Some Design Considerations for Picture Archiving and Communication Systems

J. R. Cox, G. J. Blaine, R. L. Hill, R. G. Jost, and C. D. Shum

Design considerations for picture archiving and communication systems are reviewed with special emphasis on those issues that differ from conventional network architectures. Design equations for three layers of a picture network are developed and discussed in the context of preliminary estimates of the flow of digital images between a multiplicity of picture sources, picture archives and picture viewing stations. Discussions of differences from conventional networks focuses on the local nature of the net, the availability of a wide-band transmission media with low error rates, the relative costliness of network equipment capable of taking advantage of the wide-band transmission media... **Read complete abstract on page 2.**

Recommended Citation

Some Design Considerations for Picture Archiving and Communication Systems

J. R. Cox, G. J. Blaine, R. L. Hill, R. G. Jost, and C. D. Shum

Complete Abstract:

Design considerations for picture archiving and communication systems are reviewed with special emphasis on those issues that differ from conventional network architectures. Design equations for three layers of a picture network are developed and discussed in the context of preliminary estimates of the flow of digital images between a multiplicity of picture sources, picture archives and picture viewing stations. Discussions of differences from conventional networks focuses on the local nature of the net, the availability of a wide-band transmission media with low error rates, the relative costliness of network equipment capable of taking advantage of the wide-band transmission media and the preponderance of large blocks of image data in the network traffic. The locality of network traffic flows associated with the generation, storage and display of images is a characteristic of the radiology application that can probably be used to advantage.

SOME DESIGN CONSIDERATIONS FOR PICTURE

ARCHIVING AND COMMUNICATION SYSTEMS

.

J.R. Cox, Jr., G.J. Blaine,
R.L. Hill, R.G. Jost and C.D. Shum

WUCS-83-1

Department of Computer Science

Washington University

St. Louis, Missouri 63130

Some Design Considerations for Picture Archiving and Communication Systems

J.R. Cox, Jr., G.J. Blaine, R.L. Hill, R.G. Jost, C.D. Shum
Washington University

Abstract

Design considerations for picture archiving and communication systems are reviewed with special emphasis on those issues that differ from conventional network architectures. Design equations for three layers of a picture network are developed and discussed in the context of preliminary estimates of the flow of digital images between a multiplicity of picture sources, picture archives and picture viewing stations. Discussion of differences from conventional networks focuses on the local nature of the net, the availability of a wide-band transmission media with low error rates, the relative costliness of network equipment capable of taking advantage of the wide-band transmission media and the preponderance of large blocks of image data in the network traffic. The locality of network traffic flows associated with the generation, storage and display of images is a characteristic of the radiology application that can probably be used to advantage.

Introduction

The layered approach to the design of computer networks has been developed in the contexts of distributed data processing and office automation.[1,2] We have suggested[3] an analogous approach for picture archiving and communication systems (PACS). By studying the first three layers (physical, picture link, picture network) several design equations appropriate to PACS have been obtained. These equations combined with estimates of the traffic flow between network nodes form the basis of a methodology for the prediction of the performance of various PACS designs. The recent availability of

wide-band transmission media, the requirement for rapid network response in a department of radiology and the preponderance of large blocks of image data, suggest architectures that differ markedly from those popular for distributed data processing and office automation. An architecture consisting of multiple subnets interconnected by internet bridges is suggested that takes advantage of the locality anticipated in network traffic flow while limiting the cost of network equipment.

The advent of digital radiographic images has revealed the prospect of a picture network. Early reports of the design of such systems have appeared,[4-8] optical disks capable of archiving large amounts of data are being field tested,[9-11] new magnetic recording techniques promising major increases in capacity have been reported,[12] and methods for the compression of images have been suggested.[13,14]

Advantages of PACS have been widely recognized, but the performance of present network equipment is inadequate for rapid response, high volume systems. Careful design will be required if the advantages anticipated for departments of radiology are to be achieved at reasonable cost. This paper summarizes several steps toward prediction of network performance useful in the design of a PACS.

The basic form of a PACS (Figure 1) includes three classes of nodes interconnected by a wide-band network. The first class of nodes includes all image sources: computed tomography,[15] nuclear medicine,[16] ultrasound,[17] digital subtraction angiography,[18] nuclear magnetic resonance,[19] positron emission tomography,[20,21] and eventually chest and other forms of radiology. The picture archive may be centralized, but is likely to be, at least in part, decentralized at multiple nodes located throughout the network. In either case it will be heavily dependent upon new technology in the

form of optical disks or new magnetic recording systems each providing
substantial increases in information packing density. The third class
of network nodes, the picture viewing stations will incorporate image pro-
cessing that is appropriate to the radiologist's needs. The functional
characteristics of these stations are just being defined, but rapid response
time seems to be a high priority requirement whether the image to be viewed
is in a distant archive or has already been retrieved but needs to be pro-
cessed before viewing.

PACS engineering design is in its infancy, but one conclusion is
clear. Design issues previously studied in the context of computer networks
must be reviewed. Specification for PACS are unique and new models for
the networks must be developed.

Analog Video Versus Digital Transmission of Pictures

The image sources in radiology are increasingly based on digital
technology. CT has led the way and other modalities requiring image processing
have followed. Even those imaging modalities that have traditionally used
analog techniques (film and fluoroscopy) are expected to move to digital
techniques within the next decade. Furthermore, digital techniques for
the transmission and storage of information can assure the faithful reproduction
of images stored in the archive and avoid the degradation of images character-
istic of analog methods. The choice of digital transmission between source
and archive is clear.

For the next decade the image viewing station is likely to be based
upon analog video signals applied to one or more cathode ray tubes (CRTs)
within the viewing station. Thus conversion of images stored in digital
form in the archive to analog video signals applied to the CRTs must take place
at some point between the archive and the CRTs. Uniformity of transmission

and faithfulness of reproduction argue strongly for an all digital network. Economy may well be on the side of a mixed network, however, with both digital and analog video links present.

Consider an archive node connected to a viewing station where n simultaneously displayed, but different images are presented (see Figure 2). Digital storage to refresh each of the displays can be placed at either the archive or the display end of the link with analog video transmission used in the first case and digital transmission in the second. In either case n refresh buffers are required. The cost of the transmission medium required in the two cases is a small part of the system cost and not of primary importance to this discussion. In the digital case high-speed digital modems are used for modulation and demodulation of digital signals so that they may use a limited portion of a wide-band medium. At present, the cost of two high-speed modems is substantially higher than the corresponding transmitter/receiver pair used in the analog case. Analog video transmission requires n transmitter/receiver pairs to support an n-CRT display station while digital transmission needs but a single pair of modems for the same task. Thus even though a cost differential may exist between the equivalent analog video and digital equipment, there is a value of n for which digital transmission is less expensive. Perhaps that critical value is now ten, but trends in microelectronics indicate it may eventually approach one. Analog video transmission will probably remain less expensive than high-speed digital transmission for a single CRT display station, however, for some years to come.

There are several caveats that must be stated regarding these conclusions. Some digital transmission is essential in a PACS, but lower speed links may suffice for source to archive transmissions. If analog video links are used, they have their place primarily in support of viewing stations. Thus a PACS

must be all digital or support mixed modes of transmission. If the transmission medium chosen does not allow mixed modes the utility of any analog video transmission is questionable since two distinct networks would be required.

The centralized sharing of refresh buffers in the analog video case may lead to economies not available with digital transmission. On the other hand, the volume production of digital viewing stations may eventually lead to economies not available with centralized refresh buffers. In summary, time is on the side of the all digital system, but individual situations within a network will have to be examined to determine when economy no longer favors the analog video transmission of images for display.

In what follows we analyze an all digital PACS. One may wish to consider some network nodes as distribution points for analog video signals not discussed explicitly, hereafter. For example, using a 400-MHz coaxial cable transmission plant, a subset of channels may be allocated to analog video transmission, the rest to digital transmission.

Modeling a Picture Archiving and Communication System (PACS)

Our design method adopts the layered approach to computer networks because it reduces design complexity and conforms well to modern practice. The reference model of Open Systems Interconnection (OSI) proposed by the International Standards Organization (ISO) has several layers,[22] but in the following we consider only the first three layers. The first is the physical layer corresponding exactly to ISO terminology. Since this analysis is aimed at picture archiving and communication exclusively, we vary the ISO terminology and call the next two layers the picture link layer and the picture network layer. It is in these three layers that the differences from the conventional approach to networks takes place. The remaining four, higher-level layers deal with pictures at a sufficiently high level

of abstraction that differences are small from conventional networks handling messages and viewed at the same level of abstraction.

