

W. H. Sanders, R. Martinez, Y. Alsafadi, and J. Nam, "Performance Evaluation of a Picture Archiving and Communication Network Using Stochastic Activity Networks," to appear in IEEE Transactions on Medical Imaging, March 1993.

PERFORMANCE EVALUATION OF A PICTURE ARCHIVING AND COMMUNICATION NETWORK USING STOCHASTIC ACTIVITY NETWORKS*

W. H. Sanders, R. Martinez, Y. Alsafadi, and J. Nam

Computer Engineering Research Laboratory
Department of Electrical and Computer Engineering
The University of Arizona
Tucson, AZ 85721 USA

(602) 621-6181
whs@ece.arizona.edu

ABSTRACT

The concept of picture archiving and communication systems (PACS) is now widely accepted in the medical community. In order to bring the concept to reality, however, innovative designs and implementations are needed. One such design is a fiber optic star based PACS. This PACS network is based on a multiplexed passive star local area network with wavelength-division multiplexing to provide separate logical channels for transfer of control and image data. The system consists of an *image network* (INET), for image transfer at a rate of 140 Mbps, and a *control network* (CNET), operating at 10 Mbps, for mediating the flow of image transfers. INET is a circuit switched network devoted solely to image transfer, while CNET employs the CSMA/CD protocol for bus arbitration. Before such a system can be deployed, an accurate evaluation study must be carried out to estimate its performance characteristics. Such evaluations are complicated both by the complexity of the PACS itself and the varied demands that are placed on such a system. A novel approach based on *stochastic activity networks*, a stochastic extension of Petri nets, is useful in this regard. Stochastic activity networks were used to develop a detailed model of the command and image channels. The performance of the system was then evaluated under realistic workload conditions. In particular, we were able to estimate a number of important performance variables including the image response time, command channel delay, and queue length at each type of node and the network supervisor. The results 1) show that stochastic activity networks are an appropriate model type for evaluating picture archiving and communication systems, 2) delineate the workload conditions under which PACS may effectively operate, and 3) show that even when these conditions are exceeded, the command channel load remains extremely light. Results of this type are useful both to designers of other PACS networks and those interested in this particular PACS design.

Keywords: picture archiving and communication systems, performance evaluation, local area networks, stochastic activity networks

*This work was supported in part by a grant from the Medical Systems Division, Toshiba Corporation and by the Digital Equipment Corporation Faculty Program: Incentives for Excellence.

I Introduction

The concept of *picture archiving and communication systems* (PACS) is not new. In theory, a PACS will operate in a hospital and provide a totally automated method for storing and retrieving medical images, diagnostic reports, and patient data. Early research in the design of PACS networks considered the use of off-the-shelf local area networks, such as Ethernet [1, 2]. While these early designs illustrated the utility of the PACS concept, they also pointed out the limitations of using traditional coaxial-cable LANs as PACS networks. Specifically, the need for extremely fast transfer of large digitized images made optical fiber a natural choice for the physical layer [3]. Different topologies, such as the ring, bus, and star, and media access schemes, such as CSMA/CD and the token passing method, were also investigated [4]. While these general purpose media access schemes work well for many applications, they do not seem to be well suited to PACS networks. In particular, the inherent conflict nature of CSMA/CD makes it difficult to meet response time constraints for image transfer while still delivering patient and control information in a timely fashion [5]. Similarly, the overhead associated with token passing precludes using this method, without modification, as a media access scheme.

Accordingly, it is appropriate to investigate the development of special purpose media access control protocols for PACS networks. Much work in this regard has been done at the University of Arizona [6]. For example, traffic patterns of images within the radiology department have been analyzed [7]. This work resulted in the observation that there are three types of information data that need to be considered: patient data, control information, and the images themselves. The three types are distinguished by the relative differences in the size of the data units that must be transmitted. Specifically, digitized images are very large (possibly on the order of 50 Mbits), while patient data and control information are relatively small (on the order of 1 Kbyte). This suggests a two channel approach to network design, so that a large image transfer does not interfere with the timely transfer of patient or control information. Using this approach, the first channel could be reserved for image transfers and utilize a circuit switching media access method, while the second channel could be for the transmission of patient data and control information and utilize a packet switching approach [8]. At the physical layer, either a star or ring topology could be used.

While this approach seems reasonable, an evaluation is necessary to determine, for a

given network design, what range of workload demands will result in acceptable performance. This performance is not simple to compute and depends on many factors, including the details of the protocols that govern each channel and the number and type of nodes in the network, among other factors. This paper reports the results of an evaluation of a PACS based on the star topology. The results are significant for two reasons; 1) they illustrate the suitability of a particular model type, “stochastic activity networks,” for evaluating PACS, and 2) they provide new and useful information concerning the use of star topology for a two channel PACS. Both of these issues will be addressed in detail in Sections V and VI.

The remainder of the paper is organized as follows. First, a general description of a fiber optics star based PACS and its components is presented to acquaint the reader with the particular type of network designs considered. Next, the modeling approach used is reviewed. The approach makes use of a particular extension to Petri nets known as “stochastic activity networks” (SANs), which are particularly suitable for representing the complex network designs inherent to PACS. This section is followed by a discussion of the assumptions made in constructing the model, the workload assumed, and a description of the SAN models of each of the PACS components. Next is a description of the performance measures considered. These measures include the total expected response time to deliver an image, the utilization of each channel, and the fraction of time during which collisions are present on the control channel. Finally, we present the results of the solution of the model for these variables via simulation, and offer some conclusions regarding both the appropriateness of SANs for evaluating PACS and the performance of this specific PACS design under realistic workload conditions.

II A Fiber Optics Star Based PACS Design

As outlined in the introduction, we consider a PACS with fiber optic star topology network using a passive star coupler [9]. The network consists of two sub-networks, an *image network* (INET) and a *control network* (CNET), implemented over the same physical optical fiber by using wavelength multiplexing technology. Images are transferred over INET, a circuit switched channel, at a rate of 140 Mbps, due to the available network hardware. Control and patient information are transferred over CNET, a CSMA/CD network operating at 10 Mbps. A network interface unit (NIU) is used to connect viewing workstations, imaging equipment and database archive systems to the radial fibers of the star network

