented using N.2. The simulation model is a highly detailed model which can be used to pinpoint the exact bottlenecks in the system. The basic entity for transactions in the simulation model is an image packet which is also the entity used within the system itself. We used the simulation model to do parametric studies for the system. The conclusion drawn from the parametric studies is that the DMS becomes a bottleneck if more than 4 nodes are up. This was an expected result in an active star topology since star architectures are not easily scalable. Besides, the central bottleneck in the DMS, there were some local bottlenecks, most notably the compression and decompression times in the various nodes. Most of the local bottlenecks disappeared or had a lesser effect when more nodes were running since the DMS becomes the major bottleneck. To obtain the full benefits of image compression, the performance of the compression and decompression hardware has to be improved substantially because the transmission time between the DMS and the DW's or AM's is not a bottleneck. Currently, the only benefit of image compression is a reduction of central storage requirement and image handling time within the DMS.

We also found that the measurements alone were not a good predictor of the system degradation for multinode studies since during multinode measurements the images were acquired, stored on the disks and then the transmission requests were generated simultaneously at each node. However, the measurements served to provide the parameter values for our simulation model. The only valid conclusion (regarding the system performance) that we can draw from our measurements is that the Alaska modem does not have any detrimental effect on the system performance even though it has a reasonably high priority. This is because the rate at which a packet arrives at the modem (at 4.8 kbps) is slow enough not to affect the rest of the system.

Presently, the overall system does not meet the PACS clinical requirements because of various limitations of the system and its architecture, local bottlenecks like the decompression time at the DW's, and the overall global bottleneck which is the DMS itself even though there are some components of the system that are clinically useful and widely used, e.g., the PCR. However, removal of the bottlenecks discussed in this paper, intelligent software (e.g., image preloading into the local workstations with fast parallel transfer disks) and better system integration will make the system clinically more acceptable in the future. Future PACS developers should build detailed simulation models of viable architectures and do detailed parametric studies to arrive at an optimal PACS configuration.

5. REFERENCES

1. Bakker, A. R., Stut Jr., W. J. J., DeValk, J. P. J., and Reijns, G. L., "PACS Costs: Modeling and Simulation", *SPIE Medical Imaging II*, Vol. 914, pp. 1147 - 1151, 1988.

2. DeSilva, V., Perros, H., and Stockbridge, C., "A Simulation Model of a Picture Archival and Communication System", *SPIE Medical Imaging II*, Vol. 914, pp. 940 - 946, 1988.

3. Levine, B., Randall, K., Mun, S. K., Benson, H., Meissner, M., and Sherman, A. B., "Optimization of Picture Archiving and Communication Systems Network Performance: Simulation and Validation", *SPIE Medical Imaging III*, Vol. 1093, pp. 330 - 339, 1989.

4. N.2 ISP' User's Manual. Document # 101, Version 1.14. Endot Inc., Cleveland, Ohio, 1987.

5. N.2 Ecologist User's Manual. Document # 105, Version 1.08. Endot Inc., Cleveland, Ohio, 1987.

6. N.2 Simulator User's Manual. Document # 106, Version 1.21. Endot Inc., Cleveland, Ohio, 1988.

7. N.2 General Purpose Post Processor User's Manual. Document # 111, Version 2.2. Endot Inc., Cleveland, Ohio, 1988.

8. Lee, W., Rowberg, A. H., O'Leary, Y. M., and Kim, Y., "Integrating a Radiology Information System with a Picture Archiving and Communications System", *SPIE Medical Imaging IV*, Vol. 1234, in press, 1990.