Physical layer

The design issues for this layer concern the transmission of raw bits over a coaxial cable, a microwave link or an optical fiber. Picture service will require wide-band transmission hence neither telephone lines nor twisted pair cable are considered. The physical layer provides a channel with a capacity C in bits per second (b/s) and error rate e in erroneous bits per bit transmitted. The channel can transmit D bits of raw data from source to destination in τ seconds,

$$\tau = \frac{D}{C}. \tag{1}$$

The propagation time has been neglected here because signals travel from one point in a local area network to another in times that are several orders of magnitude less than the transmission time τ. Although (1) is quite simple, it forms the foundation of the analysis to follow.

For convenience in what follows we have chosen to call a standard picture frame (spf) a picture with 512 × 512 pixels and 8 bits of intensity information for each pixel. Thus a 1024 × 1024 × 8 picture corresponds to 4 spf and a 2048 × 2048 × 8 picture corresponds to 16 spf. Simple scaling, e.g., 1.25 for 10 bits/pixel adjusts for a specific image characteristic. Other sizes of images may be used by assigning a single scale factor for each different image size. Also note that one spf corresponds to 2.1 Mb (megabit) or 268 kB (kilobyte).

The transmission time τ given by (1) is plotted in Figure 3 against the channel capacity C with D a parameter chosen to correspond to pictures with 1, 4 and 16 spf. Note transmission times less than 1 second require

channel capacity greater than 2 Mb/s for 1 spf. A 16 spf picture must have a channel capacity greater than 32 Mb/s for subsecond transmission times.

Picture link layer

The second layer takes the raw transmission facility provided by the physical layer and transforms it into a channel that can transmit pictures with a vanishingly small probability of error in spite of the finite error rate e of the physical layer. It transmits information from source to destination in a unit of data that we shall call a frame.[1] Noise may cause a frame to arrive at the destination in damaged condition. The picture link layer at the destination must detect this error and reject the frame. The picture link layer at the source must subsequently retransmit the frame and when the frame eventually arrives at the destination with no detectable errors the picture link layer there must acknowledge its acceptance.

The protocol that accomplishes this retransmission of damaged frames requires the attachment of a prefix and a suffix to the data. The purpose of the prefix is to aid in frame synchronization, to number sequentially the frames and to pass along other control information. The purpose of the suffix is to provide error checking. The destination sends to the source an acknowledgement, which may be a separate transmission or may be "piggy-backed" on a transmission bound in the reverse direction. We assume the prefix and suffix together occupy H bits and the acknowledgement A bits.

The addition of control information and the retransmission of damaged frames increases transmission time. To quantify this effect we introduce the delay factor δ,

$$\delta = \frac{\text{actual transmission time}}{\text{ideal transmission time}} - 1 , \qquad\qquad (2)$$

the fractional increase in the mean transmission time introduced by the actual channel as compared to the ideal channel. In the picture link layer the number of retransmissions n is distributed geometrically[23] with mean

$$\overline{n} = \sum_{n=0}^{\infty} n \, (1-p)^n \, p = \frac{1}{p} - 1 \tag{3}$$

where $(1-p)^n$ is the probability of exactly n retransmissions preceding a successful one and p is the probability of a successful transmission

$$p = (1-e)^H (1-e)^D (1-e)^A = (1-e)^{D_L} \tag{4}$$

assuming independence of errors. In (4) the total number of bits transmitted is D_L = H+D+A, H+D in the forward direction and A in the reverse direction.

On the average there are a total of $\overline{n}$ + 1 transmissions including the successful one. Thus the total transmission time for the picture link layer is

$$\tau_L = \frac{D_L}{C} \, (\overline{n} + 1) = \tau \frac{D_L}{pD} \tag{5}$$

The delay factor for this layer is

$$\delta_L = \frac{\tau_L}{\tau} - 1 = \frac{D_L}{pD} - 1 \simeq D_L \, e \tag{6}$$

where the approximation holds for H+A << D and D_L e << 1. Transmission of a complete spf $(2.1 \times 10^6$ bits) with e < 10^{-8} does not degrade the transmission time measurably $(\delta_L = 2.1 \times 10^{-2})$. Such an error rate seems achievable for local networks.[24,25]

In summary, easily achieved constraints (H+A << D and D_L e << 1) insure that the delay factor is small and that transmission time at the picture link layer is approximately given by (1) and Figure 3.

Picture network layer

The third layer of PACS interfaces the host computers at the picture source, archive and viewing station to the network. At the sending node, a set of images from a study is partitioned into a set of packets which contain routing and control information. The flow of picture packets into the network is adjusted to match network traffic. At the destination this layer receives picture packets from the network, monitors the flow of these packets from the network and organizes the received packets into a set of images that constitute a study.

Model applications. Modeling the picture network layer of a PACS is more difficult than the picture link layer, but some insight can be gained by study of the single network node of Figure 4a. Let us assume that the processor executes the protocol of the picture network layer and connects j inbound channels to k outbound channels each with capacity C. The processor may be a node in a store-and-forward network with j inbound and k outbound channels, may be a gateway or bridge between two networks with j=k=1, or may be a channel controller associated with a single host. In the last case the j inbound channels may represent the network traffic destined for the host and a single outbound channel (k=1) represents a DMA transfer to another memory or to disk storage. Alternatively, a single inbound channel (j=1) may represent DMA transfers and the k outbound channels represent the traffic destined for the network.

If certain assumptions and approximations, to be discussed below, are valid, the model of Figure 4a may also be applied to a single shared transmission medium that transports picture packets between all nodes in a PACS (a scheme similar to Ethernet[26]). In this case there is a single outbound channel (k=1) corresponding to the shared transmission medium. This channel is contended for among all of the j processors at the nodes of the PACS.

We assume some scheme for arbitration among several competing requests for the channel that uses only a small fraction of the channel capacity. For example, we may assume a carrier-sense protocol[26] that limits collisions to a relatively short period preceding the transmission of a packet. Similarly, we may assume a collision-free protocol[1] that uses a short, fixed period before packet transmission for the selection of the sender. Another alternative, is a ring network[27] with a worst case, but still short time for selecting the sender. In any of these three cases the time that must be allocated to arbitration or selection among competing requests does not, by itself, degrade response time since it is small compared to the time required to transmit the picture packet. Finally, we must assume a channel acquisition protocol that assures a first-come-first-serve discipline for all j processors or a protocol that approximates such a discipline.

With these assumptions the collection of j processors at the network nodes share a distributed queue that is composed of all picture packets waiting for service from the shared channel.

Response time. In all of the above cases we can consider an abstraction of the problem which includes only the source rates λ_i in studies per second, the channel capacities C in bits per second and the study size m in packets per study. We further assume that origination of studies at the i^{th} source ($1 \leq i \leq j$) can be represented by a Poisson process with rate λ_i. Although Poisson statistics may not be appropriate in specific cases, other assumptions are mathematically difficult and, at best, equally hard to justify. Calculations with Poisson statistics can give good design information, but may fail if detailed predictions are required. The study size m is assumed to be an integer valued random variable.

The transmission time τ_L found for the picture link layer (5) applies to each outbound channel and leads to a service time $m\,\tau_L$ for an M/G/1 queueing model[28,29] with Poisson input, random service times and a single server. Note that only one outbound channel can be in service at a time, but the processor may switch channels between studies.

Figure 4b gives a simple example with 3 inbound channels (j=3) and a 6.3 second sample transmission. The channel number is identified within each 2.1 Mb packet transmitted. The 10 Mb/s capacity of the outbound channel determines the service time, $\tau_L = 210$ ms.

We assume that the control information added by the picture network layer is included in the header and acknowledgement portions of the packet produced by the picture link layer. With these definitions and assumptions the utilization factor ρ for a network node such as that of Figure 4a is given by

$$\rho = \lambda \, \overline{m} \, \tau_L \tag{7}$$

where λ is the total inbound source rate $\lambda = \sum_{i=1}^{j} \lambda_i$; $\overline{m}$ is the mean study size; τ_L is the picture packet transmission time; and $\overline{m}\,\tau_L$ is the mean service time. Note that the node's utilization factor can be calculated as a busy fraction as shown in Figure 4b.