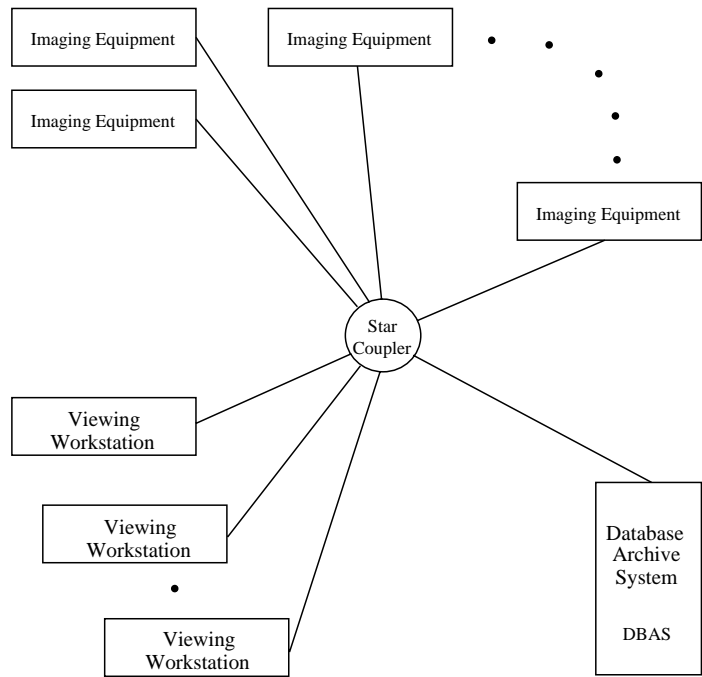


Figure 1: Star System Topology in Radiology Department

[10]. Requests to transfer images (via INET) are transmitted over CNET to the network manager, which is physically located with the database archive. Such requests are served in a first-come, first-served manner, as the image channel becomes available. After INET is acquired, an image or set of images (called a *folder*, in the following) is sent as a single entity, with no packetization or fragmentation.

The PACS we consider consists of three types of network nodes; *imaging equipment* (IE), *viewing workstations* (WS), and *database archive systems* (DBAS) (see Figure 1). Imaging equipment are image acquisition devices which acquire digital patient images. These imaging modalities include: computed tomography, nuclear medicine gamma cameras, magnetic resonance imaging, digital radiography, digital subtraction angiography, ultrasound scanners, and conventional x-ray machines. These types of imaging equipment differ in their characteristics and features, and hence generate different data formats which must be accounted for in the evaluation process.

A viewing workstation provides a viewing facility for the radiologist and referring physicians to review images acquired earlier and stored in a PACS DBAS. In order to ensure that

the user does not experience undue delay, viewing workstations must have high performance CPUs, or array processors and ultra fast memory systems [11, 12]. The data base archiving system (DBAS) stores acquired images and patient data related directly to the image. A radiologist at a viewing workstation requests image retrieval from the DBAS. Image storage and retrieval requests contain a mix of image and text data. Details regarding each component, together with the assumptions used in the modeling, will be given after the model type used is described.

III Model Representation Using SANs

Traditional methods, based on queueing networks, do not suffice for evaluating PACS network designs due to their inability to handle the complex interactions present in PACS protocols. Stochastic extensions to Petri nets [13, 14], on the other hand, are suitable for such systems due to “lower-level” model components. One stochastic extension to Petri nets, “stochastic activity networks,” is particularly suitable in this regard. They have been used to evaluate a wide range of systems ranging from a self-checking, self-correcting memory system [15] to an IEEE 802.4 token bus network for factory applications [16]. Stochastic activity networks are also well suited for PACS evaluation, and hence used in this study. While no attempt is made to fully describe stochastic activity networks here (the reader is urged to consult the relevant references for more information), we briefly review their basic aspects to aid in understanding the models developed in the next section.

Stochastic activity networks (SANs) [17, 18, 19] incorporate features of both stochastic Petri nets and queueing models. Structurally, SANs have primitives consisting of *activities*, *places*, *input gates*, and *output gates*. Activities (“transitions” in Petri net terminology [20]) are of two types: *timed* (e.g. *arr* and *send1* in Figure 2) and *instantaneous* (e.g. *sw* and *jamq* in Figure 2). Timed activities represent activities of the modeled system whose durations impact the system’s ability to perform. Instantaneous activities, on the other hand, represent system activities which, relative to the performance variable in question, complete in a negligible amount of time. Cases associated with activities (e.g. the two hollow dots on the activity *sw*) permit the realization of two types of spatial uncertainty. Uncertainty about which activities are enabled in a certain state is realized by cases associated with intervening instantaneous activities. Uncertainty about the next state assumed upon completion of a timed activity is realized by cases associated with that activity.