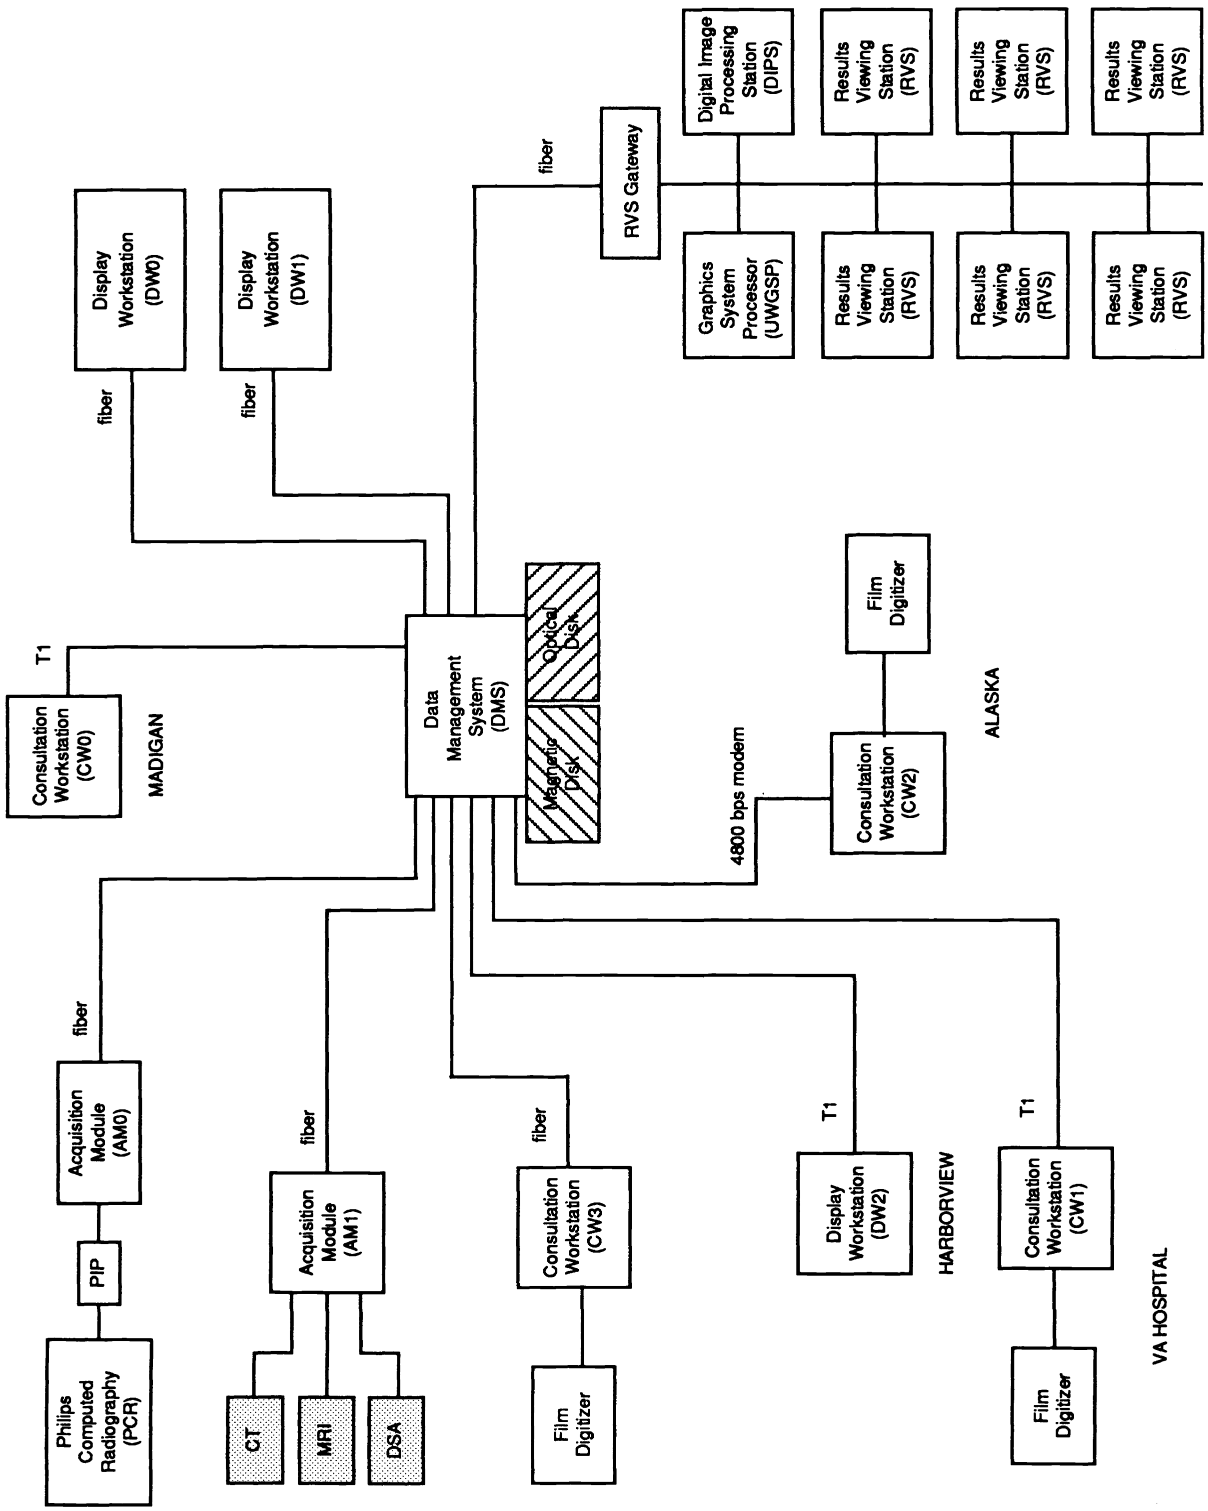

Figure 1 UW PACS BLOCK DIAGRAM

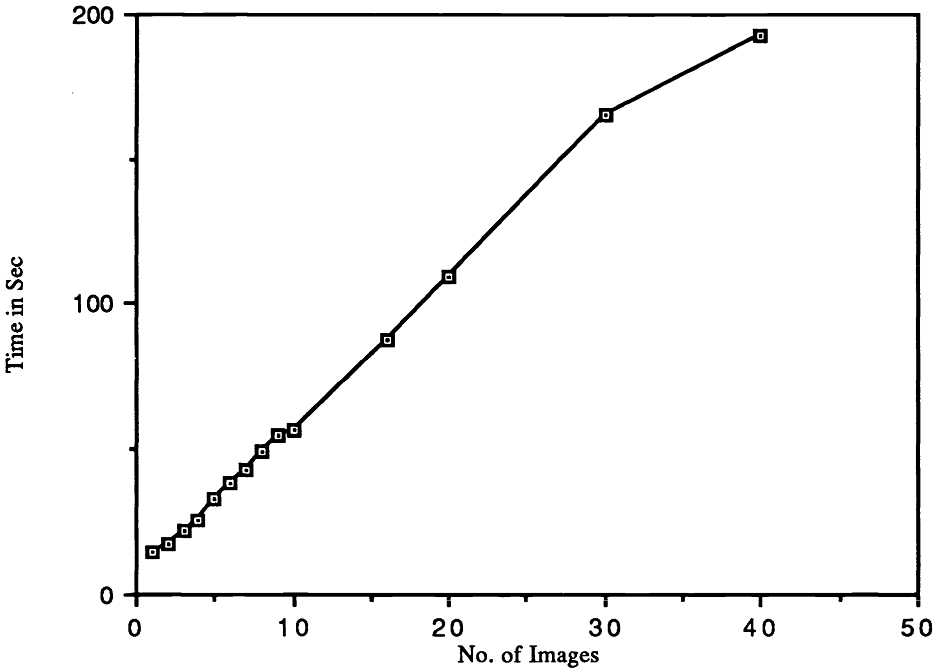