The mean response time τ_N at a node described at the picture network level is the time a study spends in the queue and in being transmitted (serviced). This time is given by the Pollaczek-Khinchine formula[28,29]

$$\tau_N = \overline{m} \, \tau_L \, (1 + \beta \, \frac{\rho}{1-\rho}) \tag{8}$$

where $\beta = \frac{1}{2} E(m^2)/(\overline{m})^2$ is half the ratio of the second moment of the study size to the square of its mean. For example, a geometric distribution[23] for the study size leads to $\beta = 1 - 1/(2\overline{m})$. Thus, for a large mean study size $\beta \simeq 1$, the maximum value possible for the assumed distribution and the same result that would have been obtained had an exponential distribution of service times been chosen. On the other hand, for a constant study size $E(m^2) = (\overline{m})^2$ and $\beta = \frac{1}{2}$. For mixtures of the two cases the coefficient β in (8) is bounded by $\frac{1}{2} \leq \beta \leq 1$.

The delay factor for the picture network layer can now be calculated,

$$\delta_N = \frac{\tau_N}{\overline{m}\tau} - 1$$

$$= \frac{\overline{m}\ \tau_L}{\overline{m}\tau} (1 + \beta \frac{\rho}{1-\rho}) - 1 = \frac{D_L}{pD} (1 + \beta \frac{\rho}{1-\rho}) - 1$$

$$= \delta_L\ (1 + \beta \frac{\rho}{1-\rho}) + \beta \frac{\rho}{1-\rho} \tag{9}$$

If the delay factor for the picture link layer is negligible ($\delta_L \simeq 0$), the network delay factor is proportional to the ratio of the busy fraction to the idle fraction for a network node.

Buffer size. The buffer size required for the processor in Figure 4a is determined by ρ and the distribution of study sizes. General results are not available, but Dor[30] tabulates values for a queueing model equivalent to a constant study size. For example, to obtain a probability of buffer overflow of 10^{-3} for $\rho = 0.2$ requires a buffer with capacity four times the mean study size. This is an optimistic result since a random distribution of study size can only increase the probability of overflow under the same conditions.

For this reason store-and-forward networks seem ill-suited for picture service. Note that smaller packets do not improve the situation. Fast response to requests for the transmission of an entire study cannot be achieved if some packets are delayed because they cannot fit into a buffer enroute. Consequently, circuit-switched, shared-transmission-medium or ring networks seem most appropriate for picture service.

By use of the design equations (1), (5), (6), (8) and (9) we can analyze many features of PACS architectures, but to reach meaningful conclusions estimates of mean study size and of mean traffic flows between network nodes are needed.

Use of the Design Equations

To illustrate the use of these design equations we summarize some data gathered recently at the Mallinckrodt Institute of Radiology (MIR).[3] About 250 thousand studies containing about 7.5 million images were gathered. Average study size was about 10 spf assuming that chest films and other conventional radiographic examinations were digitized to yield 4 spf per image (1024 × 1024 × 8). Future improvements in resolution may raise average study size to as much as 40 spf.

Network traffic flows for use in the design equations can be obtained by first estimating the total peak traffic flow λ. Although no direct studies to determine λ have been carried out at MIR, pilot data[3] indicate that at peak traffic $\lambda = 0.2$ studies per second (generated and retrieved) is a reasonable estimate for MIR. Future growth may well double λ.

Figure 5 shows the relationship between response time τ_N and channel capacity C plotted for two average study sizes ($\overline{m} = 10$ spf and $\overline{m} = 40$ spf) with traffic flow λ as a parameter. Several values between $\lambda = 0.08$ and $\lambda = 0.4$ have been plotted as well as the low traffic limit at $\lambda = 0$. Equations (1), (5) and (8) have been used to obtain Figure 5 with the assumptions: $D = 1$ spf $= 2.1$ Mb, H+A $<<$ D and e $<< 1/D_L$. Finally we chose a geometric distribution of study size and consequently $\beta \simeq 1$.

Design goals

Responsiveness. Retrieval of an entire study, not part of one, is the most concrete and challenging requirement to be imposed on the design of a PACS. Until a radiologist makes a preliminary review of an entire study, he cannot begin a detailed interpretation of a single image. The context of an image must be presented and for this reason we have concentrated on the study as the unit of transmission to be provided to the higher protocol levels by the picture network level.

Figure 5 shows that study response times of a few seconds require wide-band channels even for low traffic ($\lambda \simeq 0$). No amount of ingenuity can overcome the fundamental limitation imposed by a network with inadequate channel capacity.

The requirement for wide-band channels is most severe for retrievals. The picture sources have internal buffering so that smoothing of fluctuations in the rate of study generation can be achieved by slight delays in archiving.

A network capable of handling λ = 0.02 would probably be adequate if only image generation need be considered. Retrieval traffic is greater than source traffic, however, both in its average value and in its peak-to-average factor. Unlike archiving, retrieval of images cannot be delayed so peak retrieval traffic must be serviced without delay. This leads to traffic flows an order of magnitude greater than those associated with image generation and requiring a network capable of handling λ = 0.2. During such peak activity a shared channel PACS with C = 10 Mb/s would suffer a degradation in average response time of a factor of 1.72 from 2.1s to 3.6s, perhaps a significant degradation in performance. Furthermore, some images would be delayed much longer.

Modularity. The continued growth of radiological image production and interpretation throughout the nation is certain. Growth in the population, growth in the fraction of the population served and growth in the number of tests administered to an individual all contribute to this trend. Over the next decade substantial improvements in PACS components will be made possible by advances in microelectronics. Thus changes in PACS equipment are inevitable. These forces toward growth and change argue for a modular approach toward system design. Furthermore, the estimates of network traffic flow presented above suggest the need to distribute the load over several distinct channels.

Such an approach requires a partitioning of the network into subnets. Clearly this partitioning should be done in a way that minimizes the traffic between subnets. Considerable locality of traffic seems likely for most radiology departments. Today interpretation of radiographic studies is done almost exclusively within a single imaging modality. This may change in the future because of PACS itself, but a preferred modality exists in

most cases and collateral studies in other modalities will, even then, probably represent a small fraction of the total traffic associated with the case.

It is important to recognize that network traffic locality does not require geographic locality. At MIR, interpretation is usually done on the floor where the study was performed so geographic and traffic locality are likely to be similar. Subnets that overlap geographically are possible, however, permitting subdivision of network traffic along functional rather than geographic lines.

Subdivision of the picture archive is an important design objective that can facilitate modular growth and can prevent a traffic bottleneck at the interface between the archive and the network. The age of the study, the imaging modality and the patient identifier are three possible keys on which to classify studies into disjoint archives residing on separate subnets.

Reliability. The attainment of a reliable system is an obvious design goal. Correctly implemented modularity can help here through partial operability in the face of failure. The layered approach to system development can make protocol design and verification more systematic.

Errors introduced during transmission are likely to be infrequent if modern local area net technology is used.[24,25] Studies to be archived should be transmitted with vanishingly small probability of error. For retrieved studies a few pixel errors may be tolerated. Error detection and retransmission protocols seem well-suited to the task of error-free transmission and the delay factor expression (6) for the picture link layer indicates a negligible cost for this procedure. In fact, in an all digital PACS there seems to be little reason to take advantage of the relaxed error criterion that may be tolerated on retrievals.

Efficiency. An efficient transmission scheme minimizes the channel bandwidth needed for a specified traffic flow, responsiveness and reliability. Improvements in modem technology will yield some increase in efficiency. Networks that share the same channel are generally more efficient than those that allocate a channel to a transmitter, to a receiver or to a point-to-point link. The greatest improvements in efficiency, however, are likely to result from methods for the compression of images. Noiseless compression can yield about a 3:1 improvement in efficiency. Noisy compression is only appropriate for retrievals but may yield much greater improvements in efficiency. How much image degradation can be tolerated to achieve additional compression awaits further research. Note that the design methods described above still apply for compressed images; the value of the average study size must be reduced according to the image compression achieved. Furthermore, the time for reconstruction of the image from its compressed form must be added to the response time. The application of custom microelectronics to this problem may hold the time for reconstruction to a negligible amount.