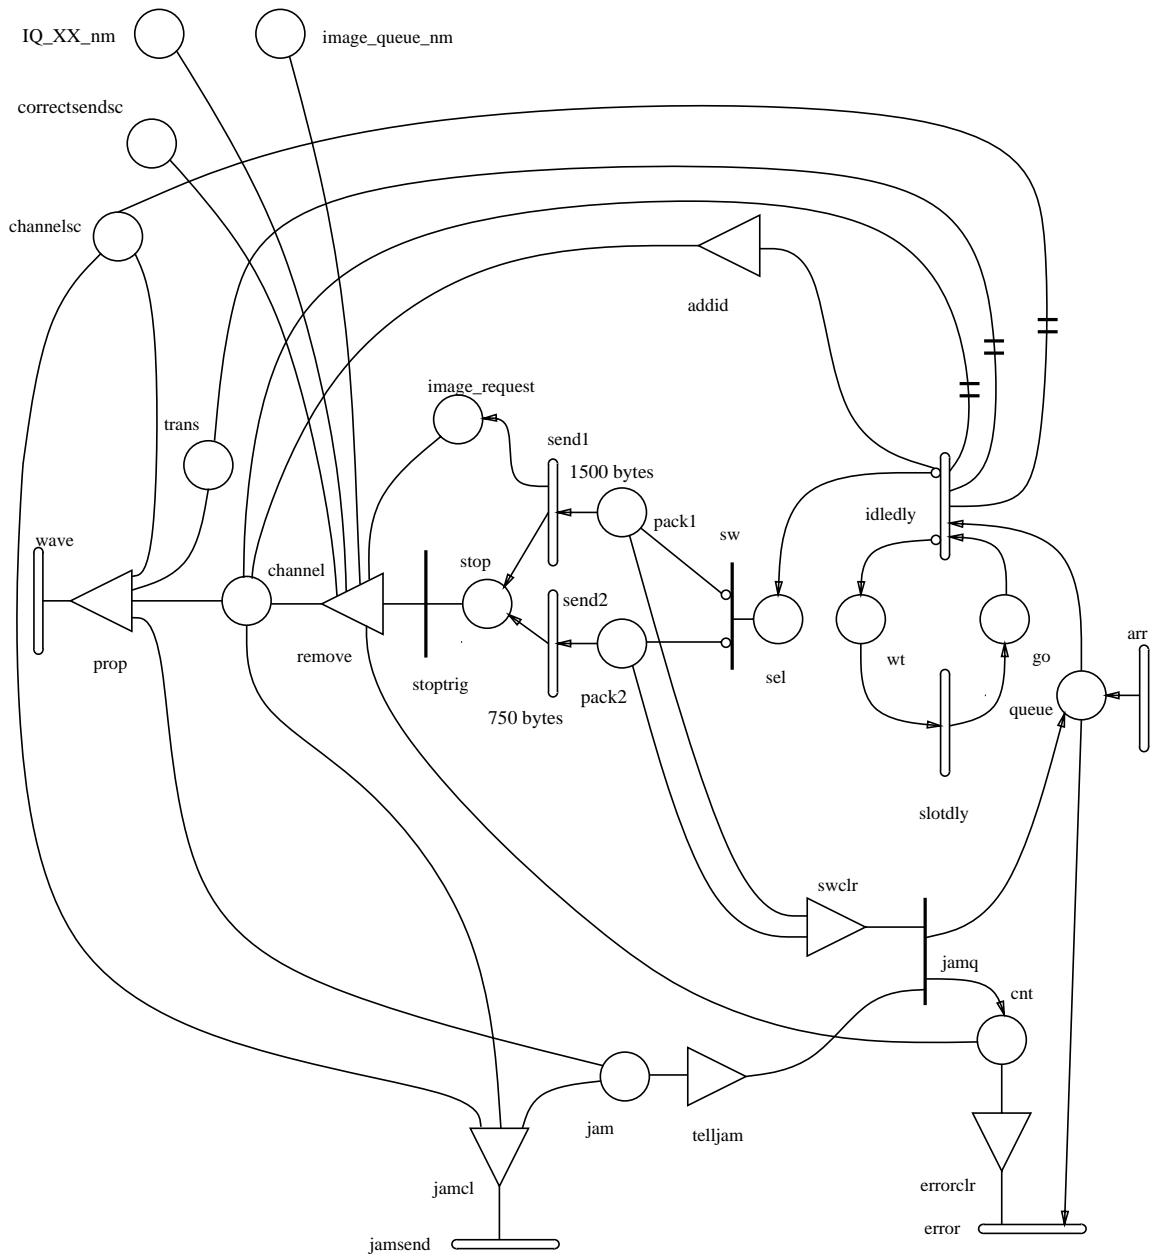


Figure 2: SAN Model of Imaging Equipment and Viewing Workstations

Places are as in Petri nets. They contain tokens, which represent the state of the system being modeled. The number of tokens in a place is referred to as the *marking* of the place. The assignment of markings to each of the places in a network is referred to as the marking of the network.

Input gates (e.g., *prop*, in Figure 2) and *output gates* (e.g., *remove*, and *addid*) permit greater flexibility than with normal Petri nets in defining enabling and completion rules. Input gates have *enabling predicates* and *functions*, while output gates have only *functions*. The enabling predicate can be either true or false and controls the enabling of an attached activity. The function describes an action (change in marking) that will occur upon completion of the activity. Activities are *enabled* if there is at least one token in each of the places directly connected to the activity and if the predicate of each associated input gate is true (i.e., *holds*).

Directed arcs are used in the graphical SAN description to denote “default” gates. In particular, arcs directed toward an activity are input gates which have a predicate of “at least one token” in the associated place, and a function which decrements the marking of the place by one. Arcs directed away from an activity are a default type of output gate, with function which add one token to the associated place. Both of these arcs function exactly as in normal Petri nets. Lines with two cross marks on them are inhibitor arcs, as in normal Petri nets.

The stochastic nature of the nets is realized by associating an *activity time distribution function* with each timed activity and a *probability distribution* with each set of cases. Generally, both distributions can depend on the global marking of the network. A *reactivation function* [18] is also associated with each timed activity. This function specifies a set of *reactivation markings* for each marking. Informally, given that an activity is activated in a specific marking, the activity is *reactivated* whenever any marking in the set of reactivation markings is reached. This provides a mechanism for restarting activities that have been activated, either with the same or different distribution. This decision is made on a per activity basis (based on the reactivation function), and is not a net-wide execution policy.

The execution of stochastic activity networks is discussed in detail in several places, including [21]. Informally, though, SANs execute in time through completions of activities that result in changes in markings. More specifically, an activity is chosen to *complete* in the current marking based on the relative priority among activities (instantaneous activities have priority over timed activities) and the activity time distributions of *enabled* activities.

A case of an activity chosen to complete is then selected based on the probability distribution for that set of cases. These two choices determine uniquely the next marking of the network, which is then obtained by executing the input gates connected to the input of the activity chosen and the output gates connected to the chosen case. This procedure is repeated by considering the activities enabled in the new marking.

Stochastic activity networks can be solved by both analysis and simulation, depending on system characteristics. Informally, SANs can be solved via analytic methods when all activity time distributions are exponential and activities are reactivated often enough to ensure that their rates depend only on the current state. When this is the case, stochastic processes exist that can be used to obtain analytic solutions for a wide class of variables characterizing both activity and marking related behavior. If this is not the case, simulation can be used to evaluate system behavior.

In order to be effectively applied to realistic systems, model construction and solution techniques require machine implementation. Both the complexity of the construction procedures and the typical sizes of resulting base models make this a necessity. To fill this need, an extensive software package called METASAN¹ [22] has been developed specifically for the construction and solution of SAN-based performability models. This package was used in the evaluation that follows.

IV PACS Models

A Component Models

The particular PACS design to be evaluated is now described in more detail and the workload under which the system is assumed to operate is specified. In particular, we assume a PACS network design of the type described in Section 2, with three network component types: imaging equipment, viewing workstations, and DBAS. As per the concept of a totally digital radiology department [6], we assume that the total PACS will consist of multiple stars. However, this investigation focuses on a single star in a radiology department. This star consists of one star coupler, many pieces of imaging equipment (2-31), three viewing workstations, and a single DBAS.

When a patient is admitted for a procedure session, a number of images may be generated. These images are stored in the local database of the imaging equipment. When

¹METASAN is a registered Trademark of the Industrial Technology Institute.

the procedures are finished, the IE sends a request for the INET in order to transfer the images to the DBAS. The request is queued at the network manager in the DBAS. When the request is at the front of the queue, the network manager sends a connection granted packet on CNET to the IE, and prepares the DBAS to receive and store the image. Upon reception of this message, the IE starts transmitting the image over INET. Upon completion of the image transfer to DBAS, the network manager closes the connection and processes another request.

It is assumed that each WS has enough memory to store the whole examination, and that the request rate of a viewing workstation is the same as for a piece of imaging equipment. The difference is the amount of data transferred. Here, a connection request means the viewing workstation is requesting the DBAS to send a number of images to it.

For the current study, we ignore the migration algorithms and consider the DBAS as a fast storage device. Moreover, we consider that there is enough buffer space at each DBAS to enable it to handle multiple retrieval and storage requests without any delay. Image retrieval time is very dependent on the type of hardware employed but typically is small, on the order of a few milliseconds. This time has been neglected in this study, but could be added simply. It is also assumed that requested images are always available in the DBAS buffer and do not need to be fetched from the archive.

The star coupler is a passive device. It is modeled as a joint point between branches of the star. The network manager controls the transfer flow over INET, using transmissions on CNET. CNET employs CSMA/CD as a method to access the channel at a transmission speed of 10 Mbps. INET is a circuit-switched channel running at 140 Mbps, used for transferring images between PACS nodes. The network manager, which resides at the DBAS, keeps track of transfer requests in a FIFO queue. The INET is allocated to one node at a time, which may transfer a number of images. Image sizes differ according to the sizes of films used for a procedure.

