THE UNIVERSITY OF MICHIGAN

COMPUTING RESEARCH LABORATORY[1]

MERIT'S EVOLUTION – STATISTICALLY SPEAKING

Eric M. Aupperle

CRL-TR-4-84

JANUARY 1984

Room 1079, East Engineering Building
Ann Arbor, Michigan 48109
USA
Tel: (313) 763-8000

[1]Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agency.

This report also appeared in *IEEE Transactions on Computers*, **Vol.C-32,NO.10 (October), 1983.**

Merit's Evolution—Statistically Speaking

ERIC M. AUPPERLE, SENIOR MEMBER, IEEE

Abstract—Merit links four Michigan universities with a packet-switching computer network. This network provides several interactive and batch services to the students, faculty, and researchers of these schools. While the network and its services are reviewed, this paper focuses on the statistical performance of Merit during its first ten operational years. Numerous graphs and tables present data including time-series showing the growth of usage, transmission-rate related measures, normalized measures on a per-connection basis, and the distribution of activity among the hosts. The data reveal that on an aggregated basis the current average packet size is nearly 50 bytes with a corresponding transmission rate of 1350 packets per hour for each virtual connection. While these are the aggregate averages, there is considerable variation of these values among the different services Merit offers. The paper concludes with explanations for these differences and general observations about Merit's usage patterns.

Index Terms—Data communications, interuniversity networking, network statistics, packet switching, regional network, resource sharing.

I. INTRODUCTION

THE Merit Computer Network is a packet-switching network developed by three Michigan universities: Michigan State University, the University of Michigan, and Wayne State University. The network became operational in July 1972 and has operated continuously since then. Throughout this period, the network's functional facilities have been enhanced and its performance refined. In January 1979 Western Michigan University became the fourth Merit member.

The first section of this paper offers an overview of the network's current characteristics and capabilities. This background information provides the context for understanding the main topic of the paper, the network's usage statistics. The next section describes statistical information that have been recorded throughout the network's operational life and how these data are processed. This is followed by several sections that present various statistical summaries and data represented in graphical and tabular forms. In the final section the author offers several observations about these network data.

II. AN OVERVIEW OF THE MERIT COMPUTER NETWORK

A. The System

The network connects eight separate university computers in southern Michigan, most of which are large combined time-sharing and batch-processing systems. The communication and processing features of these systems are portrayed in Fig. 1.

MSU: The Michigan State University (MSU) host computer is a Control Data Corporation (CDC) Cyber 170 Series 750 with 131 000 60-bit words of central memory and 500 000 words of extended core storage. It runs under a locally-modified version of CDC's NOS/BE operating system and supports nine remote-batch stations and 110 terminal ports in addition to its network ports. MSU also has a CDC 6500 with 98 000 words of central memory that does not have a direct connection to the network.

Michigan State's CDC hardware features a 60-bit word, in contrast to the 32-bit word of the Amdahl hardware at the University of Michigan and WSU and to WMU's 36-bit word on their DEC system. Thus, one of MSU's principal contributions to the network is as a supplier of extended-precision numerical and scientific computing. MSU also offers the CDC program-product versions of such languages as Fortran, Basic, and APL as well as CDC versions of SPSS and Pascal, all of which feature extensions to the corresponding implementations of these languages at the other hosts.

UM: The University of Michigan (UM) has an Amdahl 5860 with sixteen million bytes of real memory running under a locally-written virtual-memory operating system, MTS (the Michigan Terminal System). The UM host supports twelve remote batch stations and approximately 800 interactive terminal ports.

The Amdahl 5860 provides a fast and powerful network host that is also—like the WSU system—able to manipulate very large data bases. Packages such as MIDAS, an interactive statistical system, that have been written at UM to take advantage of this combination, account for a good portion of network usage at UM.

Two other UM hosts are an IBM System 370/148 used for an experimental campus-wide word-processing support service and a DEC VAX 11/780 running UNIX used for image processing and robotics research and education. These computing systems and a third, a recently added Prime 750, are primarily used for research and their usage is not reflected in data presented in this paper.

WMU: The host computer at Western Michigan University (WMU) is a DECsystem 1099 with two KL10 processors and 512K 36-bit words of main storage. It runs DEC's TOPS-10 operating system, the only host on the network to offer interactive service on an unmodified vendor system.

WSU: The principal host computer at Wayne State University (WSU), which also runs MTS, is an Amdahl 470V/6

Manuscript received January 29, 1980; revised March 28, 1983.
The author is with the Merit Computer Network, 2200 Bonisteel Blvd., Ann Arbor, MI 48109.

0018-9340/83/1000-0881$01.00

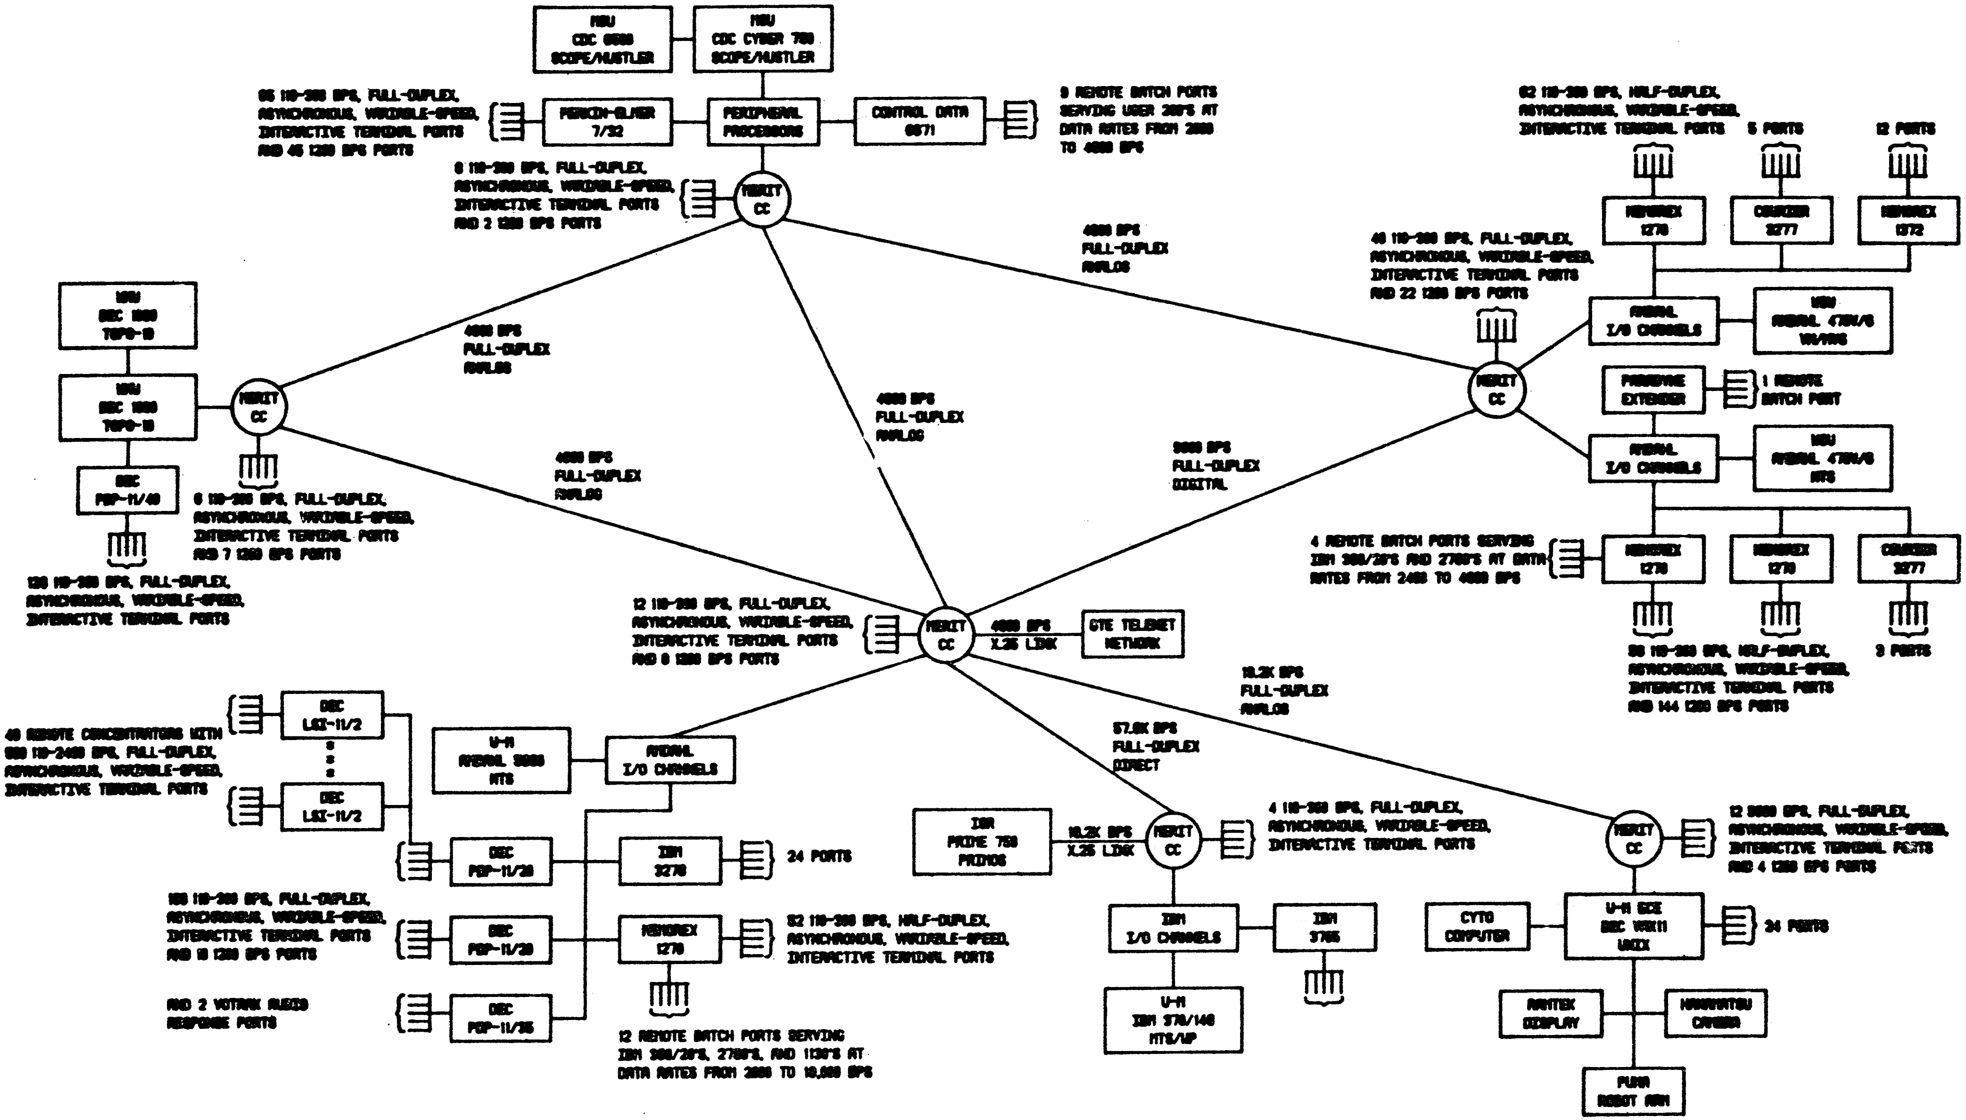

Fig. 1. System configuration of the Merit Computer Network.

with eight million bytes of solid-state memory. The system supports 282 interactive terminal ports and five remote batch stations. The WSU Computing Services Center emphasizes the production of high quality printed copy for use in data tables, graphs, and even the full texts of technical reports.