Figure 2 CT Image Transmission Time from AM1 to DMS

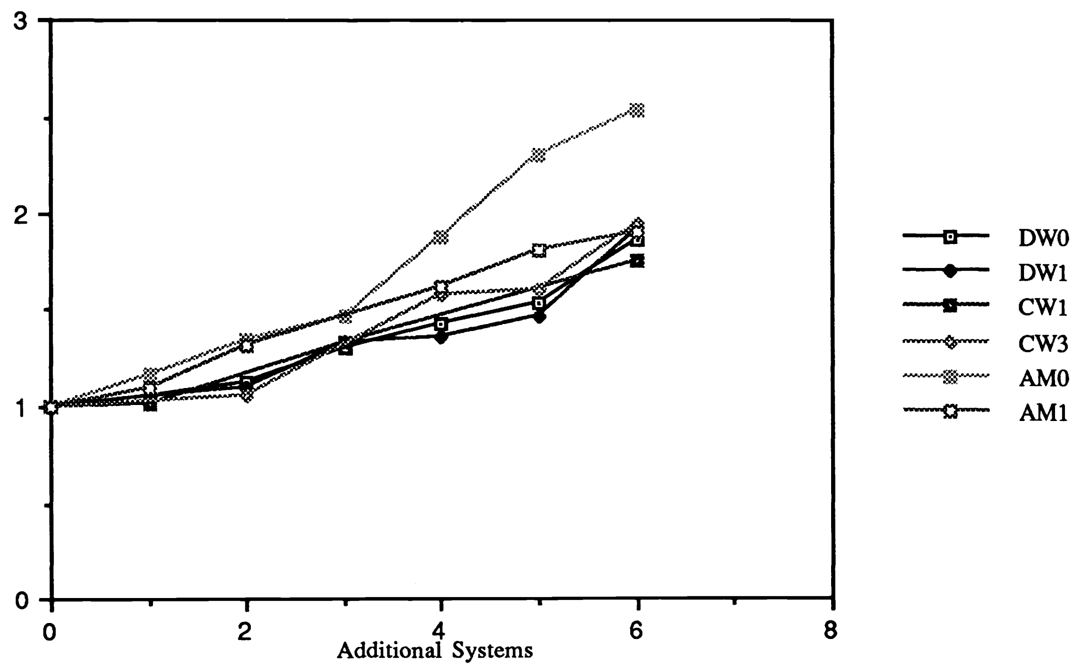