B Assumed Workload

Imaging equipment and viewing workstations are assumed to generate requests as a Poisson process, with a variable mean time between requests. There are two types of requests, the first corresponding to a request for transfer of images (connect request), and the second corresponding to a request for patient information. In this study, we assume that for imaging equipment, 94% of the requests are connect requests, and 4% are to

Table 1: Message Types Transmitted by Imaging Equipment over CNET

Message Type	Length (in bytes)	Fraction of Messages of this Type
connect request	1500	.94
patient information	750	.06

Table 2: Message Types Transmitted by viewing workstation over CNET

Message Type	Length (in bytes)	Fraction of Messages of this Type
connect request	1500	.50
patient information	750	.50

obtain patient information, as is shown in Table 1. Each connection request results in 2 or 4 image transfers, as suggested in [7]. We assume that 60% of the connection requests result in 2 image transfers and 40% of the requests result in the transfer of 4 images. The requests generated by a viewing workstation are assumed to be of the same type as those generated by imaging equipment, and are weighted evenly between connect requests and patient information requests (as is shown in Table 2).

Similarly, as with imaging equipment, each connection request can result in the transfer of multiple images (i.e. a folder of images). For this study, we assume that a radiologist requests 4, 6, 8 or 12 images per folder, according to the probabilities given in Table 3. The network manager, which resides in the DBAS, responds to all requests with a 1500 byte response packet over CNET. Multiple sizes of images are also considered. To determine

Table 3: Number of Images Transmitted per Folder Request

Number of Images	Fraction of Total Requests
4	.40
6	.40
8	.15
12	.05

Table 4: Size and Frequency of Request of Different Film Types

Film Type	Size	Length (in bits)	Fraction of Films of this Type
1	14" × 17"	2048 × 2048 × 12	.60
2	11" × 14"	1609 × 1688 × 12	.15
3	10" × 12"	1462 × 1448 × 12	.22
4	8" × 10"	1170 × 1204 × 12	.03

the sizes and relative frequencies of film types, we examined the raw data of [7]. Four film sizes were observed with the relative frequencies given in Table 4. In all cases we assume no image compression is done, and full images are sent back and forth. All the transmissions in the evaluation are done at the NIU to NIU level. The delays associated with passing a request to, or an image from, an upper layer are not considered.

We realize that when the PACS is introduced it will change the operating practices of radiology departments. An educated guess is that the data-flow volume and utilization will be increased by the PACS [6]. Furthermore, in the graph in Figure 3, which was derived from raw data of [7], we see the expected traffic over the network by year 1989 (the gray line). The thin line shows the traffic in year 1985; notice the increase. The traffic change between 1985 and 1989 is attributed to the increasing number of patients treated and the expansion of the radiology department by installing more equipment. These values were computed assuming a growth of 10% per year, as was estimated by University of Arizona Radiology Department. Traffic patterns change over the days of the week and over hours of the day. Note that, as shown in Figure 3, that a considerable volume of traffic takes place between 1 - 4 p.m., with a peak around 2:30 p.m. This peak reaches 300 Mbytes per 15 minute time interval. So, there is an average of 20 M byte/min, or equivalently, 2-3 image transfers per minute, on the average. Moreover, a request may specify more than two image transfers per imaging equipment or viewing workstation, as previously indicated.

C SAN Representation

The network structure and workload described in the previous section allow us to construct a SAN model of the PACS network and components. The model is broken down into three submodels: an imaging equipment/viewing workstation model, a CNET database in-




Figure 3: Traffic Over the Network During Hours of the Day

terface model, and an INET channel model. This organization keeps the model modular, which aids in maintenance, and when coupled with recently developed construction and solution methods [19], makes model solution more efficient. Each of these SAN models will be described in detail.

As per the assumptions of the last section, imaging equipment and viewing workstations are identical with respect to medium access method, but differ in their relative frequency of message types. The SAN models for both equipment types can thus be similar, differing only in the probabilities assigned to each message type. Figure 2 shows the SAN representation of a single viewing workstation or imaging equipment. Activity time distributions, gate functions, and predicates for this SAN can be found in Appendix A. In this model, the timed activity *arr* represents the arrival of a request from an upper layer to be transmitted over the command channel. These requests are modeled as tokens waiting in place *queue*. The activity *idledly* represents the channel sense procedure and the back off algorithm. Here, the algorithm chooses whether a transmission will occur immediately if the channel is idle, or if the NIU will back off. Two types of requests can occur: a request for connection (and subsequent transfer of images) and a request for patient information. Each type of request has a different probability (as specified in previous section) and, in turn, a different transmission time (due to the difference in packet length). Packet transmission times are represented by the activities *send1* and *send2*; the probability of sending each type depends on the probabilities of cases associated with the instantaneous activity *sw*. The command channel, as seen by this NIU, is represented by the place *channel*. When a node is transmitting, it places its identification in this place. No tokens in this place means the channel is idle.

The DBAS model is very similar to the imaging equipment/viewing workstation model. The difference, as seen in Figure 4, lies in the generation of responses to different nodes in the network. It is assumed that all nodes talk to the DBAS via packets over CNET and the DBAS answers back over CNET. Packets coming from nodes that had avoided collisions and reached the DBAS will be available in *correctsendsc*. This activates activity *arr*, which represents the “thinking” process at the DBAS before sending back the proper response.

The place *channelsc* represents the command channel as seen by the star coupler, and hence is shared among all nodes connected to this star. The activity *wave* represents the signal propagation from one end of a branch through the star coupler to the ends of other branches.