A significant and unique network resource at WSU is IBM's MVS, which runs on a second Amdahl 470V/6 and serves their administrative data processing needs. While MVS is not yet accessible interactively through the Merit network, it does provide batch service to network users. This makes available a number of program products written for IBM's OS or MVS that are not available at the other hosts.

Augmenting these host systems is the Merit Network. The technological design of the Merit Computer Network employs communications switching computers which act both as interfaces between the host computers and the network and as packet-switching nodes. One or more Communications Computers (CC) is located at each host site. The six CC's are interconnected by means of modems through leased telephone lines as shown in Fig. 1. The interfaces that the CC's present to the host systems are uniquely adapted to the requirements of each particular host; the interfaces that the CC's present to the network (and thus to each other) are identical.

One of the CC's is a Digital Equipment Corporation PDP-11/20's with 28K 16-bit words of memory. Four others, the ones on the UM and WSU campuses, are PDP-11/60's, each with at least 32K words of memory. The sixth, at WMU, is a PDP-11/40 also with 32K words of memory. The primary functions of the Communications Computers are to multiplex virtual connections over physical communications links (communicating with other nodes using the HDLC link-level protocol); to receive messages from their respective host computers for transmission on the network; to receive messages from other nodes for delivery to their hosts or for forwarding to another host; and to support local dial-up terminals.

The PDP-11 of each CC is augmented with several interfaces as shown in Fig. 2. The major ones are the host and the asynchronous and synchronous communications interfaces. The synchronous communications interfaces are designed to transfer entire messages over the telephone network directly to and from the PDP-11's memory without program intervention, once the transfer is initialized by software action. Each one is controlled by its own microcomputer. Currently, two or more of these synchronous ports are used at each node.

The synchronous interfaces are connected to AT&T Model 208A Data Sets except for the Ann Arbor-Detroit link, which uses Dataphone Digital Service, and the internal UM links which use limited-distance modems or direct connections. The data transmission is full-duplex at rates from 4800 bits/s to 57.6 kbits/s as shown in Fig. 1. The interfaces are designed to operate at rates in excess of 500 kbits/s, however, when equipped with higher data-rate modems. Each asynchronous interface provides for up to 32 local dial-up ports. Each port is equipped with either an auto-baud detect modem for operation up to 300 bits/s or with a 1200 bits/s modem for operation at that data rate.

All but one of the host interfaces are special hardware devices that allow for the parallel transmission of data records to and from the memory of the associated CC's, performing whatever memory alignment operations are required for word lengths which differ from that of the CC. The multiple-address host interfaces simplify the logical structure of multiplexing user activity over the network and thereby reduce the amount of network support software required in the host computers.

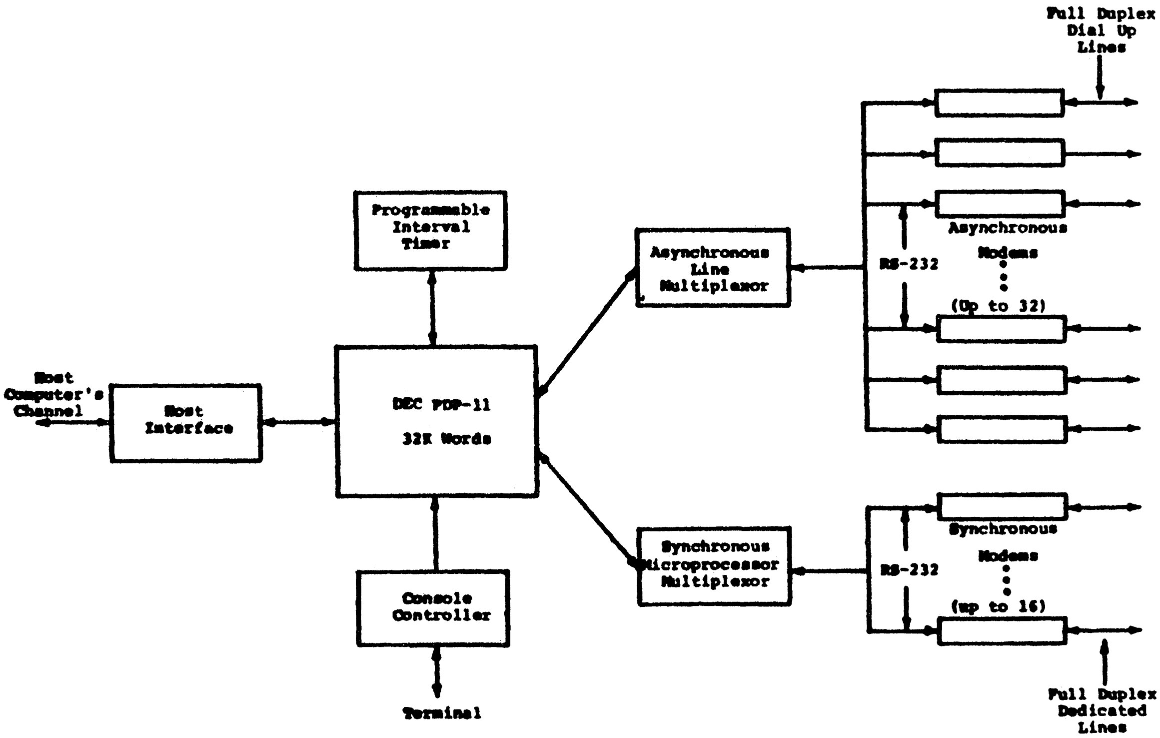

Fig. 2. The communications computer hardware.

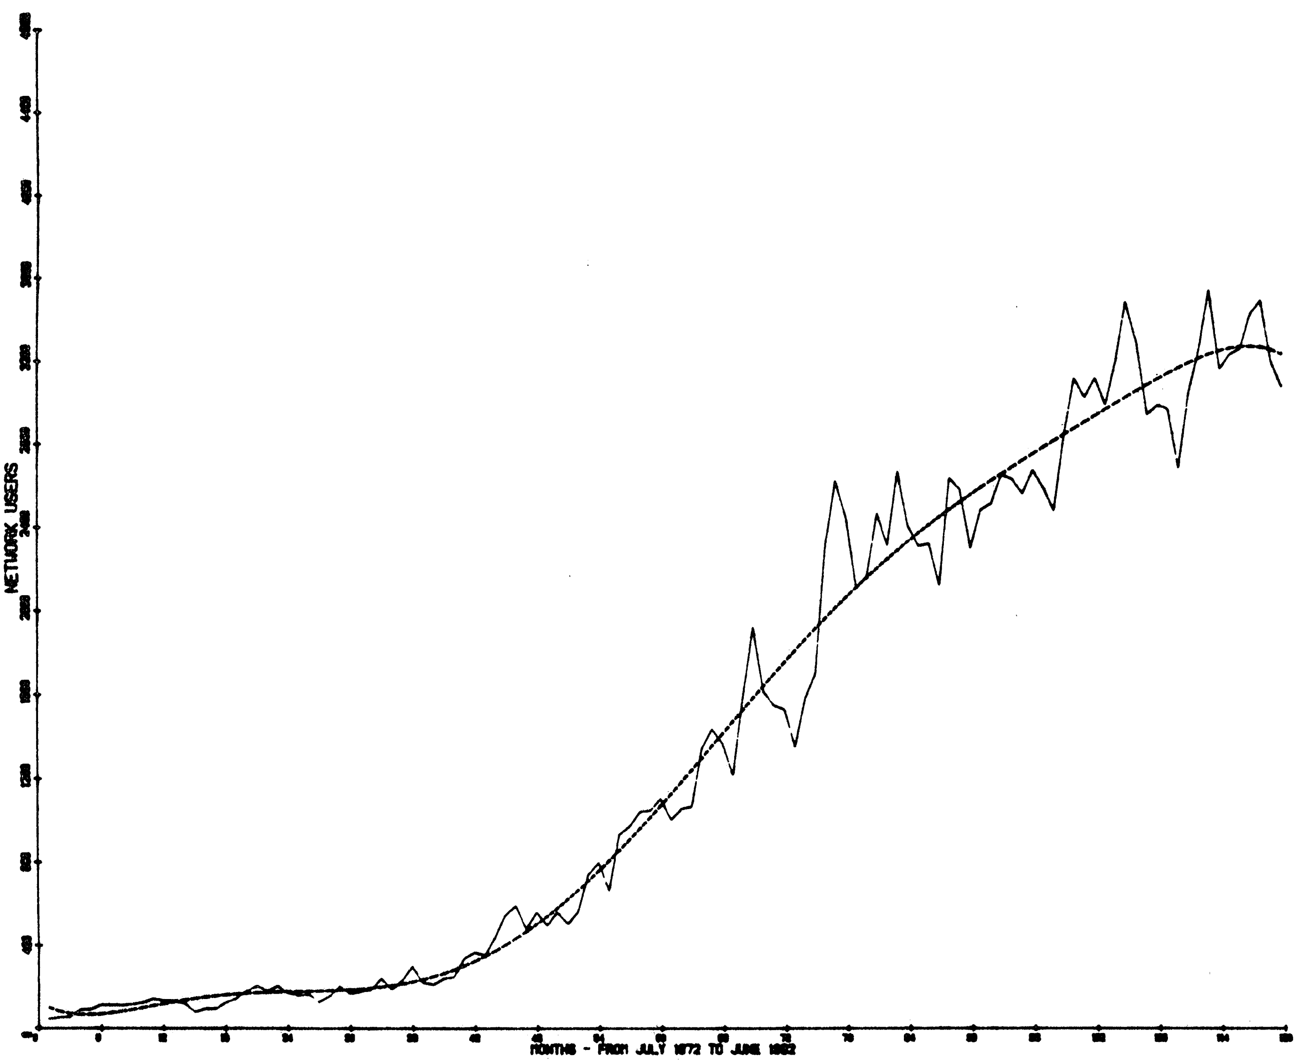

Fig. 3. Number of different network users per month.

Additional hardware details appear in [1] and [4]. The Prime 750 host is interfaced to the network through a serial X.25 communications link.

It must be emphasized that the Merit network augments the existing terminal and batch access facilities provided by each host system. The data reported in subsequent sections relate only to traffic on the network rather than to the total access activity provided by the hosts through their local front-end systems.

B. Types of Network Service

The network offers users of any one computing center access to the facilities of the others. To provide this access the computers at each center, called *hosts*, are interconnected through the network. The user communicates information between hosts through a *virtual connection*, more simply called a *connection*. This may be an information path from one host to another dedicated to the user. It may also be from Hermes, the network's terminal access facility, to a serving host. The initiating end of a connection is the *local* host and the other end is a *remote* host. The exact appearance of an interhost connection depends upon the local host's operating system, the type of service obtained from the hosts, and the type of host access employed by the user.

There are four types of network access: host-to-host interactive, direct terminal, external, and network batch. Brief descriptions of each follow.

Host-to-Host Interactive Access: This service requires that a user (or a user program) establish a connection from a local host to a selected remote host. The resources at the remote host that are made available by the network connection are the same as those available through a terminal directly attached to that host. The user (or a user program) may now use these resources and does so by issuing commands, etc., in the manner appropriate to the remote host. In general, however, the user is able to use the terminal in the manner usual to the local host. This service has four subtypes.

Classic Interactive: the interaction possible among two or more of the network hosts without the expense of normal long-distance telephone transmission. The network reduces the rate of communication error and makes it possible for users to work with the local terminal operating conventions and keyboard editing modes to which they are accustomed.