Design choices

Figure 6 is a sketch of a general PACS architecture appropriate to the goals of responsiveness, modularity and reliability. Other organizations are certainly possible, but we will describe a rationale for this architecture in the paragraphs to follow.

Two-level interconnect. The physical and functional separation of resources within most large radiology departments suggests a natural subdivision into several local subnets for CT body, CT head, musculoskeletal, chest, gastrointestinal, nuclear medicine and cardiac studies. A second

level interconnect linking these subnets imposes a penalty associated with extra handling of internet traffic, but locality of traffic within a subnet minimizes the effect of this penalty. "Bridges"[31] between the subnets and the second level are utilized that take advantage of a shared address space and a shared protocol throughout the collection of subnets. The address and protocol translation facilities normally associated with "gateways"[31] are minimized.

The problem of optimum partitioning of the PACS into subnets depends principally upon the link channel capacities, upon the mean study size and upon the traffic flow matrix. In theory a partition of all the nodes in a PACS could be selected that minimizes network delay as determined from (9), but generalized to apply to all network links and suitably weighted to account for different response-time criteria for different links. At best the elements of the traffic flow matrix are likely to be known only within an order of magnitude in any practical situation. In fact, in most cases even the number and type of nodes will be indefinitely known. Furthermore, the response-time criterion for individual links is a subjective matter requiring experience with a PACS in order to obtain reliable estimates. Thus partitioning a PACS into subnets may well be more a matter of engineering judgement than it is a matter of mathematical optimization.

Circuit-switched discipline. The large buffers required for the study-at-a-time service needed to implement a store-and-forward discipline would surely escalate the cost of a PACS. Each node would be required to store four or more studies to hold the chances of overflow to acceptable values (see section above entitled Buffer size). We propose that virtual circuits be set up in advance for all study transmissions. This decision requires

uniform channel capacity throughout the virtual circuit. Buffering for a portion of a study will be necessary at a processor serving rotating disk storage to accommodate disk seek times and network acquisition times, but the channel capacity to and from disk should also match the network. We anticipate that the overhead associated with the establishment of a virtual circuit will be small compared to the total transmission time for a study.

Subnets. The subnets themselves can be logically organized into four categories: 1) point-to-point channels, 2) shared-receive channels, 3) shared-transmit channels, and 4) shared-transmit/receive channels (in Figure 7 examples of 1), 2) and 4) are shown). The subnets shown are fully connected and for n nodes require for 1) $\frac{1}{2}n(n-1)$ channels, for 2) and 3) n channels, and for 4) one channel. All but the last subnet need not be fully connected, but if not, provision must be made for the establishment of virtual channels between nodes that are not directly connected.

We are interested in the cases for which the logical channels are obtained primarily by either frequency-division multiplex (FDM) or by time-division multiplex (TDM), but in some cases also by physical separation or by some combination of these methods. Transmission media (coaxial cable, microwave or optical fiber) capable of bandwidths of hundreds of megahertz are now available, but present day microprocessors and memories are not matched to the full bandwidth of these media. Usually FDM or TDM are utilized to provide multiple channels of narrower bandwidth.

The modems required to match the digital source or receiver to the channel are today a significant part of the system cost. In the fully connected point-to-point case $n(n-1)$ modems are required, n-1 for each node. If the modems can switch between frequency bands, time-slots, or physical channels, only a single modem is required for each of the n nodes. For n small,

this _agile_ modem may be more expensive than the n-1 _static_ modems assigned to each node in the first scheme. For n sufficiently large the agile modem must become more economical.

In the shared-receive (shared-transmit) case n modems are required. They need to be able to receive (transmit) from multiple channels, but transmit (receive) on only one. This compromise may be attractive in the FDM and TDM cases. The shared-transmit/receive case utilizes only a single channel so static modems can be used. Note, however, that to accommodate traffic comparable to the other two schemes an n-fold increase in bandwidth of the channel and the modems is required.

It is not clear which of these logical organizations is best for a particular circumstance. For example, consider the situation for which a modem's cost is related to its bandwidth by a power law. If modem cost rises sharply with bandwidth (exponent greater than unity) and ample bandwidth is available on the transmission medium, point-to-point channels may be the most economical choice. This does not necessarily imply $n(n-1)$ distinct physical links, only $n(n-1)$ logical links, perhaps provided by separate FDM channels. On the other hand, modem cost that rises slowly with increasing bandwidth (exponent less than unity) would argue for a single shared transmit/receive channel.

Figure 8 shows the present relationship between cost and bandwidth for FDM modems. Commercial products in the bandwidth range of 2-20 Mb/s are not yet being produced in large volume, but can be expected to become increasingly available as the wide-band modem technology matures. Note, however, the break in the curve at 1 Mb/s. In both cases the exponent is less than one with the steeper slope corresponding to an exponent of about ½. The two-sloped

shape is likely to endure, but the break point may migrate toward higher bandwidths. Above the break point, low production volume and heavy design costs conspire to force prices up. Below the break point well-developed markets, production economies and competition force prices down. Under these circumstances it may be true that a compromise approach like shared-receive will be most economical.

Conclusions

There is no way to predict precisely what the number of radiology studies will be in the years to come, but certainly radiology volume is increasing. Particularly striking is the rate of growth of the digital imaging modalities, and we believe this trend toward digital image growth will continue. But to obtain the maximum clinical advantage of storing images in digital form, a digital imaging network is required to distribute images so that all of a patient's clinical image data is available for immediate review.

A proposed architecture has been suggested to meet the design goals of such a network. It has been designed against the background of a set of design equations for PACS and estimates of current image generation and retrieval activity at MIR.

Some of the conclusions of this analysis include the following:

1. Systems with great bandwidth will be required to handle the prompt retrieval of multi-image studies during peak periods of retrieval activity. Although early retrieval of studies in anticipation of batch review may appear to reduce network bandwidth requirements, interactive image retrieval with response times comparable to those of any interactive computer system sets a more demanding requirement on network bandwidth. Maximum retrieval response times of a few seconds mandate high-speed transmission.

2. Protocols for error detection and retransmission can be accomplished at little cost given currently achievable bit error rates.

3. Of the various network architectures considered, a store-and-forward discipline seems ill-suited to PACS design. By comparison, local networks with a shared transmission medium have some important advantages for picture service.

4. Modular network architectures, capable of reacting to growth and changes in traffic patterns are essential. Partitioning the network into subnets will reduce the utilization factor and decrease response time by distributing network traffic.

5. Many further developments are necessary including the physical and conceptual design of an archiving system capable of supporting optimal response times, the evaluation of suitable data compression algorithms, the design of inexpensive image display stations, the availability of inexpensive, wide-band modems, and the design of efficient network interface units to buffer data as it is transferred between the network, source, and archive.

The layered approach to the design of computer networks has been applied to a picture archiving and communication system for radiology. The first three layers (physical, picture link, and picture network) have been analyzed. The resulting equations have been applied to typical data obtained from the Mallinckrodt Institute of Radiology. The two-level interconnect described seems well-suited to the problem, but much work remains to be done.

Acknowledgements

discussions. We also thank the students in our 1980 graduate seminar on computer networks. They contributed much to our thinking through their study of PACS as a term project.

References

1. A. S. Tannenbaum, <u>Computer Networks</u>, Prentice-Hall, New Jersey, 1981.

2. P. E. Green, Jr., editor, <u>Computer Networks Architectures and Protocols</u>, Plenum Press, New York, 1982.

3. J. R. Cox, Jr., G. J. Blaine, R. L. Hill, and R. G. Jost, "Study of a Distributed Picture Archiving and Communication System for Radiology," <u>Proc. of the 1982 PACS Conference</u>, Vol. 318, 1982.

4. D. Meyer-Ebrect, "The Management and Processing of Medical Pictures: An Architecture for Systems and Processing Devices," <u>Proc. of the Workshop Picture Data Description and Management</u>, IEEE, 1980, pp. 202-206.

5. S. J. Dwyer, A. W. Templeton, W. H. Anderson, and M. A. Tarlton, "A Distributed Diagnostic Image Information System," <u>Proc. of the 1982 PACS Conference</u>, SPIE, Vol. 318, 1982.

6. R. L. Arenson, D. E. Morton, and J. W. London, "Fiber Optic Communication System for Medical Images," <u>Proc. of the 1982 PACS Conference</u>, Vol. 318, 1982.