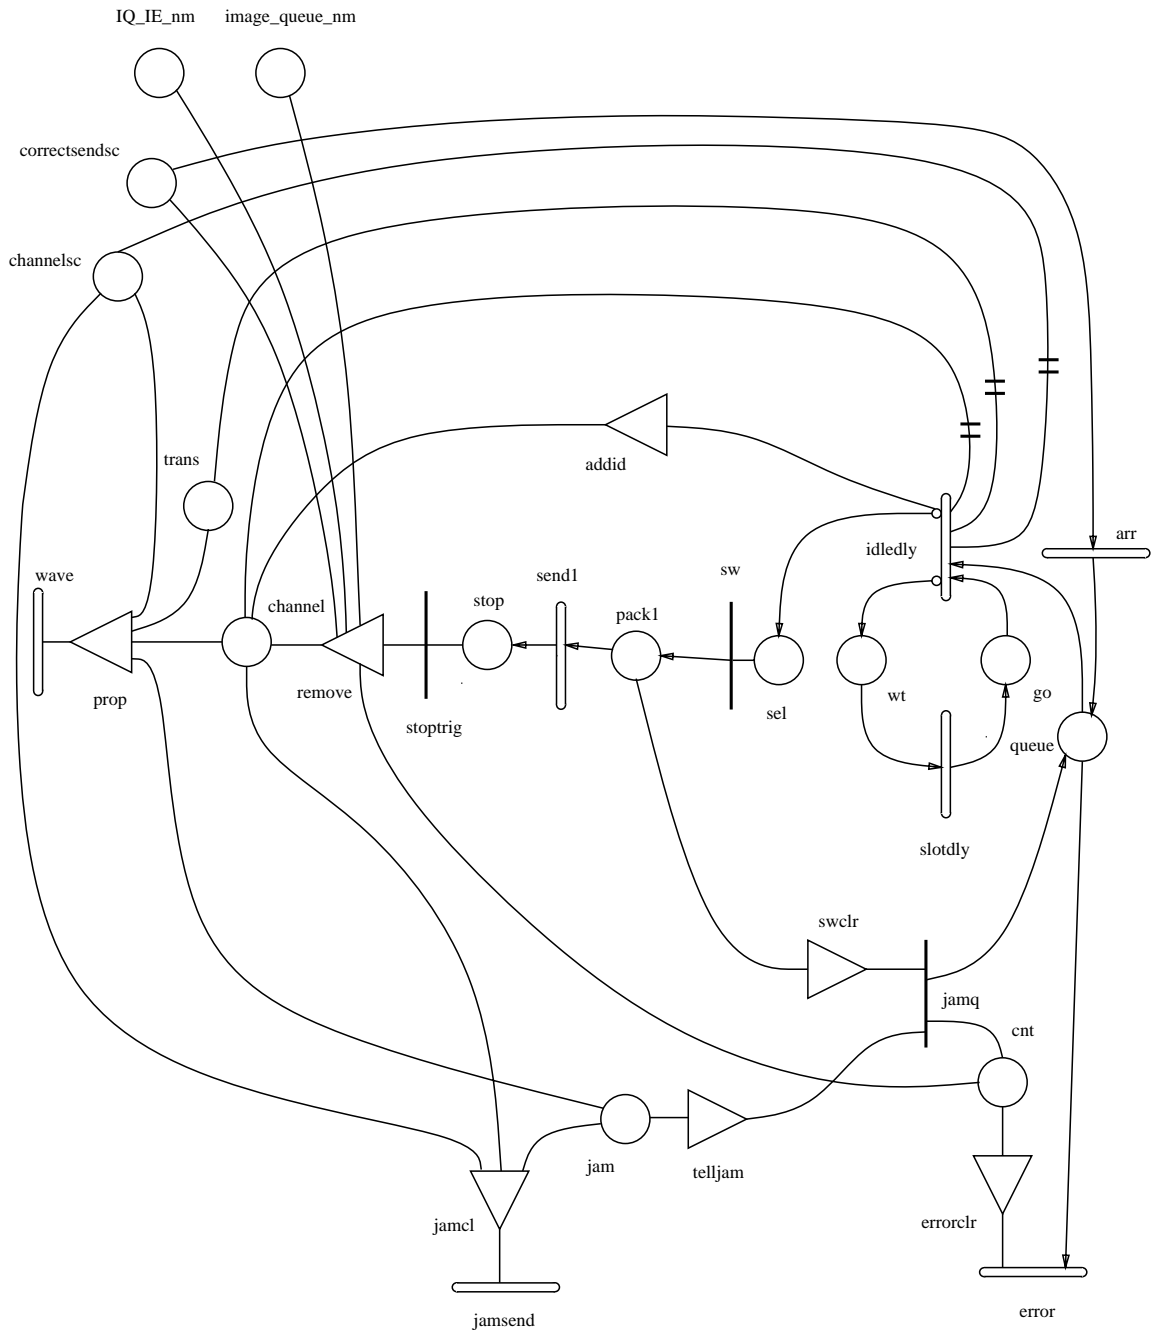


Figure 4: SAN Model of the DBAS

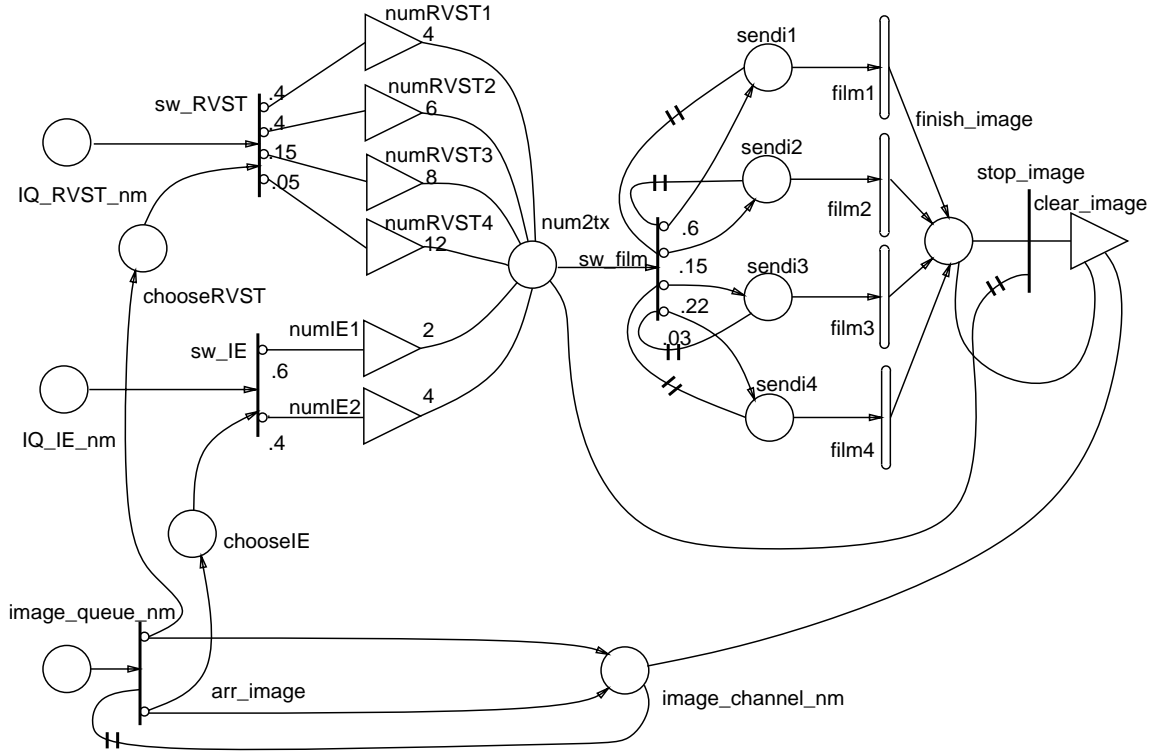


Figure 5: SAN Model of the Image Channel

Image transfer requests are queued at place *image_queue_nm* in the network manager, as shown in Figure 5. Each request involves the transfer of a number of images that depends on whether the request was from an imaging equipment or a viewing workstation. The case probabilities associated with the instantaneous activities *sw_RVST* and *sw_IE* determine the number of images transferred per request. Similarly, the case probabilities associated with the instantaneous activity *sw_film* will determine the size of film transferred. The activities *film1*, *film2*, *film3* and *film4* represent the transmission time of each film type on the image channel. The place *image_channel_nm* represents the image channel. If it has tokens in it, the channel is busy transferring images, otherwise it is idle. The output gate *clear_image* is responsible for resetting the image channel after a request is completed.