Enhanced Interactive: which means simply that data

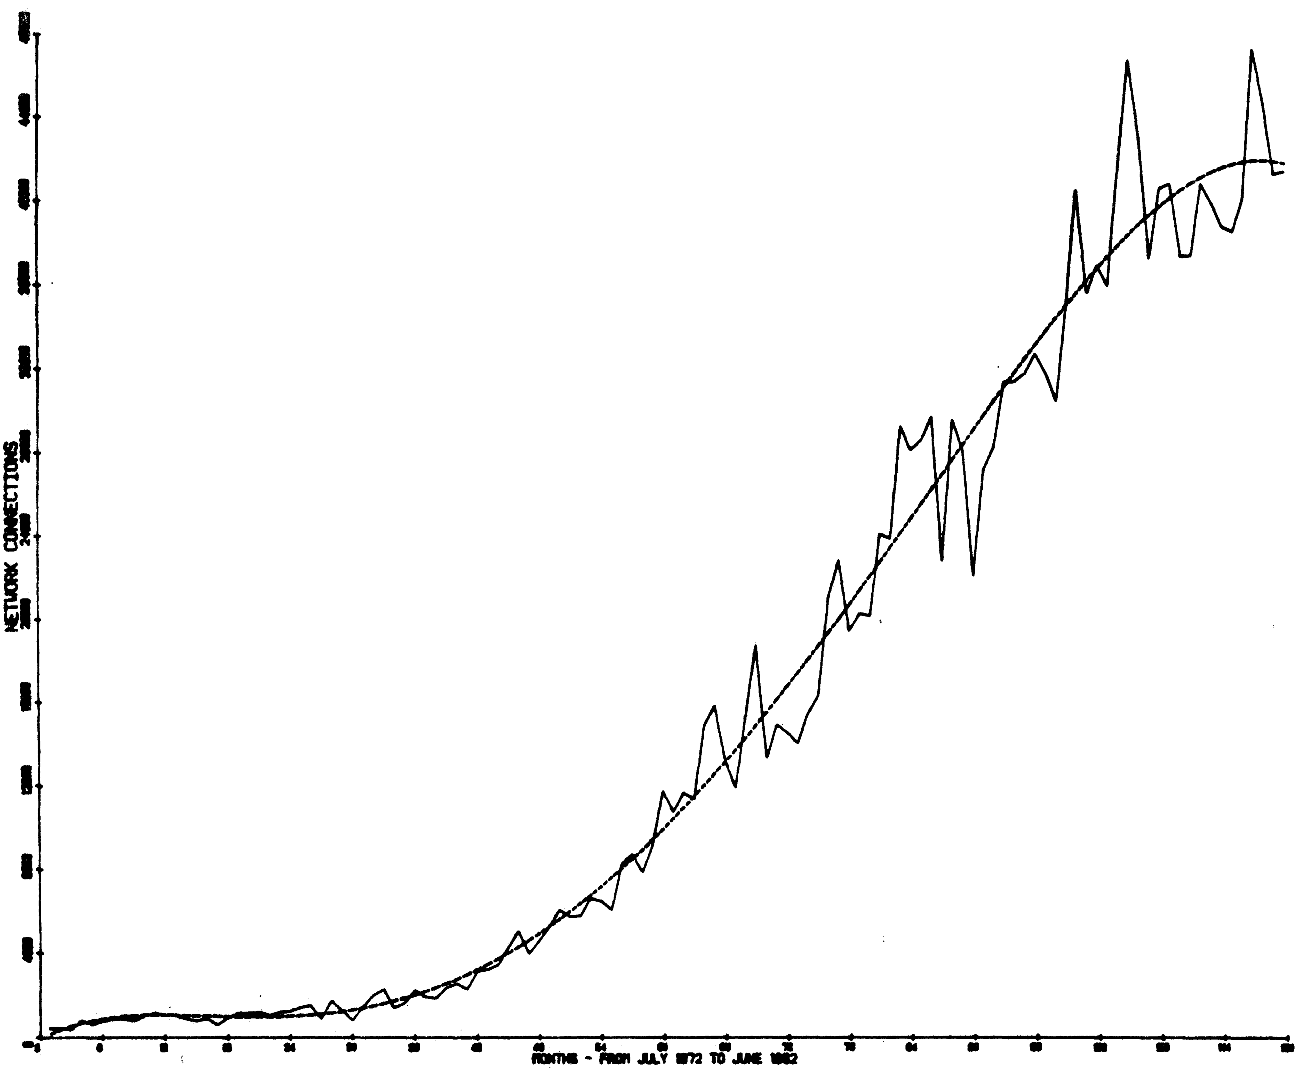

Fig. 4. Number of successful network connections.

bases from one of the host computers can be used with an executing program at another host, or on more than one additional host. It is also possible to display or print out data from a network task at any of the hosts.

File Transfer: which allows data to be copied from one host to another, whether from disk to disk, disk to tape, disk to printer, or in any of several other combinations.

Interprocess Communication: the process of running programs on one or more host computers concurrently; programs may be employed simultaneously or sequentially, and they may transfer data from one host to another.

Direct Terminal Access: This is the capability for dialing directly into the network from a terminal and accessing a remote host without going through a local host. This facility is provided by Hermes, Merit's network-to-terminal interface, which runs in each of the individual Communications Computers.

External Access: This is implemented through an X.25 link operating at 4800 bits/s between the GTE Telenet[1] network and the Merit Network. This service currently allows up to 32 users to concurrently access any of the Merit hosts through GTE Telenet.

[1] "GTE Telenet" is a registered service mark of GTE Telenet Communications Corporation.

Network Batch Access: This service allows a user to submit an execute, print, or punch job at any of the network's hosts. The job may be, and usually is, initiated from a host other than the one where it is to be processed. Establishing the connections for transmitting the job and for retrieving any output is done by the Network Batch Service.

Remote Job Entry: allows any batch job to be submitted either in the form of a card deck or from a terminal. The system where the job is to be processed need not be operating at the time. Network Batch Service will arrange for the job's transmission when the requisite host computers are operating.

Batch File Transfer: allows data to be copied from one host to another, e.g., from disk to disk, disk to tape, or disk to printer in batch mode. This form of file transfer is often cheaper and faster than interactive file copying.

III. A Description of the Network's Statistical Logging and Processing Procedures

Each time a virtual connection either enters or leaves a

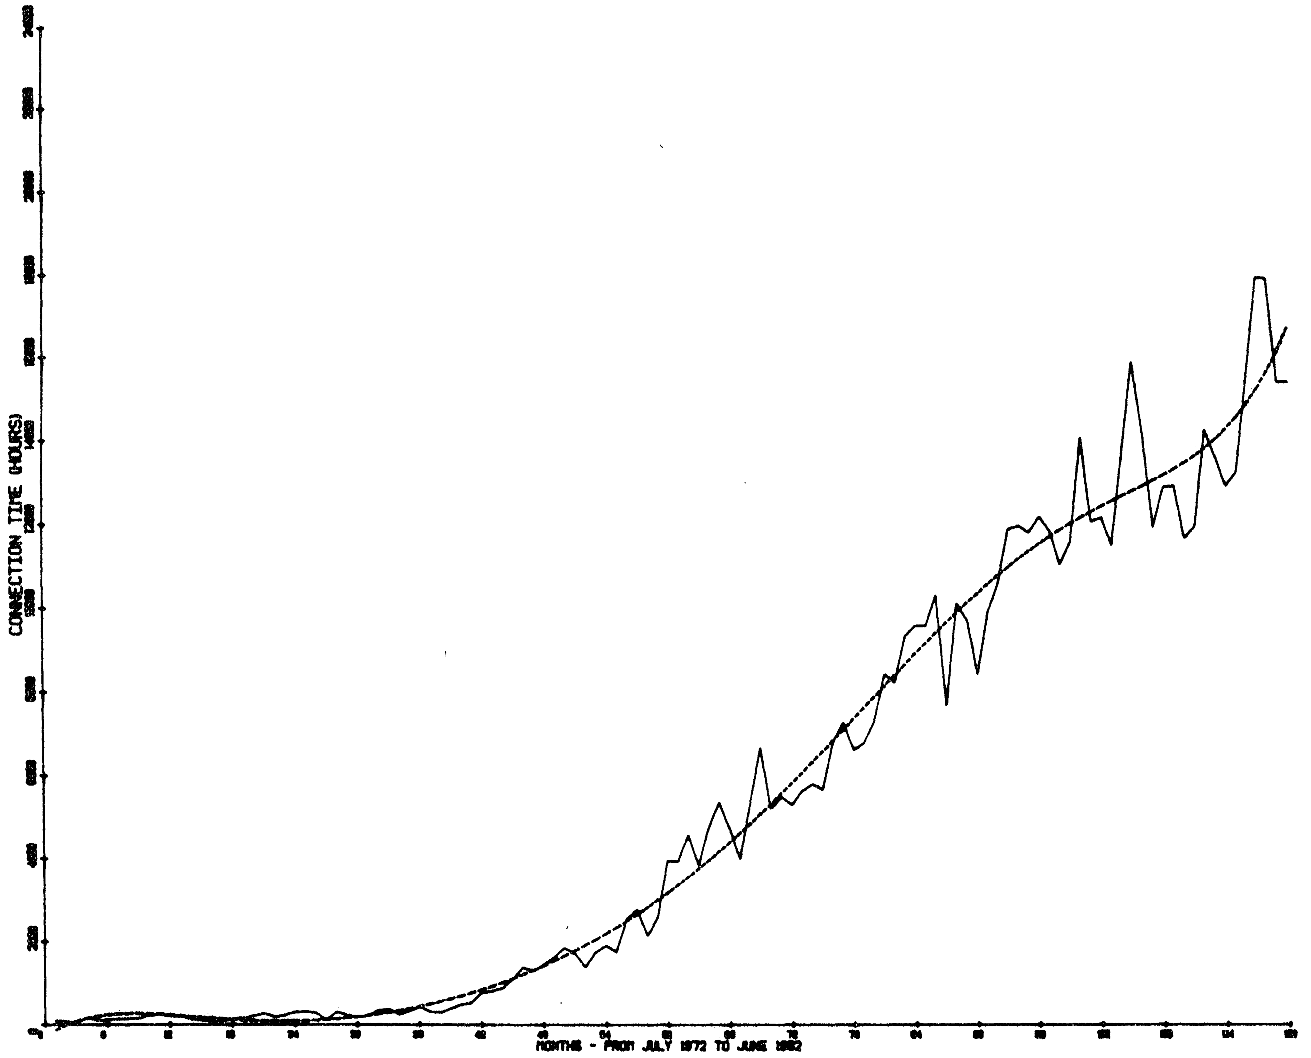

Fig. 5. Elapsed connection time (hours).

serving host, data about this connection are accumulated by the host computer. When a connection finally terminates, the connection's statistics are logged by the host. The recorded information includes the following:

- date and beginning time of connection
- duration of connection
- connection's initiating and receiving hosts
- connection type
- a code indicating whether the connection was successfully opened or the reason for its failure
- both the number of packets and the number of bytes sent and received during the existence of the connection
- user's account ID.

These data are continuously accumulated in on-line host files. Once each month the contents of these files are stored on tape and the files emptied. The monthly aggregates from all hosts are combined and processed by a comprehensive set of statistical analysis programs. The results of these analysis programs are used to prepare various network usage level reports.

Most of the statistical record parameters listed above are self-explanatory. An important one that is not, is the connection type. Network connection types are associated with the various kinds of network services. Those types relevant to this report are:

- Host-to-Host Interactive Access—IT
- Direct Terminal Access (Hermes)—TL
- External Access (Merit-Telenet)—X3
- Network Batch Access—EX and PR

The IT, TL, and X3 types represent interactive connections. They differ in that an IT connection is between a pair of serving hosts while the TL and X3 types are between a terminal and a host. The Hermes and Merit-Telenet internetwork service connections are further distinguishable by their initiating pseudohost[2] ID's. For batch service, the network differentiates between connections carrying jobs for remote execution, EX, and connections transmitting jobs for printing, PR.