7. M. P. Capp, "The Photoelectronic Radiology Department," <u>Proc. Conference on Digital Radiology</u>, SPIE, Vol. 314, 1981.

8. A. J. Duerinckx and E. J. Pisa, "Filmless Picture Archiving and Communication in Diagnostic Radiology," <u>Proc. of the 1982 PACS Conference</u>, Vol. 318, 1982.

9. J. A. Vos, "Megadoc, a Modular System for Electronic Document Handling," <u>Philips Technical Review</u>, Vol. 39, No. 12, 1980.

10. D. G. Herzog, "Optical Disc Approaches as Mass Storage," <u>IEEE</u>, 1980, pp. 38-41.

11. J. Nadan, G. Blom, S. Chandra, G. Kenney, and R. McFarlane, "Recent Advances in Digital Optical Recording," <u>Proc. of the 1982 PACS Conference</u>, SPIE, Vol. 318, 1982.

12. R. L. Benton, "Storage Technology for the 1980s," <u>Proc. of the 1982 PACS Conference</u>, SPIE, Vol. 318, 1982, pp. 60-64.

13. M. J. Flynn, J. S. Newman, and R. M. Mares, "Digital Storage Compression for Diagnostic Images," <u>Proc. of the 1981 Conf. on Application of Optical Instrumentation in Medicine</u>, 1981, pp. 357-368.

14. J. G. Dunham, R. L. Hill, G. J. Blaine, D. L. Snyder, and R. G. Jost, "Compression for PACS in Radiology," <u>Proc. of the 1983 PACS Conference</u>, SPIE, Vol. 418, 1983.

15. K. T. Lee, S. S. Sagel, and R. J. Stanley, <u>Computed Body Tomography</u>, Raven Press, New York, 1983.

16. B. L. Holman and J. A. Parker, <u>Computer-Assisted Cardiac Nuclear Medicine</u>, Little, Brown and Co., Boston, 1981.

17. D. A. Sarti and W. F. Sample, <u>Diagnostic Ultrasound</u>, G. K. Hall, Boston, 1980.

18. C. A. Mistretta, A. B. Crummy, C. M. Strother, and J. F. Sackett, <u>Digital Subtraction Arteriography</u>, Year Book Medical Publishers, Inc., Chicago, 1982.

19. C. L. Partain, A. E. James, F. D. Rollo, and R. R. Price, <u>Nuclear Magnetic Resonance Imaging</u>, W. B. Saunders, Philadelphia, 1983.

20. M. M. Ter-Pogossian, M. E. Raichle, and B. E. Sobel, "Positron-Emission Tomography," Scientific American, Vol. 243, No. 4, 1980.

21. Proc. of International Workshop on Time-of-Flight Tomography, Washington University, St. Louis, MO., IEEE Computer Society Press, Maryland, 1982.

22. H. Zimmerman, "A Standard Layer Model," Computer Network Architectures, Plenum Press, New York, 1982, Ch. 2.

23. M. Abromowitz and I. A. Stegun, Handbook of Mathematical Functions, Dover, 1965, p. 929.

24. E. G. Rawson and R. M. Metcalfe, "Fibernet: Multimode Optical Fibers for Local Computer Networks," IEEE Trans. Commun., Vol. COM-26, 1978, pp. 983-990.

25. W. Chou, editor, Computer Communications: Volume I Principles, Prentice-Hall, Inc., New Jersey, 1983, p. 276.

26. R. M. Metcalfe and D. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," Communications of ACM, Vol. 19, No. 7, July 1976, pp. 395-404.

27. C. Tropper, Local Computer Networks Technologies, Academic Press, New York, 1981, Ch. 2.

28. R. B. Cooper, Introduction to Queueing Theory, Macmillan, 1972.

29. L. Kleinrock, Queueing Systems, Vol. 1, Theory, John Wiley, 1975.

30. N. M. Dor, "Guide to the Length of Buffer Storage Required for Random (Poisson) Input and Constant Output Rates," IEEE Trans. Elec. Comp., 1967, pp. 683-684.

31. D. D. Clark, K. T. Pogram, and D. P. Reed, "An Introduction to Local Area Networks," Proc. IEEE, 1978, pp. 1497-1517.
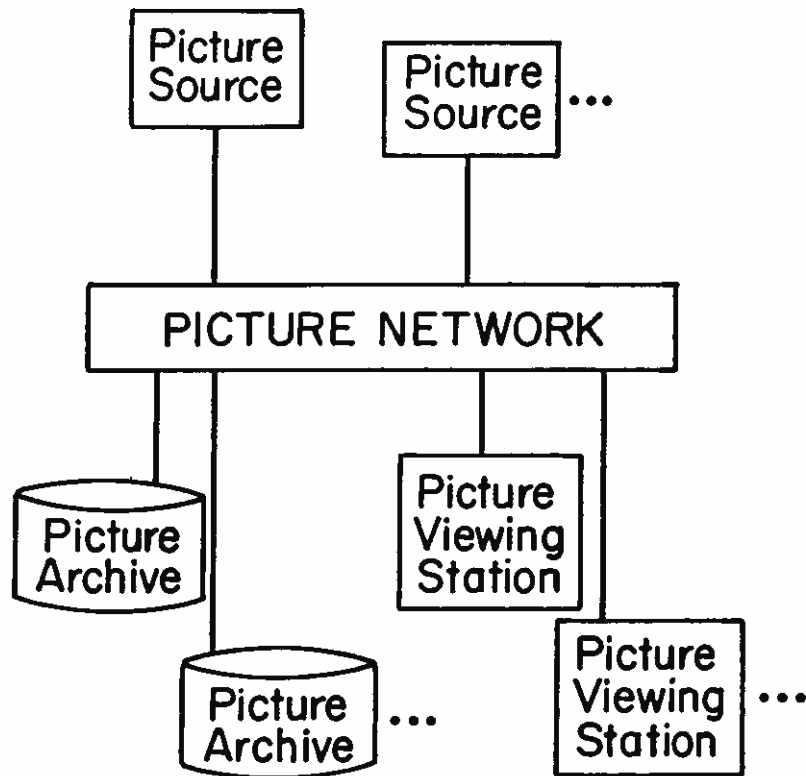
Figure 1.