D Model Execution

To illustrate the functioning of the model, we now trace the propagation of a request through the system. A request generation is signaled by the completion of activity *arr*,

and results in the addition of a token to place *queue*, as seen in Figure 2. This action, together with the presence of an idle command channel, enables activity *idledly* and initiates processing of the request. The case probabilities associated with *idledly* then determine whether transmission will be attempted or deferred. These probabilities are determined by the number of collisions the current frame has been involved in. If a transmission is attempted, the type of packet transmitted is selected by the cases at the instantaneous activity *sw*. Simultaneously, the output gate *addid* places the NIU identification in the place *channel*. The transmitted signal propagates along the fiber to the star coupler, represented by the completion of the activity *wave*. Once the signal reaches the center, the output gate *prop* places the identification of the NIU in *channelsc*, which represents the point where all branches of fiber intersect. The signal then propagates along the other branches of fiber to all other NIUs. When the activity *send* finishes, i.e. the packet was delivered correctly, the output gate *remove* terminates the transmission and places the frame in the destination queue.

A collision will occur on the command channel if a propagating signal finds another signal at either the star coupler or at another node. In this case, the activity *jamsend* is activated, resulting in the transmission of a jam signal over the command channel. The current transmission is then aborted, the retry counter *cnt* is incremented, and the message is queued again for a later transmission attempt. The number of tokens in place *cnt* controls the probabilities of cases of activity *idledly*, simulating the protocol backoff algorithm. If the number of tokens in *cnt* exceeds 16, the message is rejected and an *error* activity reports the incident to upper layers.

The network manager maintains two queues, corresponding to requests for transfer of images and requests for transfer of patient information. A request may require a transfer of several images over the image channel and a response over the command channel, or it may require only a response over the command channel. An image transfer request is queued as a token in *image_queue_nm*. Once the *image_queue_nm* is empty, i.e. the channel is idle, a request for an image transfer is serviced and a token is placed in the *image_channel_nm* declaring the establishment of a circuit and signaling that the channel is busy. The places *chooseIE* and *chooseRVST* are used with the cases associated with *arr_image* to ensure fair handling to requests coming from imaging equipment and viewing workstations. Tokens in the places *chooseRVST* and *IQ_RVST_nm* trigger the instantaneous activity *sw_RVST* and put the number of images to be transmitted in *num2tx*. For each image, selection of a

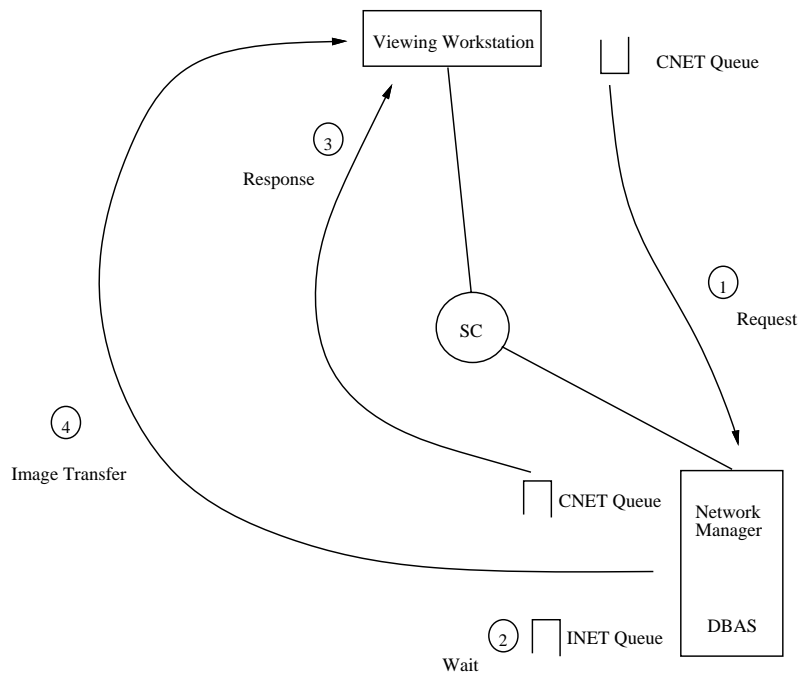


Figure 6: The Process of Transferring an Image

case of sw_film corresponds to a selection of a film type to transmit (these case probabilities were chosen according to statistics of types of film traffic). For example, if activity $film1$ is activated, it means that an image of the first type is being transmitted over the image channel and that the channel will remain busy until this activity terminates. When all the images corresponding to one request are transferred, i.e. no tokens remain in place $num2tx$, the output gate $clear_image$ clears the place $image_channel_nm$ in order to allow the processing of a new request from $image_queue_nm$.

E Performance Variables

In the context of this study, we estimate total response time, CNET queue length, INET queue length, and percentage utilization of INET and CNET. The derivation of each of these variables is discussed in more detail below. In particular, we define total response time as the time between issuing a request for a patient's folder and the arrival of the first image. In order to estimate this time, one must understand the constituent PACS actions which must be executed to achieve an image transfer. The process of transferring a group of images from the database to a viewing workstation is shown in Figure 6. The total response

time is made up of several delays. First, it involves the waiting time in a CNET queue at the requesting NIU. Second, there is the packet transmission time over CNET. In addition, there is the waiting time at the INET queue until the image network is available. When the image channel does become free, an indication of this fact must be sent over CNET, adding to the total transfer time. Finally, when the response packet reaches the requesting NIU, it starts transmitting the image, adding the image transmission time to the total response time.

To determine the average delay associated with each action, one needs to estimate the average length of the queue of requests for the action and the time between completions of service requests, and then apply Little's result [23]. In a SAN model, a queue is represented by a place and frames in that queue are represented by tokens in the place. The average queue length can then be determined by monitoring the average marking (number of tokens) in the place over time. The arrival of frames to a queue is modeled by the completion of an activity. In the case of INET, the instantaneous activity *arr_image* completes each time the transmission of a group of images begins. Thus, since the rate of arrival in a work conserving system is equal to the rate of departure in steady state, the mean time between completions of transmission of a group of images is the mean time between completions of *arr_image*, which can be easily estimated. Furthermore, for the case of CNET, the mean time between arrivals of requests (arrivals of tokens to place *queue*) is the reciprocal of the rate of the exponential activity *arr*.