One of the statistical programs run monthly separates the aggregate data records into two groups: records produced by a connection leaving a host, called Outbound, and records

[2] Each Hermes site and the Merit-Telenet interconnection are treated as hosts by the network. Hence, each has its own host ID just as the computing center hosts do. The modifier "pseudo" is used to identify noncomputing-center hosts.

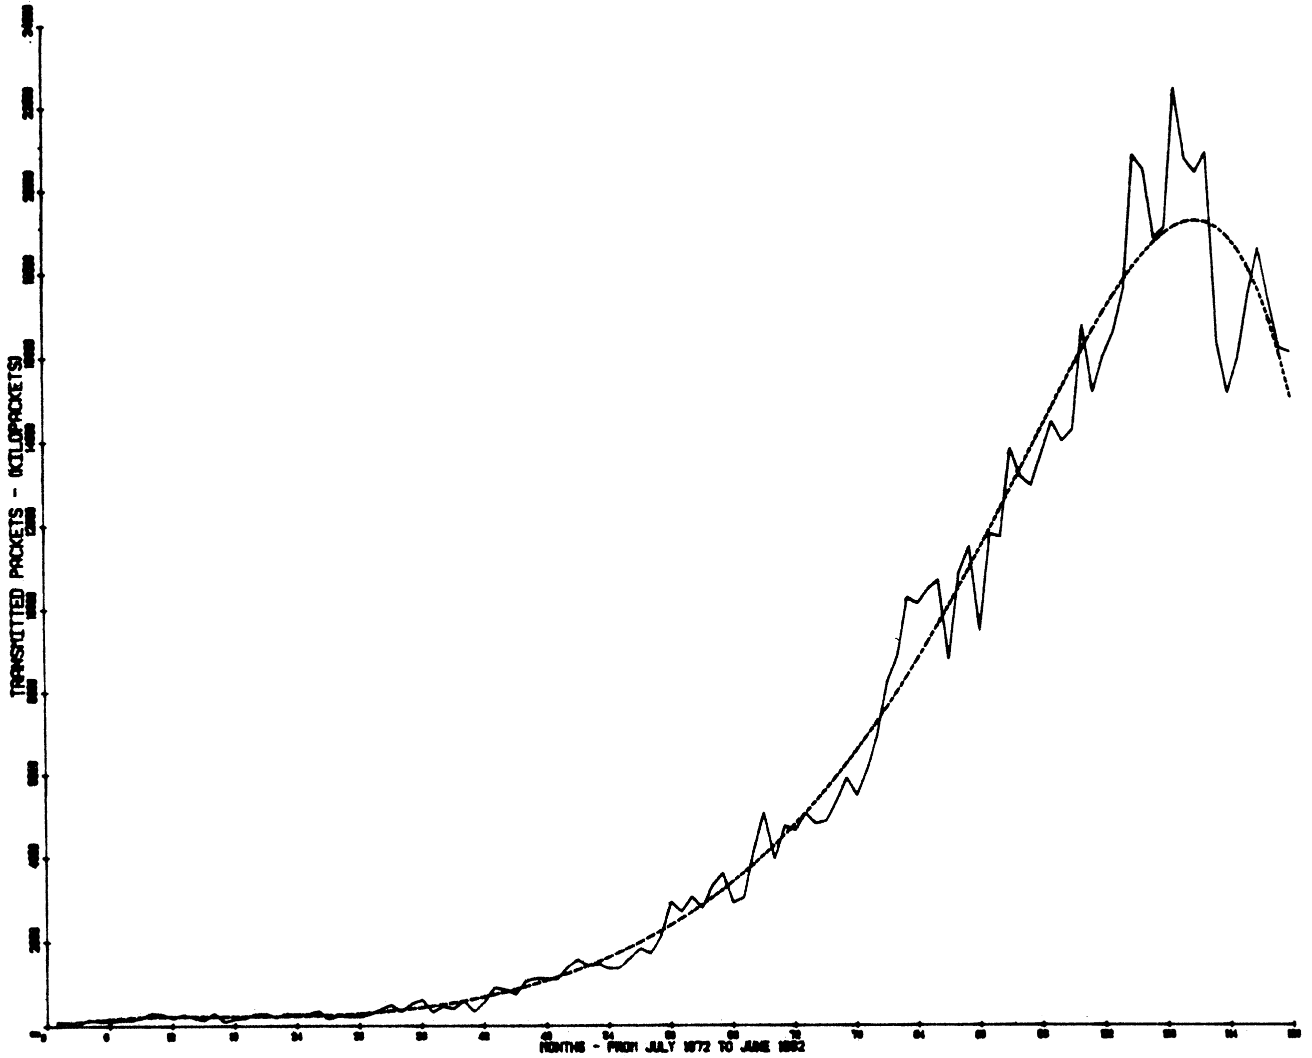

Fig. 6. Number of kilopackets transmitted.

produced by a connection coming into a host, the Inbound connections. These groups are processed separately but in essentially the same manner. The results produced for each group include:

- A listing of the individual user ID's.
- Statistical measures for connect time and transmitted bytes and packets for each connection type.
- An analysis of the distribution of connections between all pairs of network hosts.

For outbound connections, an analysis of the failure modes.

Only a small portion of the available information produced in these monthly summaries is presented in this paper.

IV. Growth of the Network

One purpose of the network's statistics is to record and report its use with time. This is done by maintaining a set of time-series, based on the monthly summaries, that describe total network traffic. Five of these time-series are as follows:

Number of Network Users—the number of different computing ID's successfully opening connections each month.

Successful Network Connections—the total number of successfully-opened connections each month.

Elapsed Connection Time—the total number of hours during which all successfully-opened connections remained connected during each month.

Transmitted Packets—the total number of packets transmitted over all successfully-opened connections each month, expressed in kilopackets. Each packet may contain up to 240 bytes of user data, but is no longer than necessary.

Transmitted Characters—the total number of characters, as eight-bit bytes, transmitted over all successfully-opened connections each month, expressed in megabytes.

Plots depicting each of these time-series are shown in Figs. 3–7. The horizontal axis in each plot represents a ten-year span beginning in July 1972. Grid or tick marks appear at six-month intervals. The scale on the vertical axis varies from one plot to another. In each plot the solid curve presents Inbound data while the dashed line represents a seventh-degree polynomial curve fitted through the respective data points.

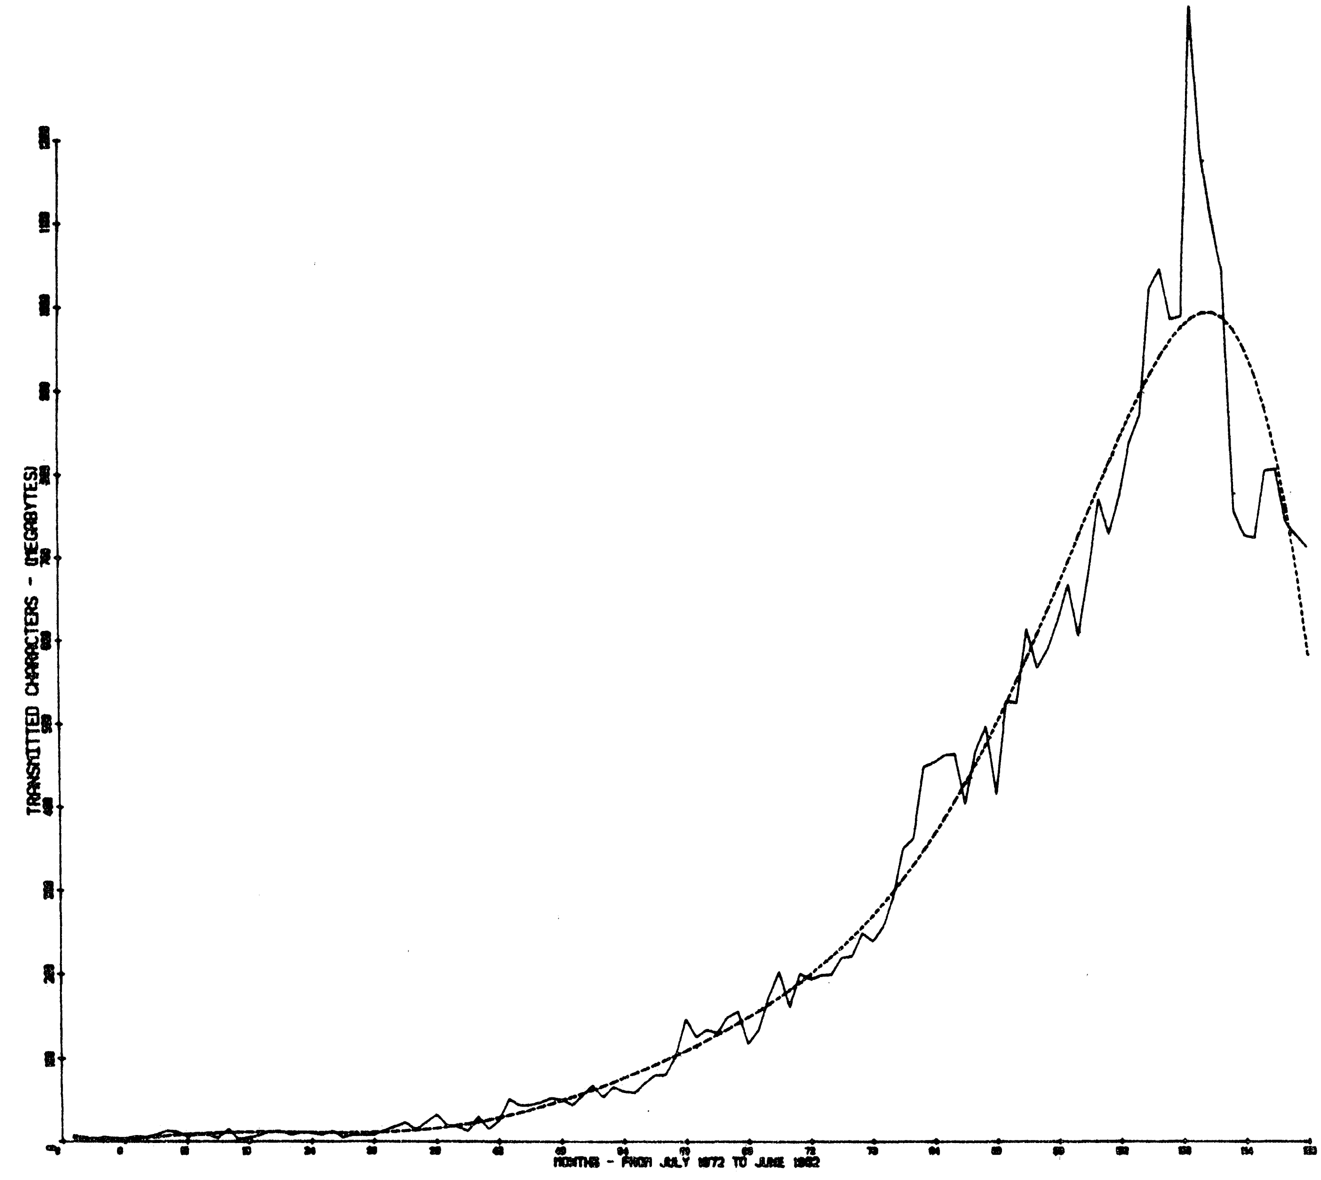

Fig. 7. Number of megabytes transmitted.