Figure 1. Picture archiving and communication system (PACS) block

diagram. The picture network must accommodate sources

representing a variety of imaging modalities, support

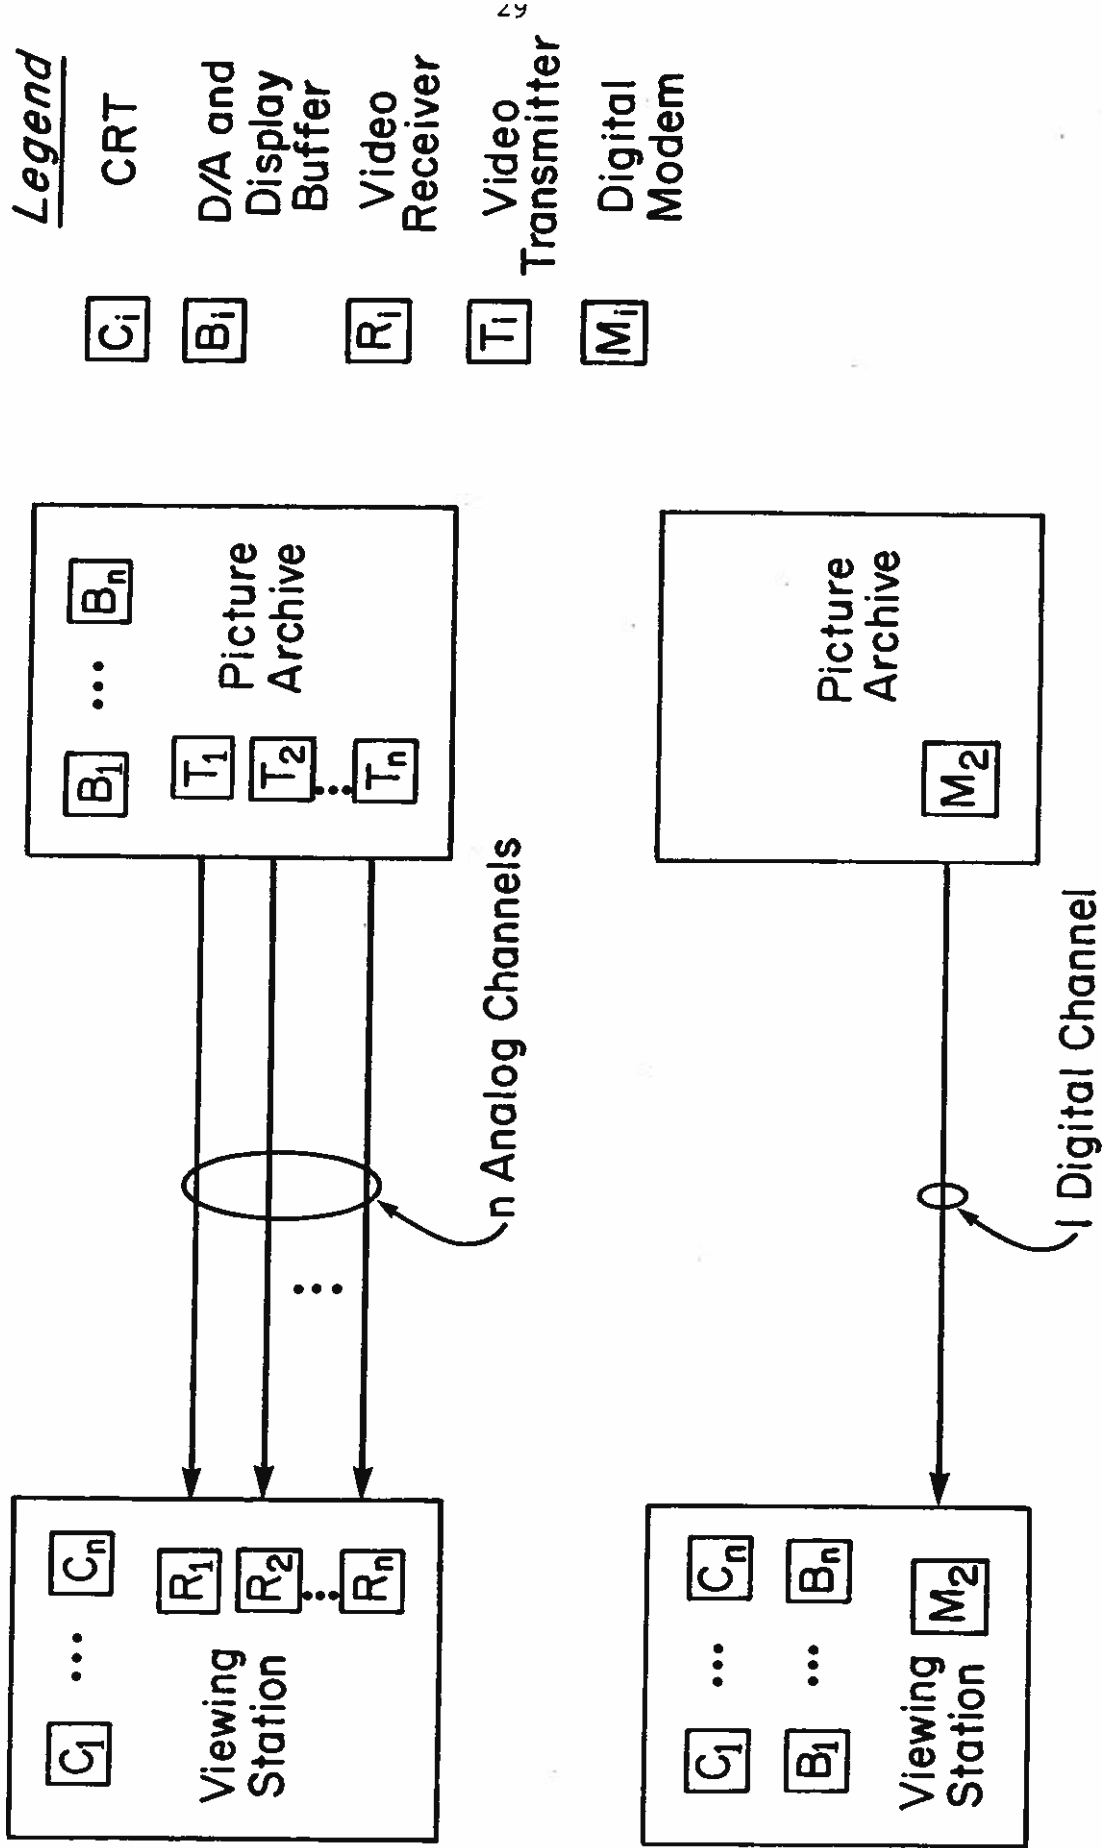various classes of viewing stations and facilitate system

growth.

Legend

CRT — C_i

D/A and Display Buffer — B_i

Video Receiver — R_i

Video Transmitter — T_i

Digital Modem — M_i

Picture Archive

B_1 ... B_n

T_1 T_2 ... T_n

Picture Archive

M_2

---

n Analog Channels

...

1 Digital Channel

Viewing Station

C_1 ... C_n

R_1 R_2 ... R_n

Viewing Station

C_1 ... C_n

B_1 ... B_n

M_2

Figure2.

Figure 2. Comparison of analog and digital methods for displaying

n images at a viewing station remote from the picture archive.

The viewing station may be thought of as an electronic

view box.
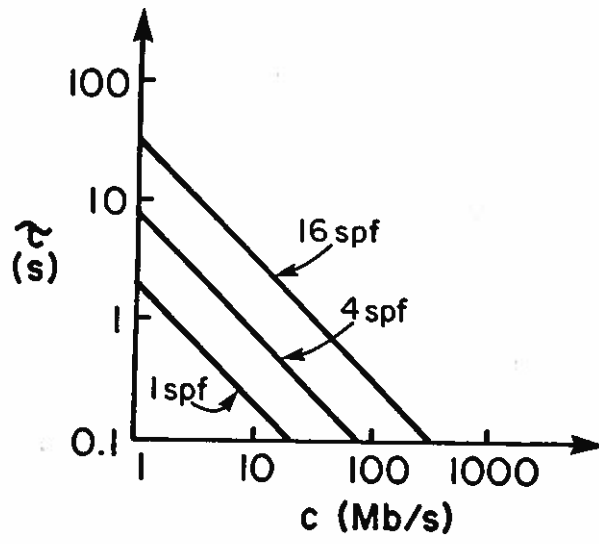
Figure 3.

Figure 3. Transmission time in seconds versus channel capacity in
megabits per second for the physical layer. Examples
are plotted for selected picture sizes where one standard
picture frame (spf) corresponds to 512 x 512 pixels of 8
bits each; 4 spf corresponds to 1024 x 1024 pixels of 8
bits each; and 16 spf corresponds to 2048 x 2048 pixels of
8 bits each.
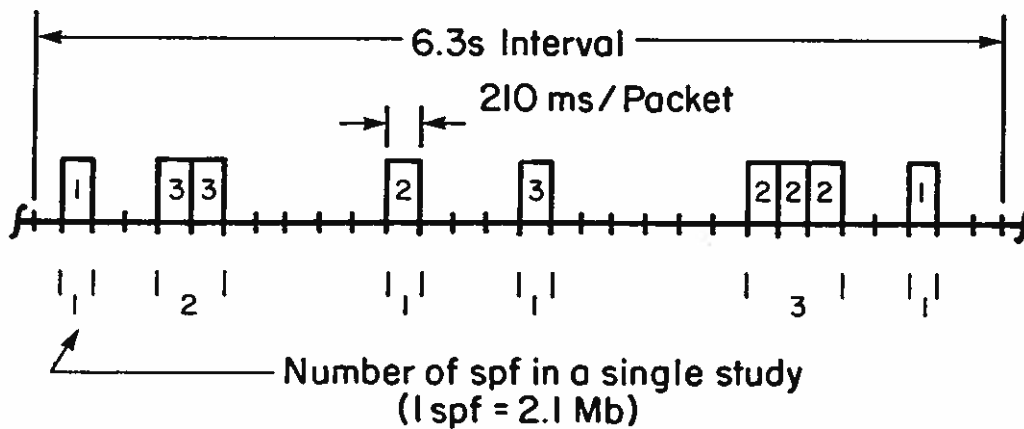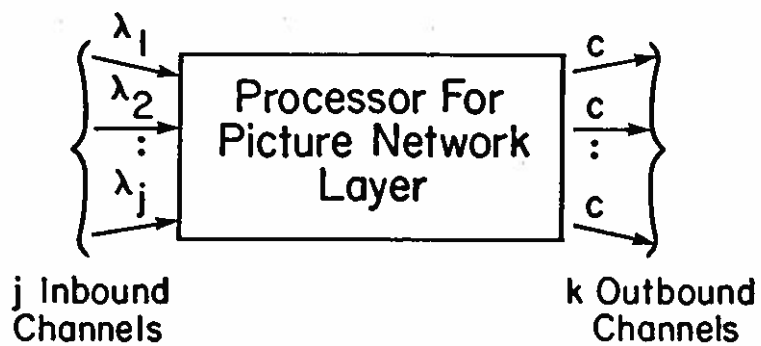
Figure 4

Figure 4.   a) Single network node modeled at the picture network layer.