While total response time may be the most important measure of performance to a user, the utilization of the PACS is an important measure to a system designer. Utilization is measured as the percentage of time the channel is busy, i.e. the fraction of time place *image_channel_nm* contains a token. The utilization of INET can be used to determine the workload at which the system becomes unstable. Furthermore, the utilization percentage of CNET shows the feasibility of using CSMA/CD as a channel arbitration scheme. CNET utilization is the fraction of total time that a valid packet is present on the channel, i.e. the fraction of time a valid node identification is present in place *channelsc*. An excessive collision rate will affect the system's overall performance by inducing delays over CNET. The fraction of time a collision is present on the channel can be found by estimating the fraction of time a jam signal is present on CNET. All CNET performance estimates are made at the place *channelsc*, since the state of the channel at the star coupler reflects the prevailing state over the star branches.

V Discussion of Results

The results of solution of the stochastic activity network PACS model are reported in this section. Specifically, we estimate the expected total response time to deliver a $14'' \times 17''$ image, the expected delay over CNET, the expected delay over INET, channel utilization of CNET, collision rate over CNET, and channel utilization of INET.

METASAN was used as the evaluation tool, and simulation was chosen as the solution method [22]. Multiple simulation runs were conducted, where the mean time between requests (for patient information or image transfers) was varied from 20 seconds to 70 seconds. Furthermore, the number of active nodes in the one star system was varied. In each scenario, there was a single DBAS, three viewing workstations, and some number of imaging equipment attached to the star coupler. This allowed us to see the effect of increasing the number of nodes on a single star network. The number of imaging equipment was varied from one to thirty-one. All results are reported at a confidence level of 90% with an interval with a relative half width of 0.1. The system was started cold, i.e. empty, and the initial portion of each run was discarded to avoid using the transient period in calculating the estimator. The batch size used was made large enough to assume independence between batch means, and large enough number of batches was used to avoid problems due to non-normality of the batch samples.

As per our previous discussion, different delays were estimated in order to calculate the total response time to deliver a $14'' \times 17''$ image. It was observed that a message waits an average of $1200 \mu s$ at a NIU CNET queue before being transmitted over the command channel. This average wait time remained fairly constant when the mean time between messages (per station) was varied from 20 seconds to 70 seconds, and when the number of nodes was varied from 5 to 35. However, the queueing delay at the INET queue jumps sharply when the mean time between messages decreases to less than 30 seconds, as depicted in Figure 7. The total time to transfer an $11'' \times 17''$ x-ray image (largest image size) is shown in Figure 8. The results are generally good when loads are moderate. For example, a 25 node system is capable of delivering the first image within 1.5 seconds while the interarrival time between requests is more than 25 seconds.

CNET had an extremely low utilization, with a utilization percentage of less than 0.03% in all cases, in spite of a possibly large packet size (1500 bytes). This is due to the relatively large time between CNET requests. Through all the different simulation runs conducted, no




Figure 7: Expected Delay in Image Channel Queue




Figure 8: Expected Total Delay for Delivery of a 14" × 17" Image

collisions were observed. This is due to the limited length of fiber the signal has to travel. The maximum distance a signal has to travel is from one end of a star branch, through the passive coupler to the end of another star branch. The time required to do this was assumed to be $10 \mu s$, assuming, very conservatively, that the signals in the cable travel at two-thirds the speed of light, and each branch is one kilometer in length. Even with 20 seconds as the mean time between requests, the probability of collision was thus extremely small. This suggests that CNET can be used for transferring more data types or greater amounts of data without affecting overall system performance.

Image channel utilization, as a function of number of nodes and mean time between requests, is shown in Figure 9. The image channel utilization increased as expected when the number of nodes and frequency of requests increased. For example, for a thirty node system, image channel utilization was less than 70% when the mean time between requests was more than 30 seconds, as shown in Figure 9. If this mean time between requests decreases to 25 seconds the utilization jumps to around 90% for a thirty node system. Furthermore, a decrease in this interarrival time to 20 seconds causes the thirty node system to be unstable.

VI Conclusions

As stated in the introduction, the objective of this study was two-fold, to illustrate the usefulness of stochastic activity networks in evaluating picture archiving and communication systems and to investigate the performance of a particular PACS design. The results presented in the discussion section demonstrate our success in both respects. In particular, with regard to the first objective, we showed that one could represent both a PACS and its workload demands. In all cases, system and workload characteristics were accurately represented with respect to the stated objectives. Moreover, the use of stochastic activity networks permitted an efficient solution for desired system characteristics via simulation, as illustrated by the results we obtained.

The results themselves speak to the quality of the particular PACS design evaluated. Although the system design, as would any, has limitations with regard to workload and number of operating nodes, it performed well over a wide range of situations. Collisions were shown not to be a problem, since CNET utilization is very light. The total delay time remained good until INET utilization became extremely large. In summary, the results show the utility of the two channel approach to designing PACS networks, and, additionally, that




Figure 9: Utilization of the Image Channel

this passive star based design performs very well.

Future work should investigate the possibility of carrying more traffic over CNET, such as voice data. Moreover, a distributed database system spread among different stars should be evaluated. In the distributed database system, databases have finite buffer space with a significant storage/retrieval overhead. Furthermore, different migration algorithms can be evaluated and their impact on the system's overall performance can be studied. Finally, different methods to interconnect multiple stars should be studied carefully. The interconnection of stars can be achieved via a backbone network (FDDI, ISDN) or via gateways.

REFERENCES

- [1] I. Obayashi, Y. Takahashi, Y. Tani, M. Suzuki, "Network analysis of picture archiving and communication system," in *ISMII Proc.*, pp 109-119, 1984.
- [2] U. M. Maydell, P. S. Gill and W. A. Davis, "Design criteria for an interactive image computer system," in *Proc. ELECTRONICOM '85*, pp. 282-285, Toronto, Oct. 1985.
- [3] R. Martinez, C. Archiwamety and M. Nemat, "Simulation and design of ACR-NEMA standard for fiber optics network image transfer," in *Proc. SPIE 1987, Medical Imaging, Vol. 767*, February 1987.
- [4] A. D. A. Massar, J. P. J. de Valk, G. L. Reijns, A. R. Bakker, "Simulation of an image network in a medical image information system," in *Proc. SPIE 1985, Medical Imaging Processing, Vol. 593*.
- [5] W. A. Davis, U. M. Maydell and P. S. Gill, "Simulation of a PACS in a hospital environment," in *Proc. SPIE 1986, PACS IV, Vol. 626*, 1986.
- [6] W. J. Dallas et al., "A prototype totally digital radiology department: conception and initiation," in *Proc. SPIE 1987, PACS V, Vol. 767*, February 1987.
- [7] T. Ozeki, M. Suzuki, H. Umino, R. Martinez, C. Archiwamety, M. Nemat, "Investigation and analysis of the requirements for a picture archiving and communication system (PACS) at the University of Arizona Health Science Center," in *Proc. SPIE 1987, Medical Imaging, Vol. 767*, February 1987.
- [8] R. Martinez and M. Nematbakhsh, "Design and performance evaluation of a high speed fiber optic integrated computer network for picture archiving and communications system," in *Proc. SPIE 1989, Medical Imaging*, Newport Beach, CA, February 1989.
- [9] E. Nishihara, K. Tawara, K. Komatsu, T. Okikawa, K. Kosos, "High speed image transfer for PACS," in *Proc. SPIE 1987, PACS V, Vol. 767*, February 1987.
- [10] K. Saito, "Present and future of diagnostic imaging systems," in *Medical Review No.20*, pp. 47-57, 1987.