While these curves, with one notable exception, show progressively increasing usage, it must be recognized that the network has matured appreciably during this ten-year period. New services have been introduced and existing ones improved. Certainly the current network is a much more powerful and convenient one than the original 1972 network. It also involves more hosts and three more communications computers.

The previous point can be better understood by briefly tracing the network's development history; for more details see [2]. Merit first offered only host-to-host interactive service, IT connections, between two of its nodes in July 1972. This service was extended to include a third node in October 1972. In October 1973 the original 2000 bits/s internode telephone lines were replaced with 4800 bits/s service. Batch service first appeared in October 1974 between two nodes. It was not fully extended to the third node until January 1977. Hermes service began in a limited way at one node in September of 1975. This introduced the TL connection type. By May 1977, Hermes was available at all three of the original nodes. The Merit-Telenet internetwork tie started in December 1976 using links to Hermes ports. In early 1981 the X.25 service was introduced. The WMU node began to provide Hermes service in August 1979. Their host became available for interactive access in January 1981. The impact of this development sequence is reflected in data shown in the five plots.

The growth exception previously cited is the heavy data traffic in Summer 1981 as seen in Figs. 6 and 7. It was due to a large flow of print jobs from the UM to WSU to a Xerox 9700 laser printer. This exceptional traffic ceased when the UM Computing Center acquired its own Xerox 9700.

A second way to represent the network's growth is to show the same data used in the previous plots in a histogram format. Fig. 8 shows the monthly average number of users for each of the ten years measured from Inbound connections. Figs. 9-12 show yearly totals for each of the other network usage measures.

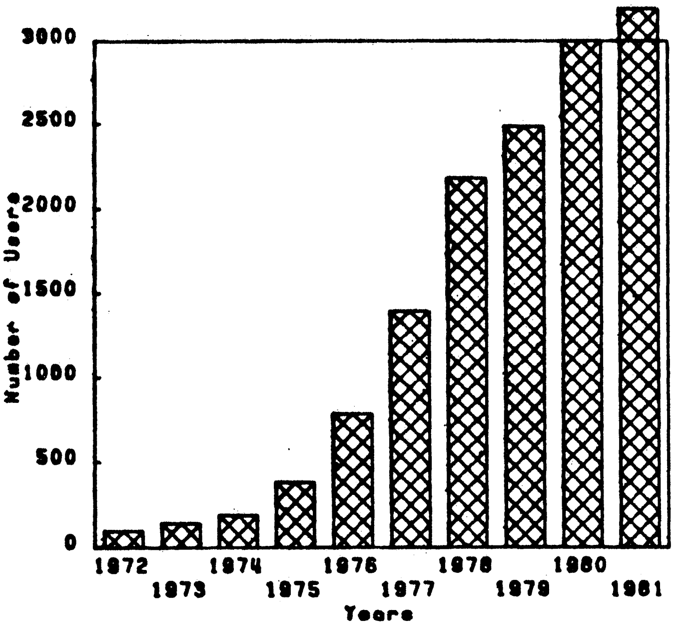

Fig. 8. Yearly averages of number of different network users per month.

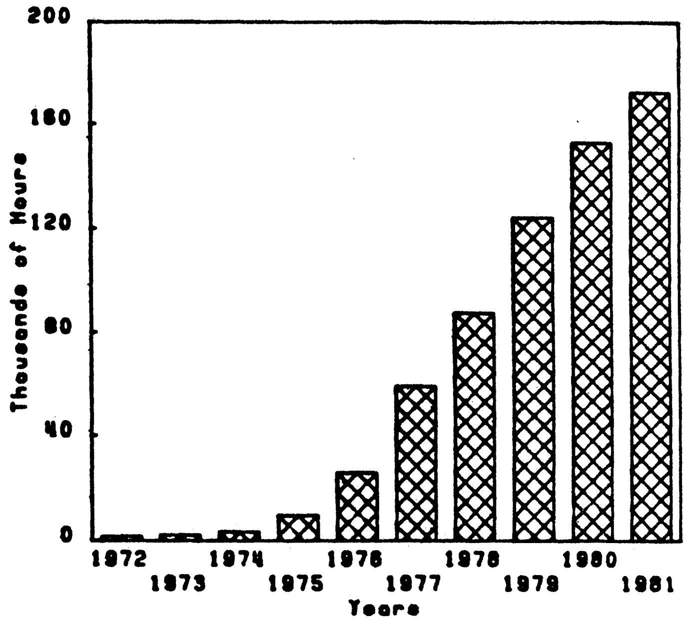

Fig. 10. Elapsed connect time per year (hours).

V. Normalized and Rate-Related Network Statistics

An alternative way to treat the data presented in the last section is to divide the three traffic measures—connect time, packets, and bytes—by the corresponding number of monthly connections. This normalization yields average values of these key network performance measures on a per-connection basis.

The next set of plots, Figs. 13-15, shows Merit's normalized Inbound time-series data for the ten-year period July 1972 to June 1982. Dashed seventh-degree polynomial curves again are drawn through the data. In contrast with Section IV, connect time is here expressed in minutes rather than hours. Similarly, the values for transmitted packets are given in units of single packets rather than kilopackets and those for transmitted characters in kilobytes instead of megabytes. Yearly averages for the same data are given as histograms in Figs. 16-18.

It is interesting to note the initially nearly constant but more recently increasing value for packets per connection, the steady and somewhat more rapidly increasing value for kilobytes per connection, and the apparently bimodal characteristics of time per connection. These changes are due to a pair of causes that affect these aggregate statistics in opposite ways. As is subsequently demonstrated in Section VI, the batch-type connections—EX and PR—are characterized by very short average connect times and very large values for packets and bytes

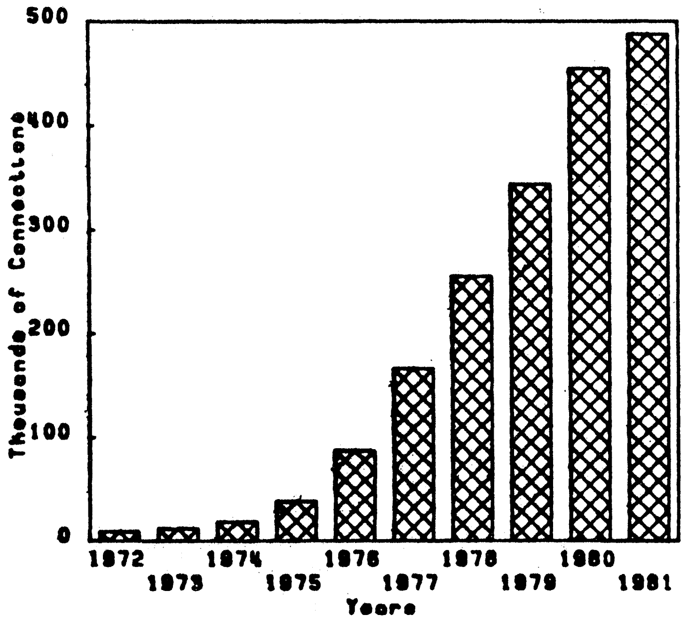

Fig. 9. Number of successful connections per year.

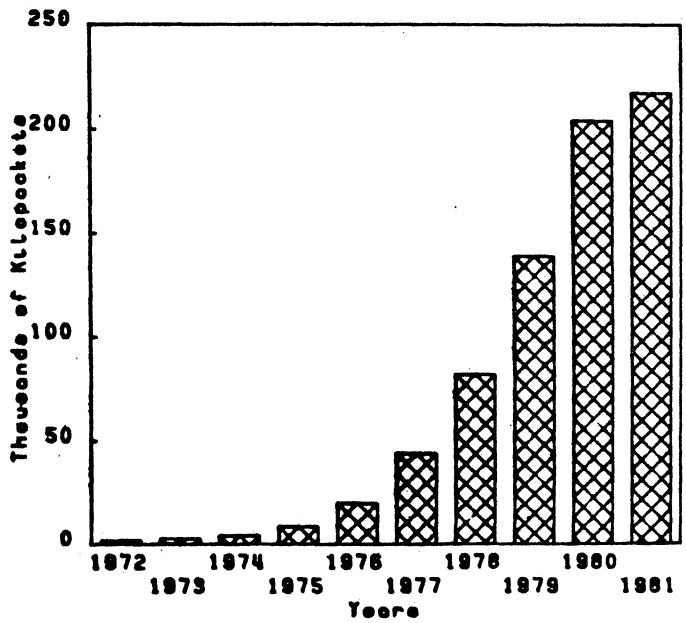

Fig. 11. Number of kilopackets transmitted per year.

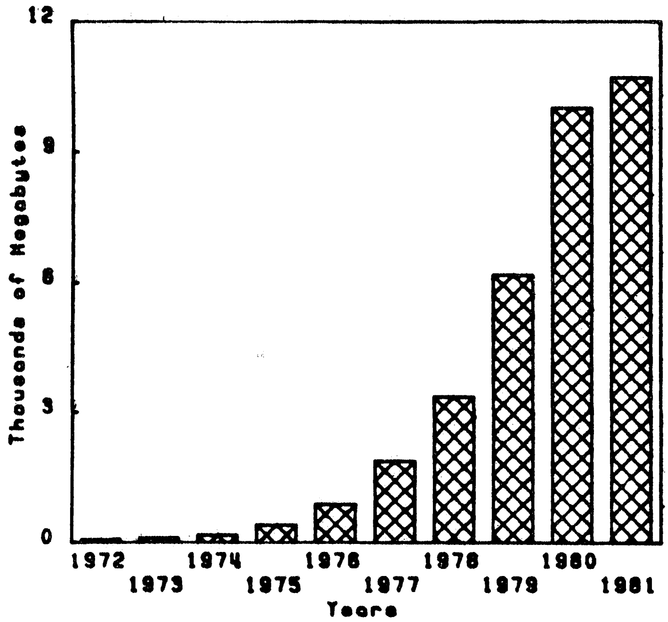

Fig. 12. Number of megabytes transmitted per year.

transmitted per connection. In contrast, the Hermes (TL) and Telenet (X3) type of connections exhibits just the inverse set of characteristics. There are now significantly more of these direct, terminal-initiated connections than any other type. This explains the increase in the average connect time starting in the latter half of 1975 and apparently ending in 1977 as the network's mix of connect type usage began to stabilize. Although the number of batch connections is fewer, the great amount of data that they carry explains the rise in the other two measures. The observed effect of the relatively heavy print traffic in the Summer 1981 is consistent with this explanation.

As with the normalized statistics, it is possible to compute other ratios from the basic Inbound data. For example, the two rate parameters—bytes per second and packets per hour—and the number of bytes per packet are useful to network designers. The time-series plots and linear regression curves for these three parameters appear in Figs. 19–21. Yearly averages for them are shown as histograms in Figs. 22–24.

The large spike which appears in Figs. 19 and 20 for October 1973 is attributable to extensive testing of the network's transmission capabilities by its staff. While this data point is atypical, it also demonstrates that much higher throughput rates are possible than those which occur naturally in the network. The generally declining average values for bytes per second and packets per hour until mid-1977 reflect the influence of Hermes connections, with their lower rate of data transmission. The rise in these parameters since then is due to three factors. One is the increasing use of the network's batch service to transfer large data files, a second is the fact that more interactive text-processing and computer conferencing are being done, and a third is that more 1200 bits/s terminals now use the network.

A remarkable fact is that the number of bytes per packet has remained almost constant over this ten-year period. No good explanation has been proposed for this consistency over Merit's relatively varied operational circumstances. This fact should interest network designers, some of whom have been known to spend hours debating what a network's maximum packet length should be.