The node processor services j inbound channels which contend

for use of the k outbound channels.   b) Example transmission

over a channel with C = 10 Mb/s, j=3 and τ_L = 210 ms.   The

channel number j is indicated within each 1 spf packet.

For the 6.3s transmission interval shown, the average study

size $\overline{m}$ = 1.5 spf, $\lambda_1 = \lambda_2 = \lambda_3 = 2/6.3$ and $\lambda = \lambda_1 + \lambda_2 + \lambda_3$

= 6/6.3 studies/s.   The utilization factor $\rho = \lambda\ \overline{m}\ \tau_L = (\frac{6}{6.3})$

$(1.5)(.21)$ = 0.3.   This is the same result that would be

obtained by calculating the fraction of time the channel
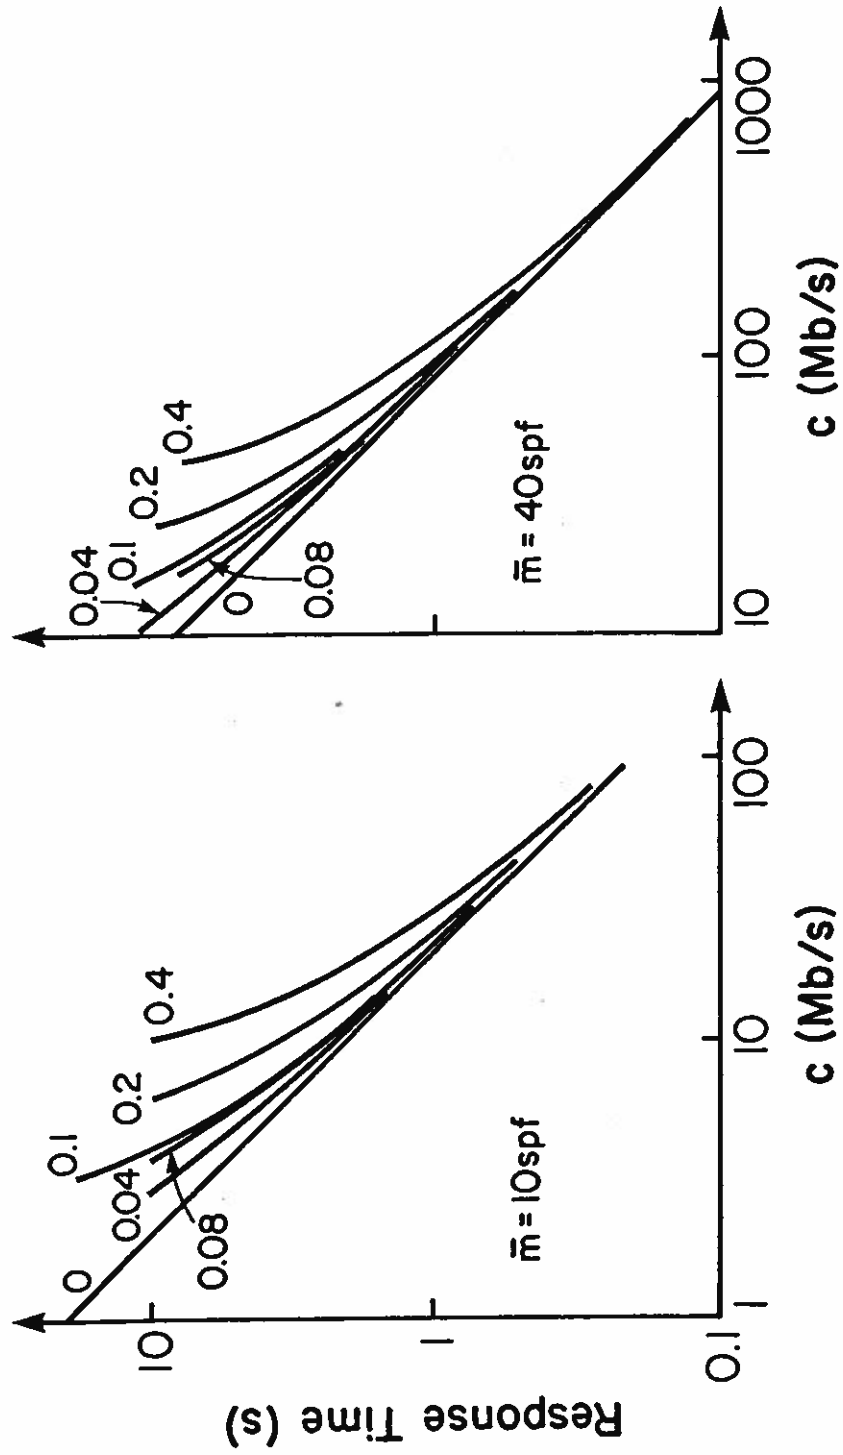
is busy, $\rho = \frac{9}{30}$ = 0.3.

Figure 5.

Figure 5.   Response time of picture network level versus channel

capacity for various values of source rate λ in studies per

second.   Figure 5a is for a mean study size of 10 spf (about

20 Mb) while Figure 5b is for a mean study size of 40 spf

(about 80 Mb).   Peak values of study generation rate are

now estimated at $\lambda = .04$ while peak values of the sum of

the study generation and interpretation rates are estimated

at $\lambda = .2$.   Future levels of activity may push rates to
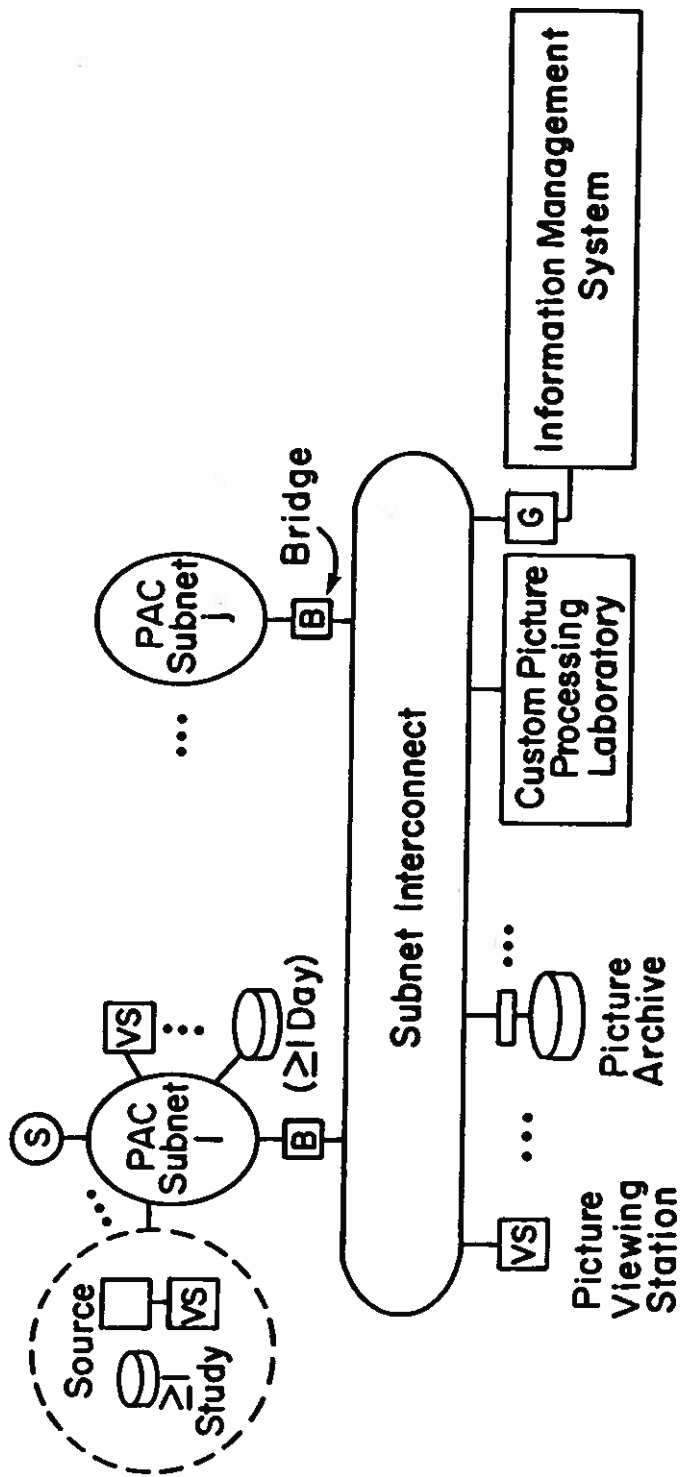
$\lambda = .4$ and beyond.