- [11] B. Roberts, M. Kakegawa, M. Nishikawa, D. Oikawa, "Toshiba TDF-500 high resolution viewing and analysis system," in *Proc. SPIE 1988, Medical Imaging II, Vol. 914*, February 1988.
- [12] R. Vercillo, H. D. Fisher III, R.D. Lamoreaux, "Digital imaging review console," in *Proc. SPIE 1987, PACS V, Vol. 767*, February 1987.
- [13] M. K. Molloy, "Performance analysis using stochastic Petri nets," in *IEEE Transactions on Computers, Vol. C-31, pp.913-917*, Sept. 1982.
- [14] S. Natkin, "Reseaux de Petri stochastiques," These de Docteur-Inge'nieur, CNAM-PARIS, June 1980.
- [15] J. F. Meyer and L. Wei, "Influence of workload on error recovery in random access memories," in *IEEE Transactions on Computers, Vol. 37, No. 4*, April 1988.
- [16] J. F Meyer, K. H. Muraldihar and W.H. Sanders, "Performability of a token bus network under transient fault conditions," in *Proc. 19th Int. Symp. on Fault-tolerant computing*, Chicago, June 1989.
- [17] A. Movaghar and J. F. Meyer, "Performability modeling with stochastic activity networks," in *Proc. 1984 Real-Time Systems Symp.*, Austin, TX, Dec. 1984.
- [18] J. F. Meyer, A. Movaghar, and W. H. Sanders, "Stochastic activity networks: structure, behavior, and application," in *Proc. International Workshop on Timed Petri Nets*, Torino, Italy, July 1985, pp. 106–115.
- [19] J. A. Couvillion, R. Freire, R. Johnson, W. D. Oball II, M. A. Qureshi, M. Rai, W. H. Sanders and J. Tvedt , "Performability Modeling with UltraSAN," *IEEE Software*, pp. 69–80. Sep. 1991.
- [20] J. L. Peterson, *Petri Net Theory and the Modeling of Systems*, Englewood Cliffs: Prentice-Hall, 1981.
- [21] W. H. Sanders, "Construction and solution of performability models based on stochastic activity networks," Computing Research Laboratory Technical Report CRL-TR-9-88, The University of Michigan, Ann Arbor, MI, August 1988.
- [22] W. H. Sanders and J. F. Meyer, "METASAN: a performability evaluation tool based on stochastic activity networks," in *Proc. ACM-IEEE Comp. Soc. 1986 Fall Joint Comp. Conf.*, Dallas, TX, Nov. 1986.
- [23] M. Molloy, "Fundamentals of performance modeling," Macmillan Publishing Company, 1989.

APPENDIX A

Table 5: Activity Parameters for IE/WS SAN

Activity	Rate	Probability	
		case 1	case 2
<i>arr</i>	$\exp(\lambda)$	-	-
<i>error</i>	determ(ERR)	-	-
<i>idledly</i>	determ(IDLE)	binary backoff	binary backoff
<i>jamq</i>	inst	1	-
<i>jamsend</i>	determ(JAM)	-	-
<i>send1</i>	determ(1500)	1	-
<i>send2</i>	determ(750)	1	-
<i>slotdly</i>	determ(SLOT)	1	-
<i>stoptrig</i>	inst	1	-
<i>sw</i>	inst	PACK1	PACK2
<i>wave</i>	determ(DIST)	-	-

Table 6: Input Gate Parameters for IE/WS SAN

Gate	Enabling Predicate	Function
<i>erreclr</i>	$\text{MARK}(cnt) \geq \text{CNT}$	$\text{MARK}(cnt) = 0;$
<i>jamcl</i>	$(\text{MARK}(channel) > 0 \text{ and } \text{MARK}(jam) > 0)$	$\text{MARK}(channel) = 0;$ $\text{MARK}(jam) = 0;$ $\text{MARK}(channelsc) = 0;$
<i>prop</i>	$(\text{MARK}(channelsc) == 1 \text{ and } \text{MARK}(jam) == 0)$ or $(\text{MARK}(channel) == 1 \text{ and } \text{MARK}(jam) == 0)$ or $(\text{MARK}(channel) > 1 \text{ and } \text{MARK}(trans) == 0)$ or $(\text{MARK}(channel) == 0 \text{ and } \text{MARK}(trans) == 1)$	controls signal propagation on CNET
<i>swclr</i>	$(\text{MARK}(pack1) > 0 \text{ or } \text{MARK}(pack2) > 0)$	$\text{MARK}(pack1) = 0$ $\text{MARK}(pack2) = 0$
<i>telljam</i>	$\text{MARK}(jam) > 0$	identity

Table 7: Output Gate Parameters for IE/WS SAN

Gate	Function
<i>addid</i>	MARK(<i>channel</i>) = id;
<i>remove</i>	MARK(<i>channel</i>) = 0; MARK(<i>cnt</i>) = 0; MARK(<i>correctsendsc</i>) ++; if (MARK(<i>image_request</i>) > 0) { MARK(<i>image_request</i>) = 0 MARK(<i>image_queue_nm</i>) ++; MARK(<i>IQ_RVSR_nm</i>) ++; }

Table 8: Activity Parameters for Image Channel SAN

Activity	Rate	Probability			
		case 1	case 2	case 3	case 4
<i>arr_image</i>	inst	proportion of number of IE requests	proportion of number of WS requests	-	-
<i>film1</i>	determ(PFILM1)	1	-	-	-
<i>film2</i>	determ(PFILM2)	1	-	-	-
<i>film3</i>	determ(PFILM3)	1	-	-	-
<i>film4</i>	determ(PFILM4)	1	-	-	-
<i>stop_image</i>	inst	1	-	-	-
<i>sw_IE</i>	inst	0.6	0.4	-	-
<i>sw_RVST</i>	inst	0.4	0.4	0.15	0.05
<i>sw_film</i>	inst	0.6	0.15	0.22	0.03

Table 9: Output Gate Parameters for Image Channel SAN

Gate	Function
<i>clear_image</i>	MARK(<i>image_channel_nm</i>) = 0; MARK(<i>finish_image</i>) = 0;
<i>numIE1</i>	MARK(<i>num2tx</i>) = 2;
<i>numIE2</i>	MARK(<i>num2tx</i>) = 4;
<i>numRVST1</i>	MARK(<i>num2tx</i>) = 4;
<i>numRVST2</i>	MARK(<i>num2tx</i>) = 6;
<i>numRVST3</i>	MARK(<i>num2tx</i>) = 8;
<i>numRVST4</i>	MARK(<i>num2tx</i>) = 12;