VI. Node- and Type-Specific Data

In the two previous sections data were presented in an aggregate form, combining all nodes and connection types. In this section, the emphasis shifts to the distribution of network activity among the nodes and to the differences in the various measures when analyzed by their respective connection types. Only data from our last fiscal year, July 1981 through June 1982, are used here: these data most accurately portray the network's recent utilization and best represent its diversity of services. Only averages based on these twelve months of Inbound data appear; the monthly fluctuations are not shown.

The distribution of usage varies considerably among the hosts. This may be seen by examining the bar graph in Fig. 25. It shows the relative number of connections, elapsed hours of connection time, kilopackets, and megabytes per host. The relative newness and still limited network access of the Western Michigan host is reflected by their comparatively low usage. In contrast, over half of the network's traffic flows to Wayne State University's two hosts.

Like the distribution among the hosts, the distribution of usage among the five connection types is also very unbalanced. Fig. 26 shows that terminal connections dominate all others. This bar chart clearly demonstrates that conventional terminal access to hosts is by far the most-used Merit network service; it even accounts for about 40 percent of the total data transmitted.

The next topic addressed is how connect time, packets, and bytes per connection vary with connection type. One expects connect time to be greater for interactive than for batch connections. But how much greater? How will packets and bytes vary among the several connection types? These questions are answered in the next series of figures.

Fig. 27 shows the variation of connect time per connection as a function of the five connection types. In this figure a solid line joins the average values but has no other significance. As expected, the host-to-host interactive, terminal, and external connections have much greater average connect times than do the batch services. In fact, the average connect time of a print connection is only 2.6 minutes compared with more than 29 minutes for a terminal connection. The difference between terminal and external connect time is interesting to note. It may be due to the higher cost of using external access.

The shaded box surrounding each average conveys two additional pieces of information. The vertical extremes correspond to plus and minus one standard deviation from the average value. The width of each box is proportional to the number of connections for its connection type.

The next five figures (Figs. 28–32) also use the graphical format described just above. They show respectively the packets and kilobytes per connection and the number of bytes per packet, bytes per second, and kilopackets per hour.

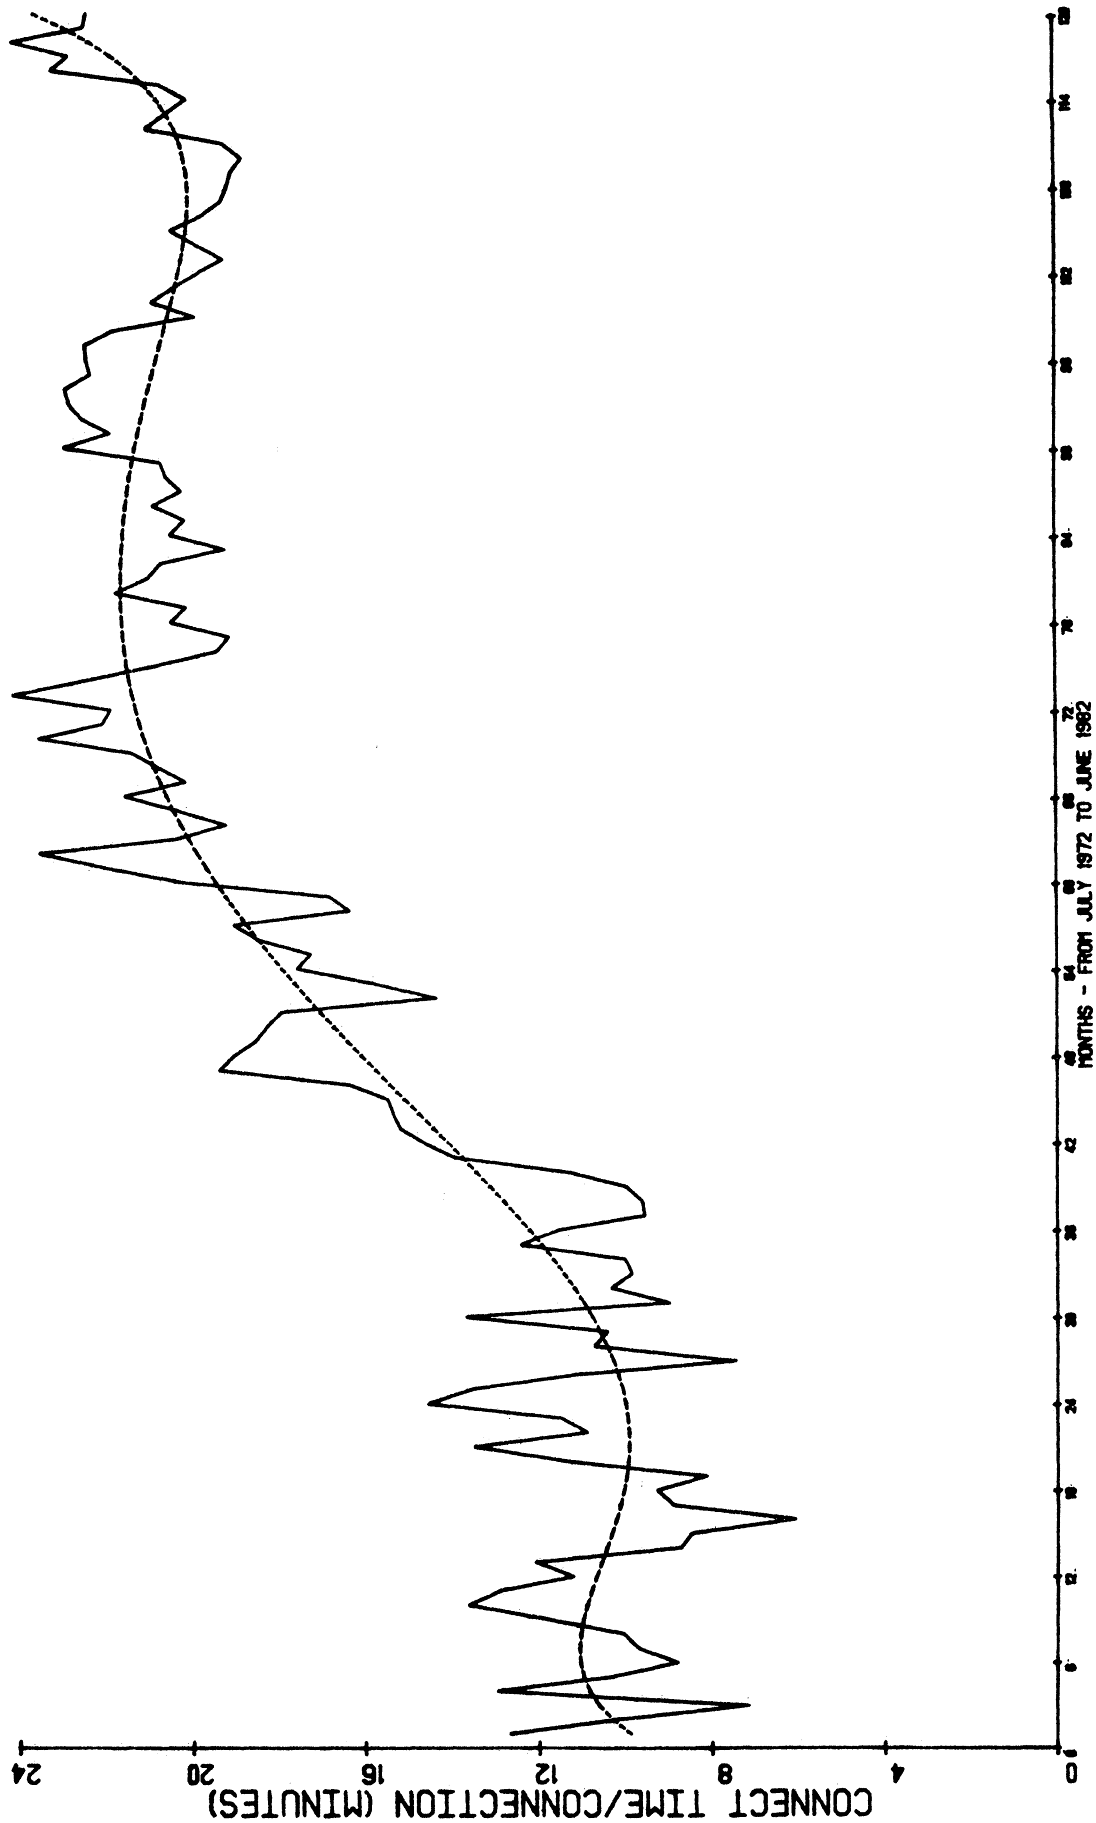

Fig. 13. Average connect time per connection (minutes).