Figure 6

Figure 6. Block diagram of a general two-level PACS architecture.

39

Point-to-point
Channels

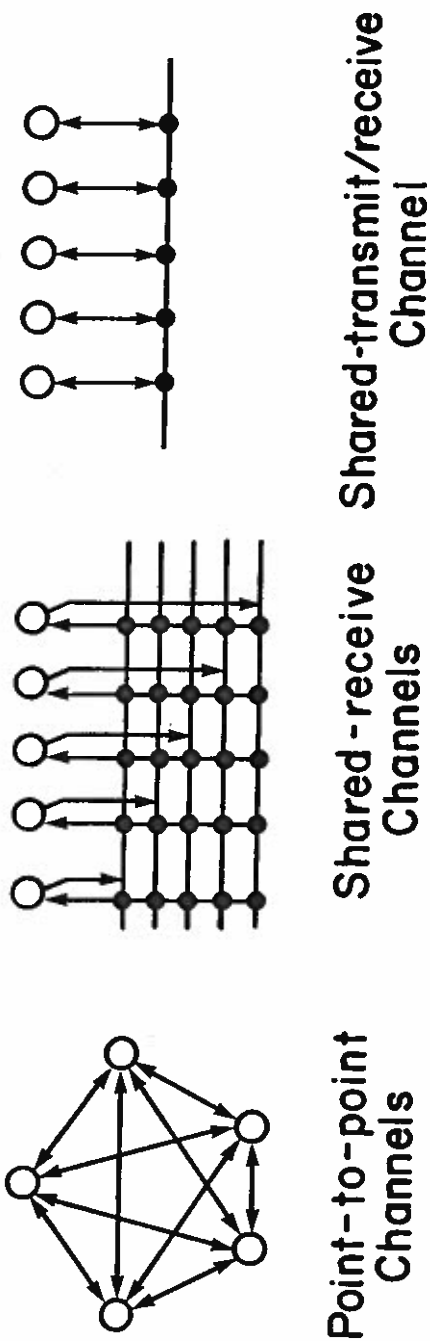Shared-receive
Channels

Shared-transmit/receive
Channel

Figure 7.

Figure 7. Categories of subnet logical organization: point-to-point, shared-receive, shared-transmit and shared-transmit/receive. The logical channels may be achieved by partitioning in frequency, time or space. Only fully connected subnets are shown. Shared-transmit is not shown, but is the same as the shared-receive with directions of packet flow reversed.
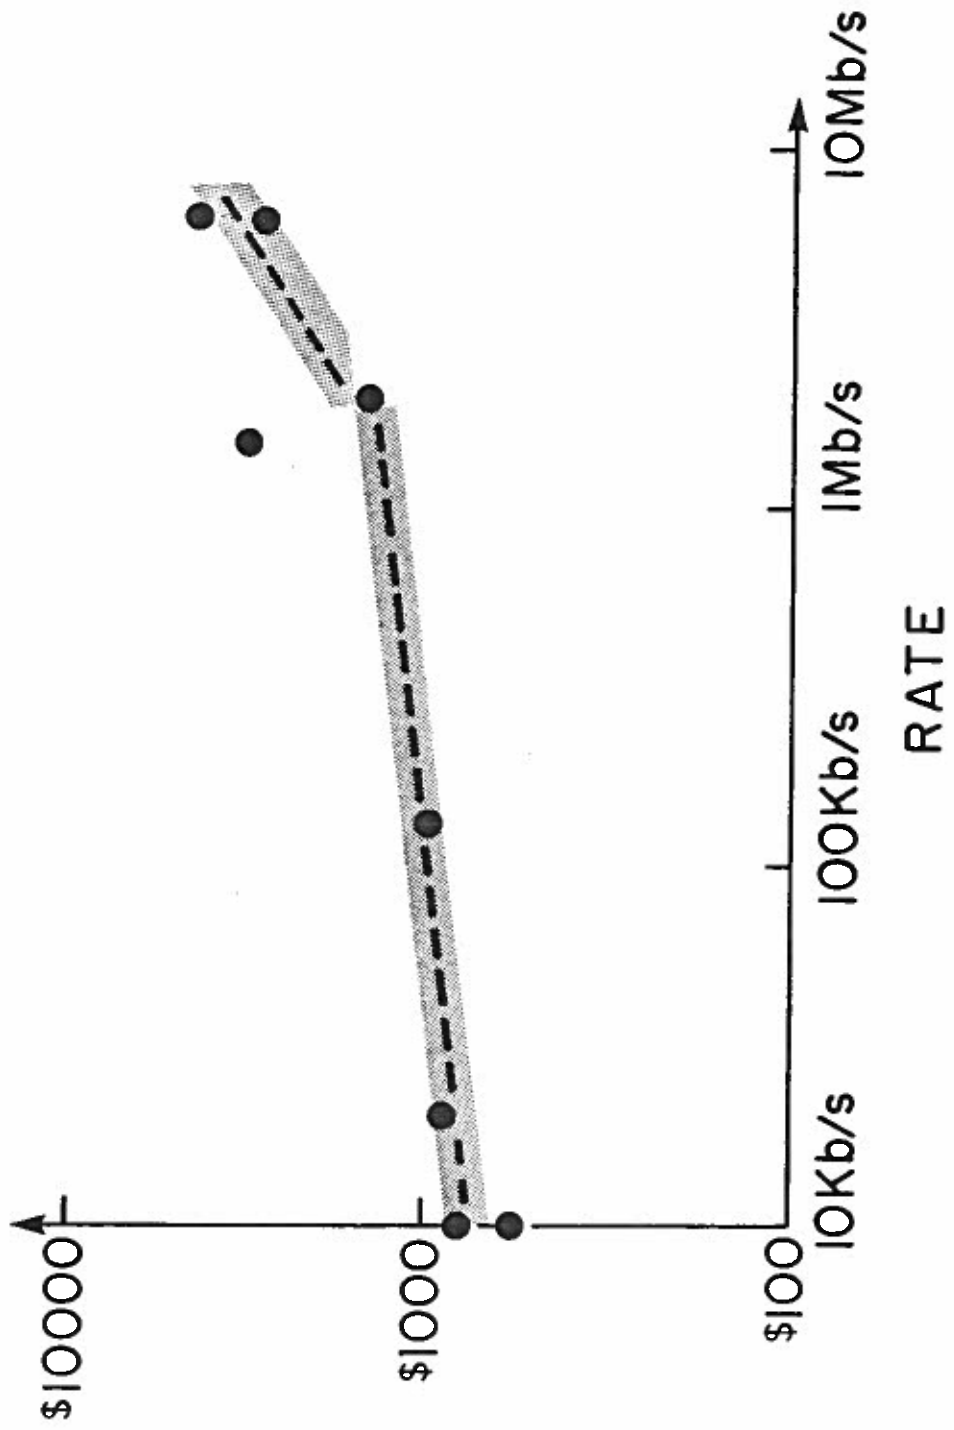
Figure 8.

Figure 8. Approximate relationship between cost and digital bandwidth
for FDM modems.