Figure 3 Expected Increase in Service Time as a Function of Additional Systems

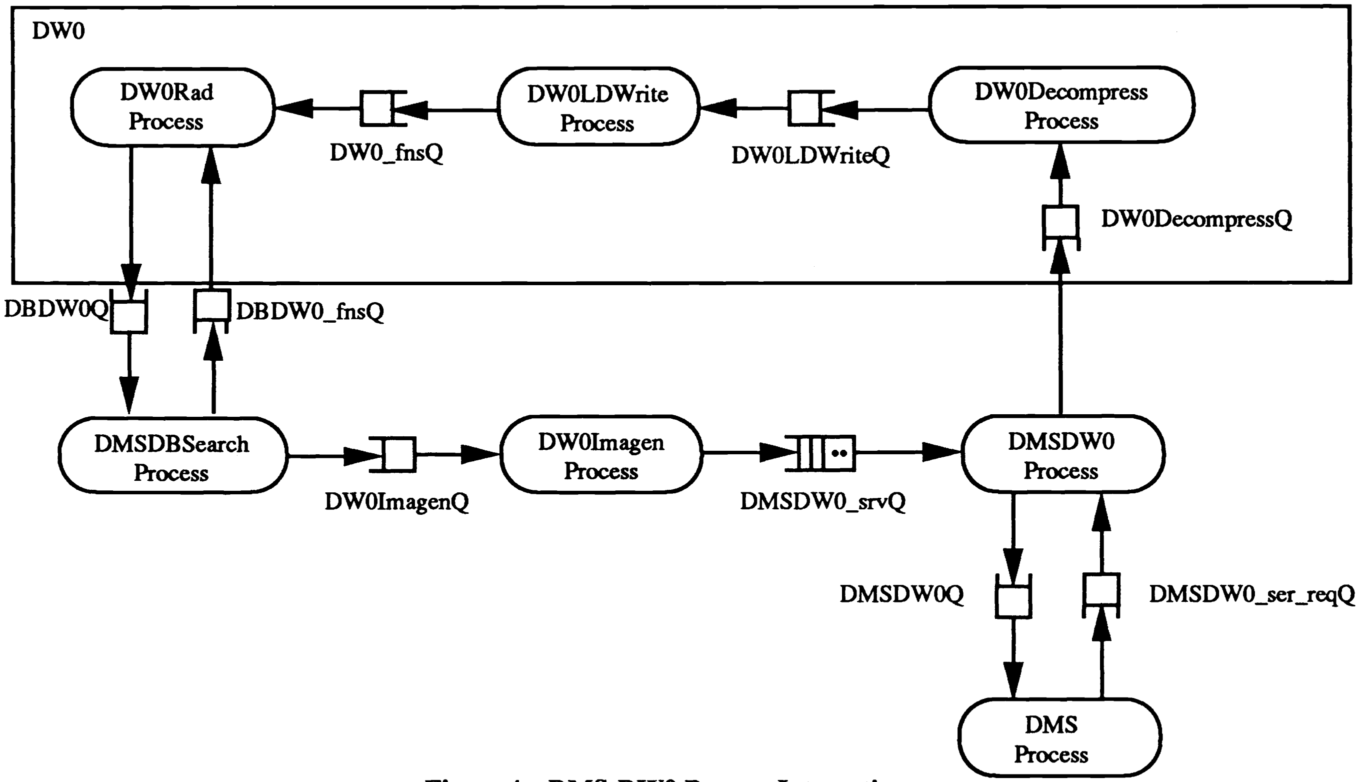