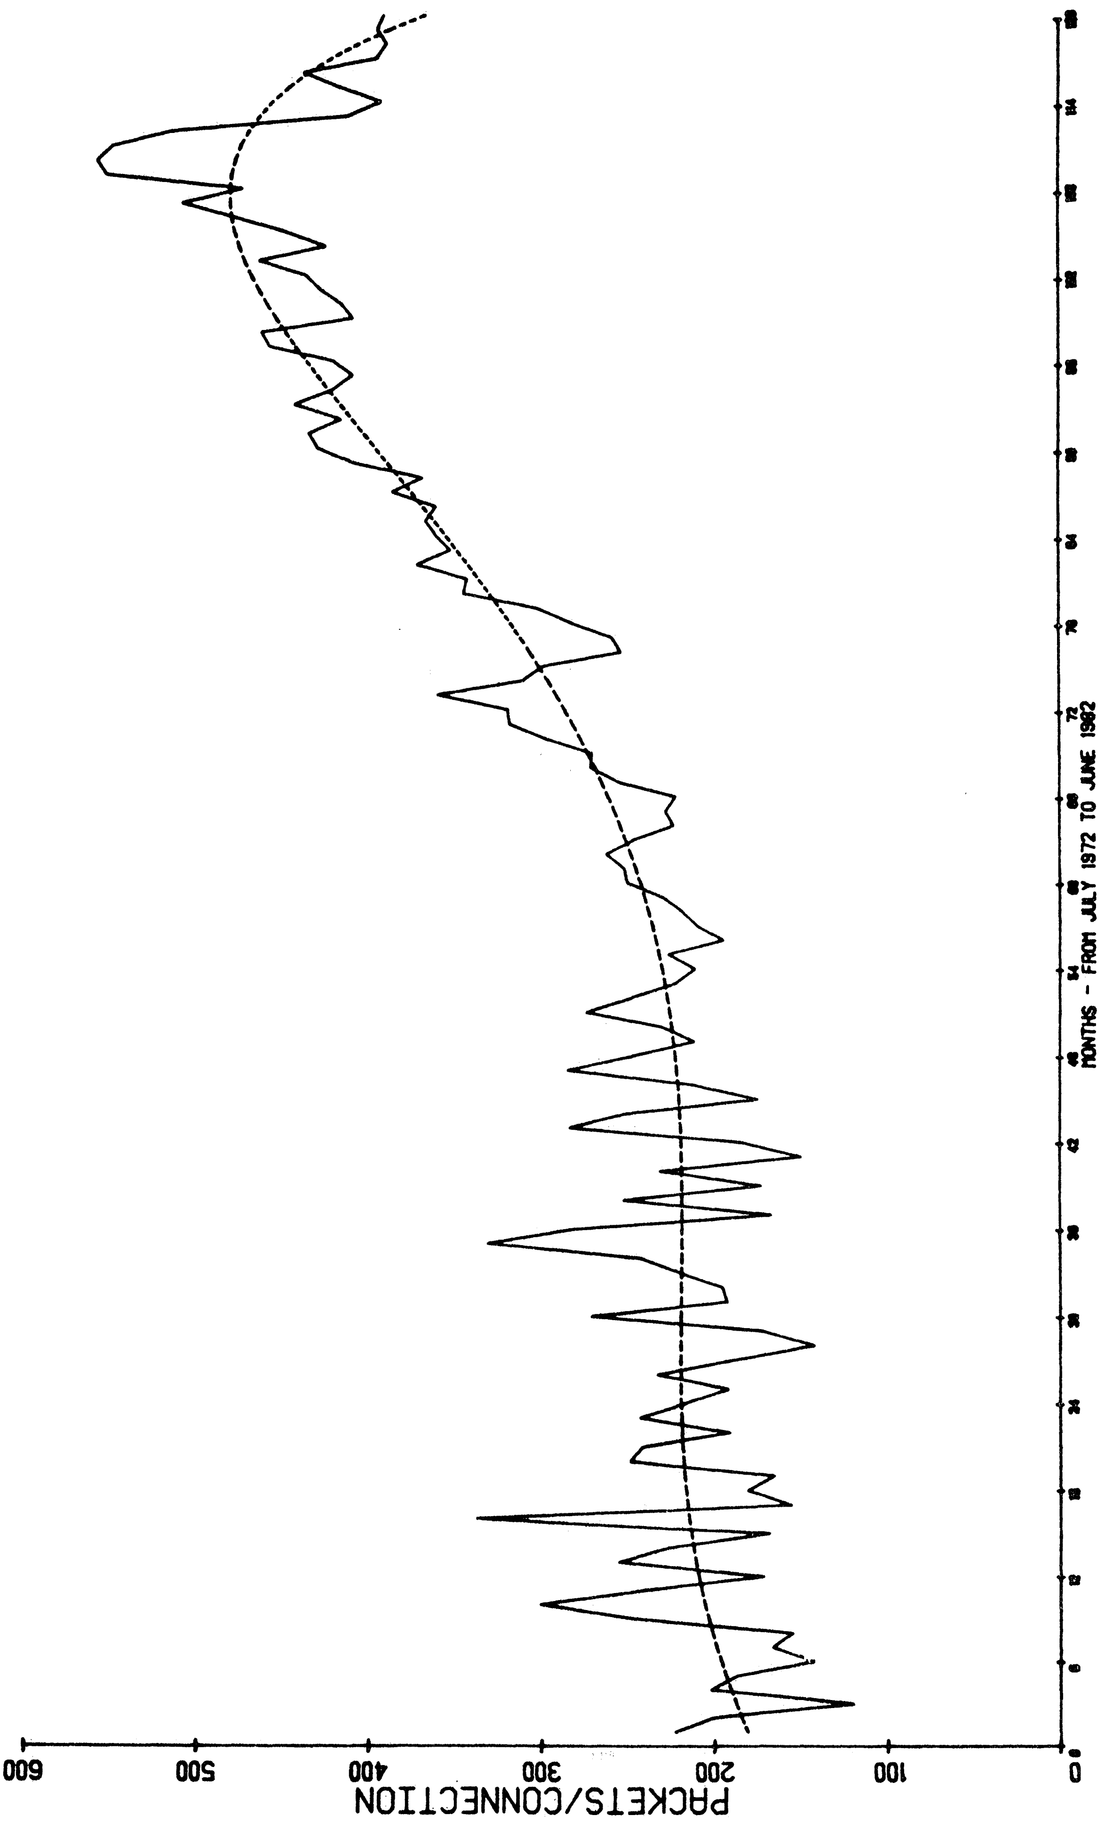

Fig. 14. Average number of packets per connection.

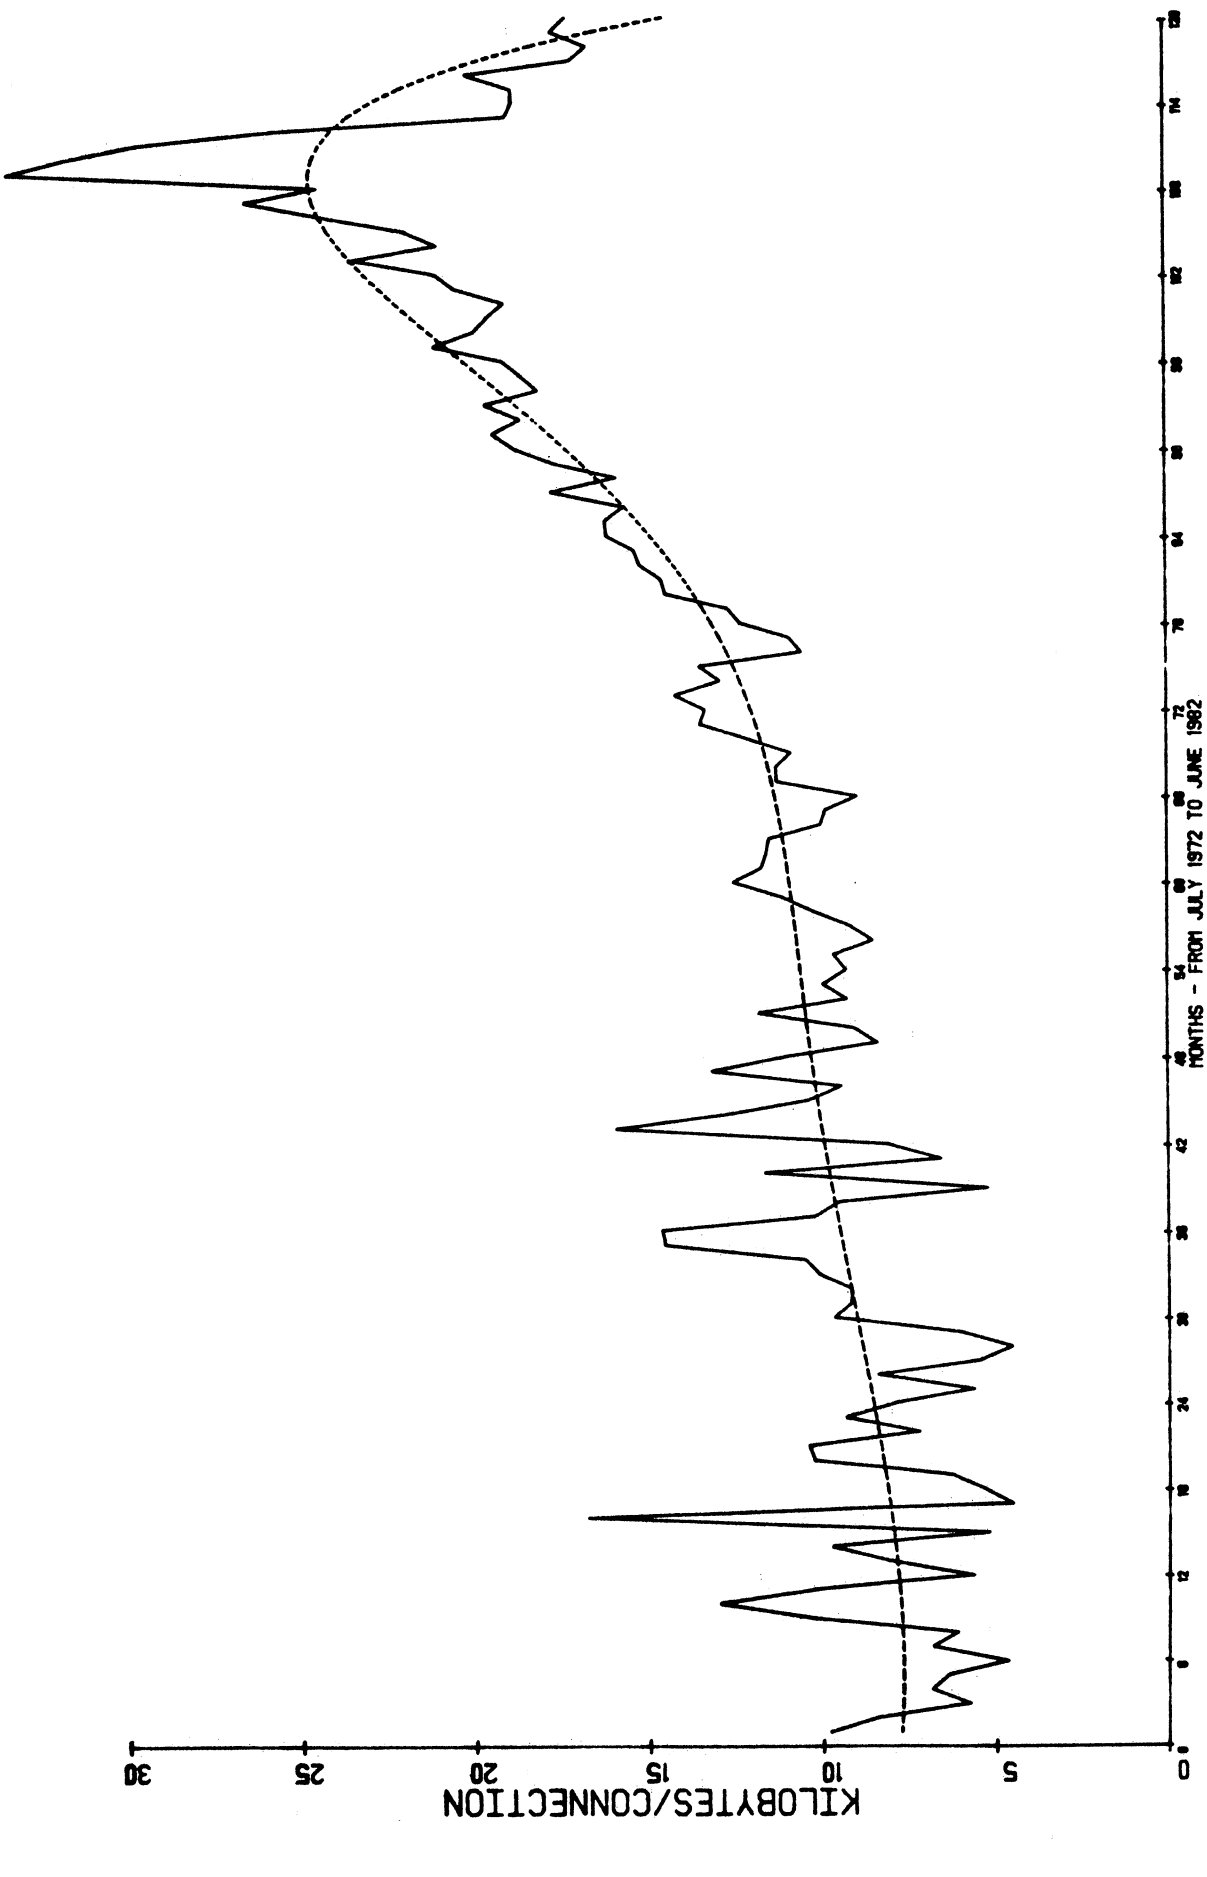

Fig. 15. Average number of kilobytes per connection.

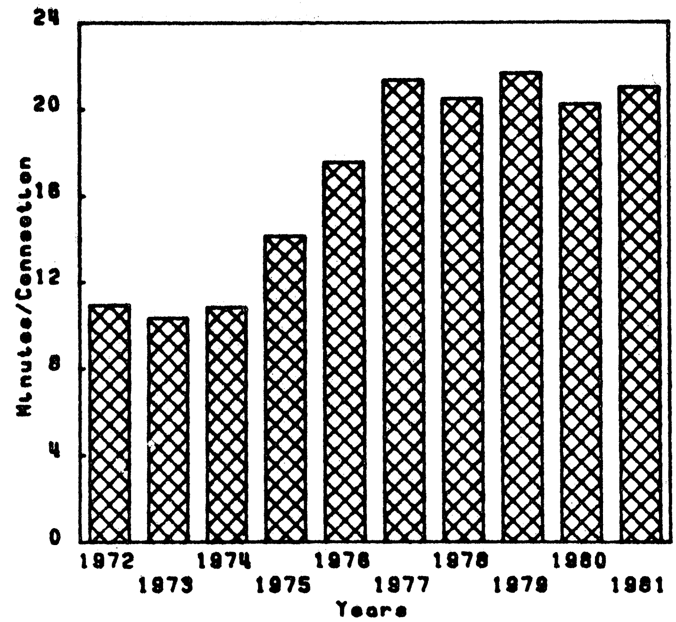

Fig. 16. Yearly average of connect time per connection (minutes).