Figure 4 DMS-DW0 Process Interactions

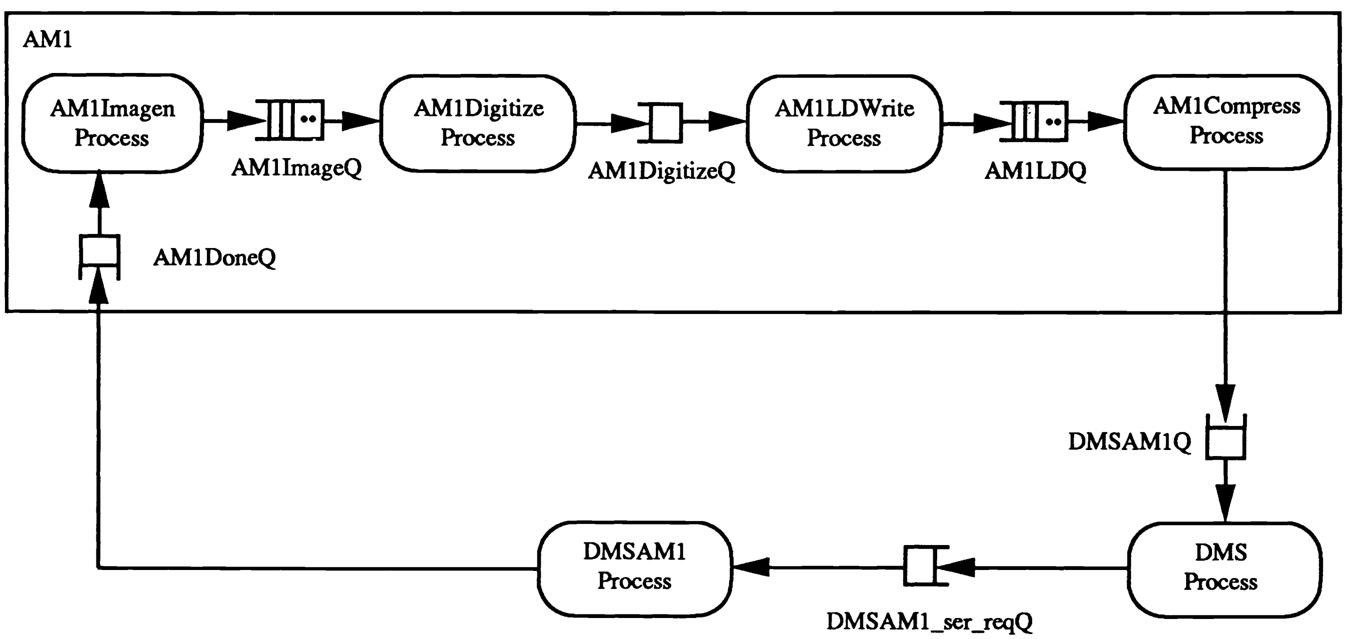

Figure 5 DMS-AM1 Process Interactions

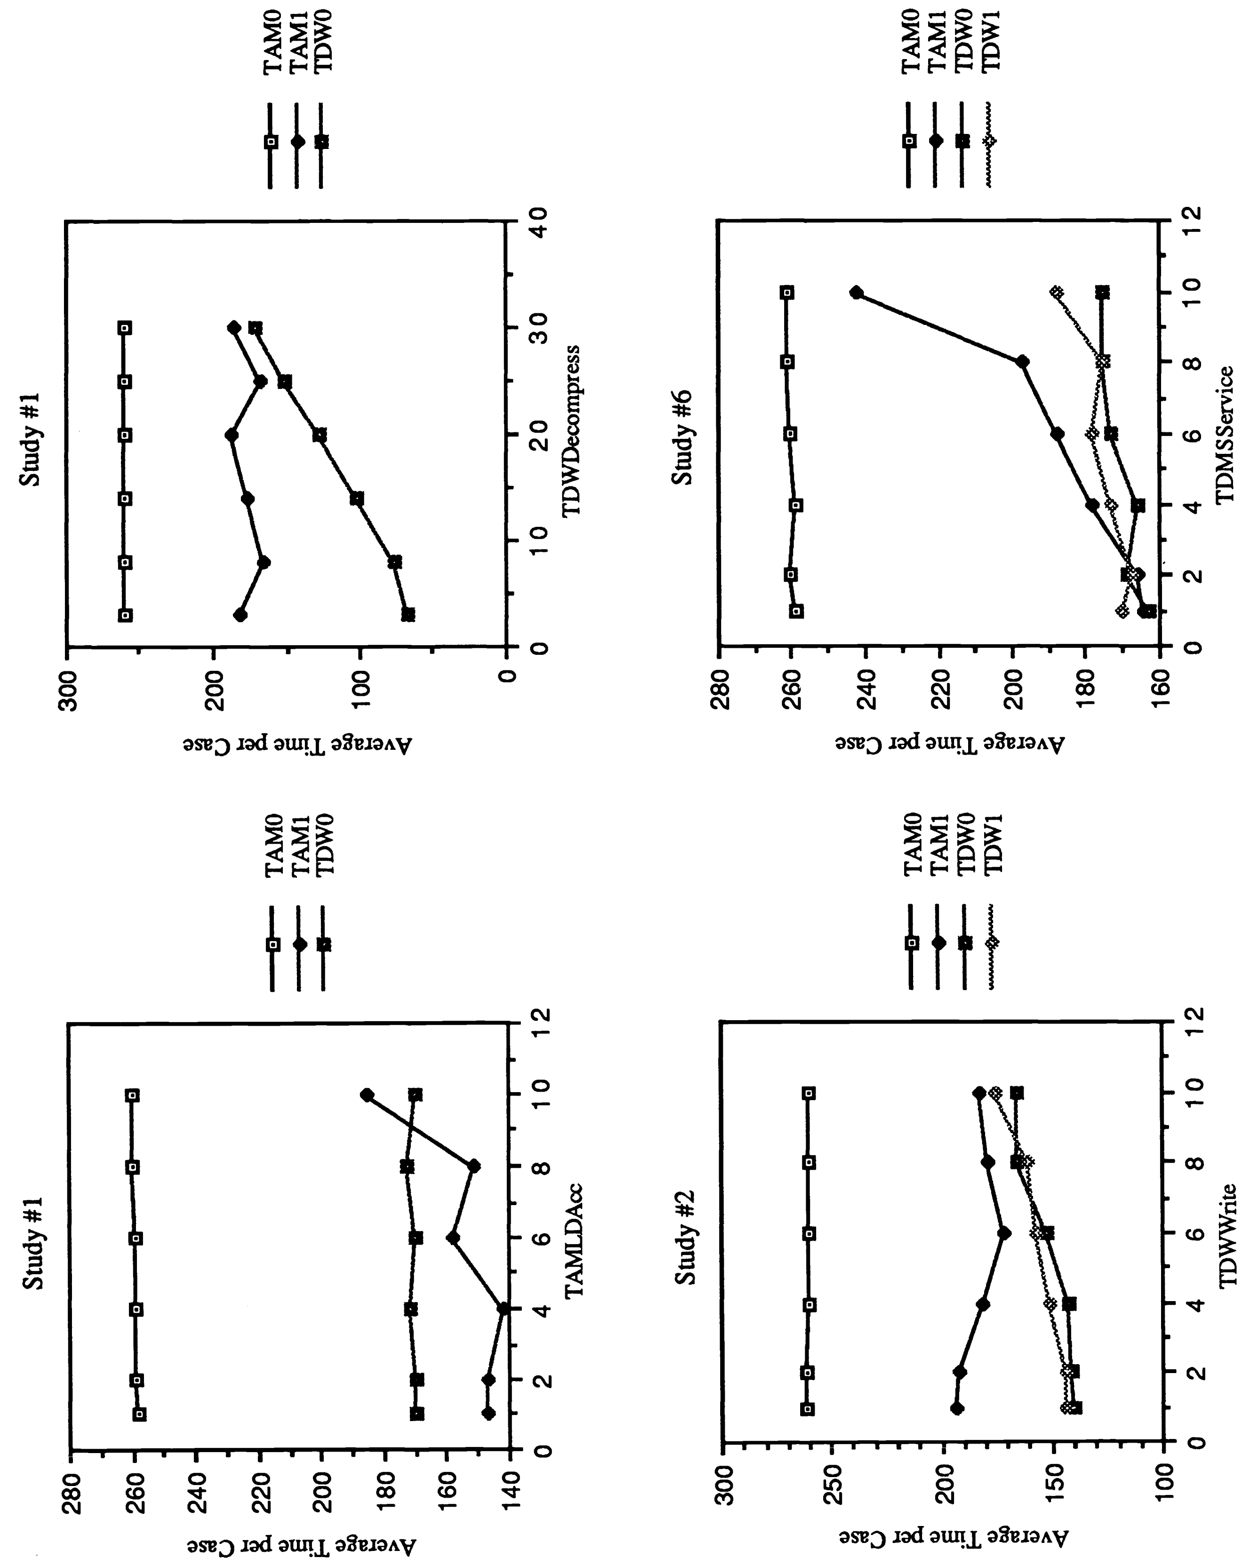

Figure 6 Sample Simulation Results