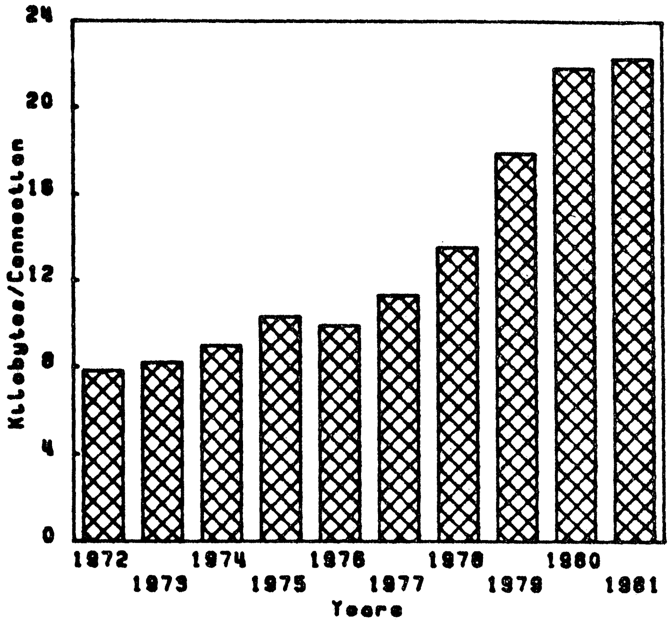

Fig. 18. Yearly average of kilobytes per connection.

That significantly more data is carried by individual batch connections, as revealed in Figs. 28 and 29, is readily explainable. Since the network makes it convenient to route output to a remote printer, most users prefer to use this service for large print files rather than copying such output to their terminals. Similarly, for reasons of both economy and convenience users use the Execute connection type to copy large files between hosts. This latter service has been made even more cost-effective and easy with the introduction of a special batch file-copy utility program which packs data into maximum (240-byte) packets. This utility program accounts for the increase in the kilobytes per connection and bytes per packet

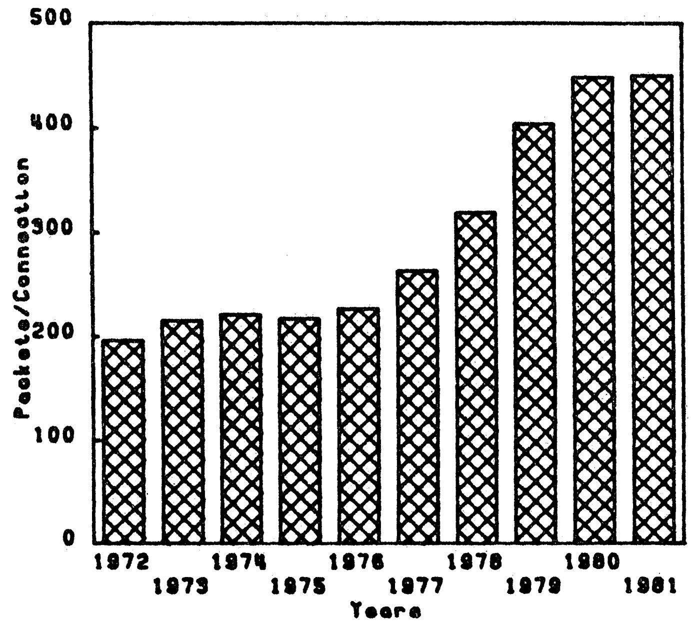

Fig. 17. Yearly average of packets per connection.

ratios for Execute connections compared with those observed in earlier years.

The batch transmission rates shown in Figs. 31 and 32 are limited primarily by the data rates on the network links. While interactive and batch connections have equal priority within the network, the hosts limit the number of concurrent connections of each class. Typically each host will accept only one batch connection from each other host but will accept many concurrent interactive connections. This form of regulation has allowed batch traffic to move at maximum data rates without resulting in any network-caused response delays for interactive traffic. The significant difference between the interactive and batch average data rates makes this mix possible.

VII. Individual Parameter Variations

This section lists the distribution of several network measures based on the individual monthly statistical records rather than on the monthly aggregates. These data best describe the variation among the network's connection types. The following tables contain Inbound data for one month. The first or All row lists the parameter values for all of the individual connection types treated as a group. Since some minor connection types are not described in this paper, the All totals slightly exceed the sum of the individually-listed connection types.

VIII. Concluding Remarks

What are the purposes and reasons for networking? There are many, but the four most commonly given are: remote access, load sharing, resource sharing, and process sharing. Each of these offers either economic or other advantages.

The first of these, remote access, means utilizing distant computing services through some form of communication system. Most commercial computer timesharing companies

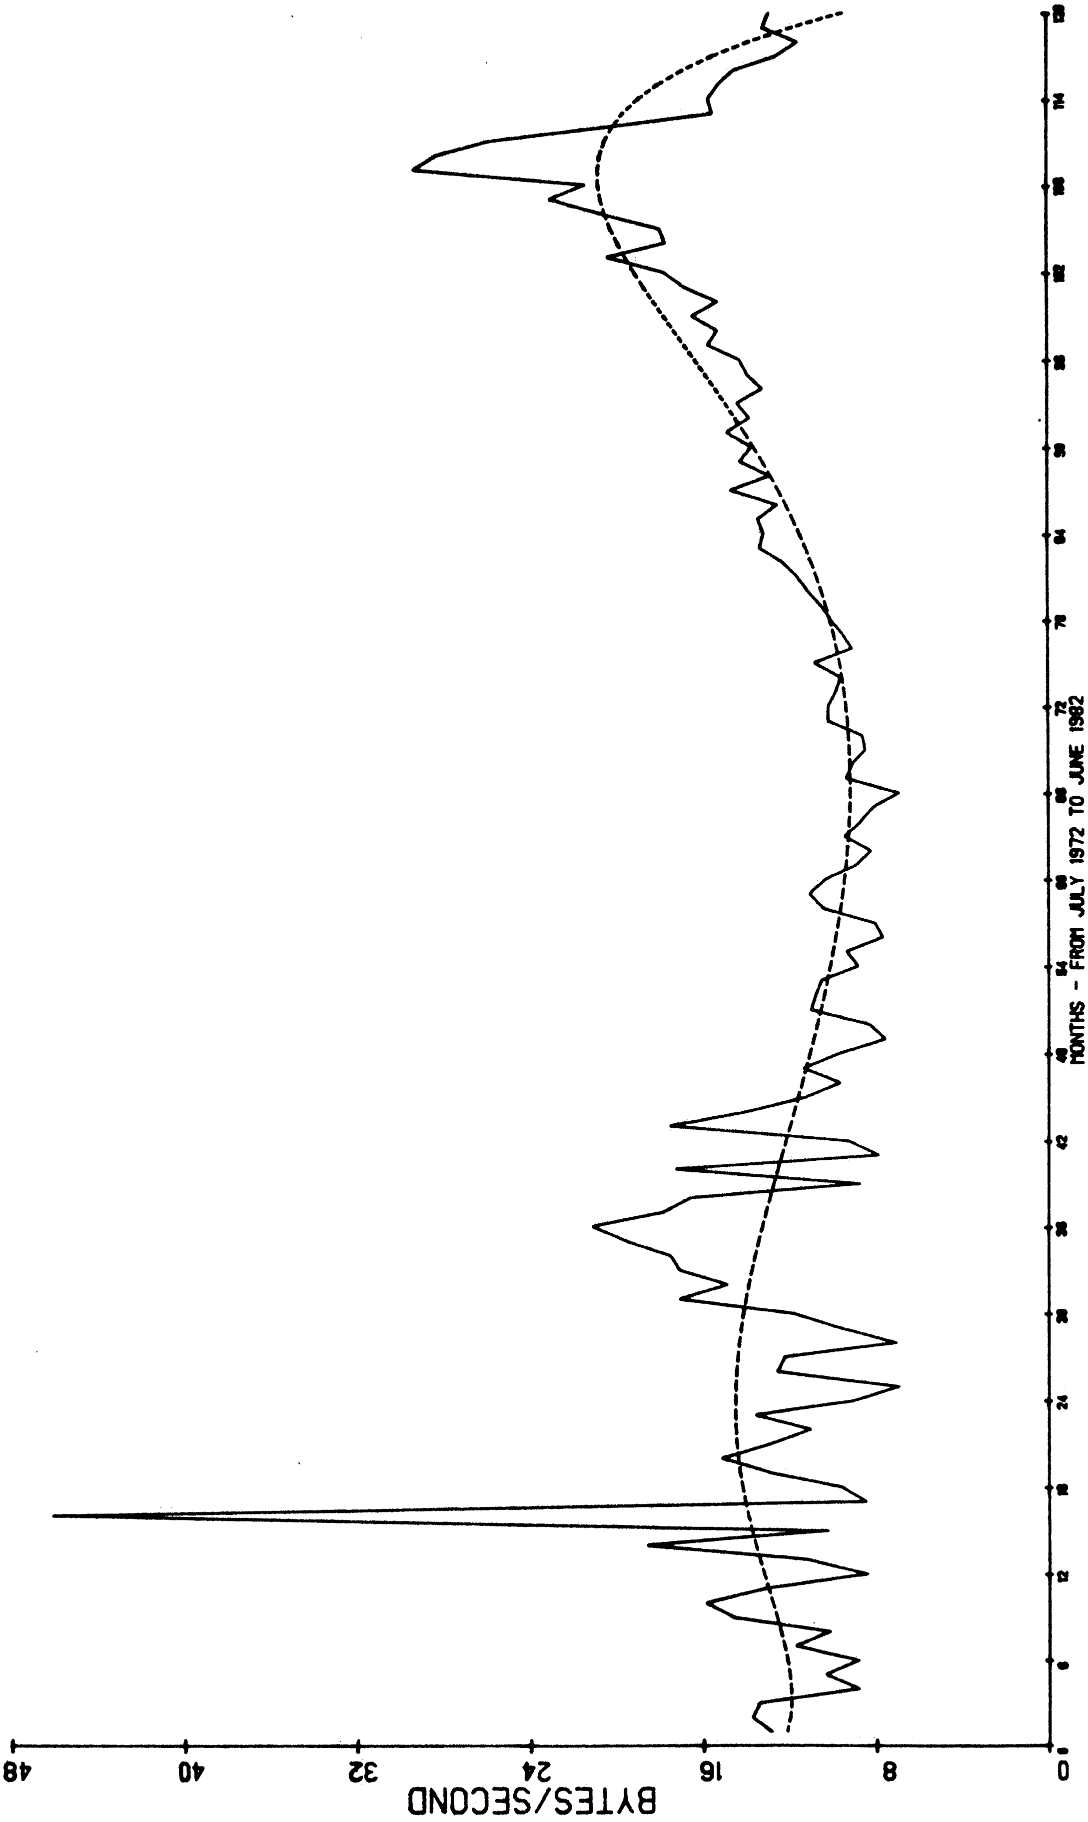

Fig. 19. Average number of bytes per second.

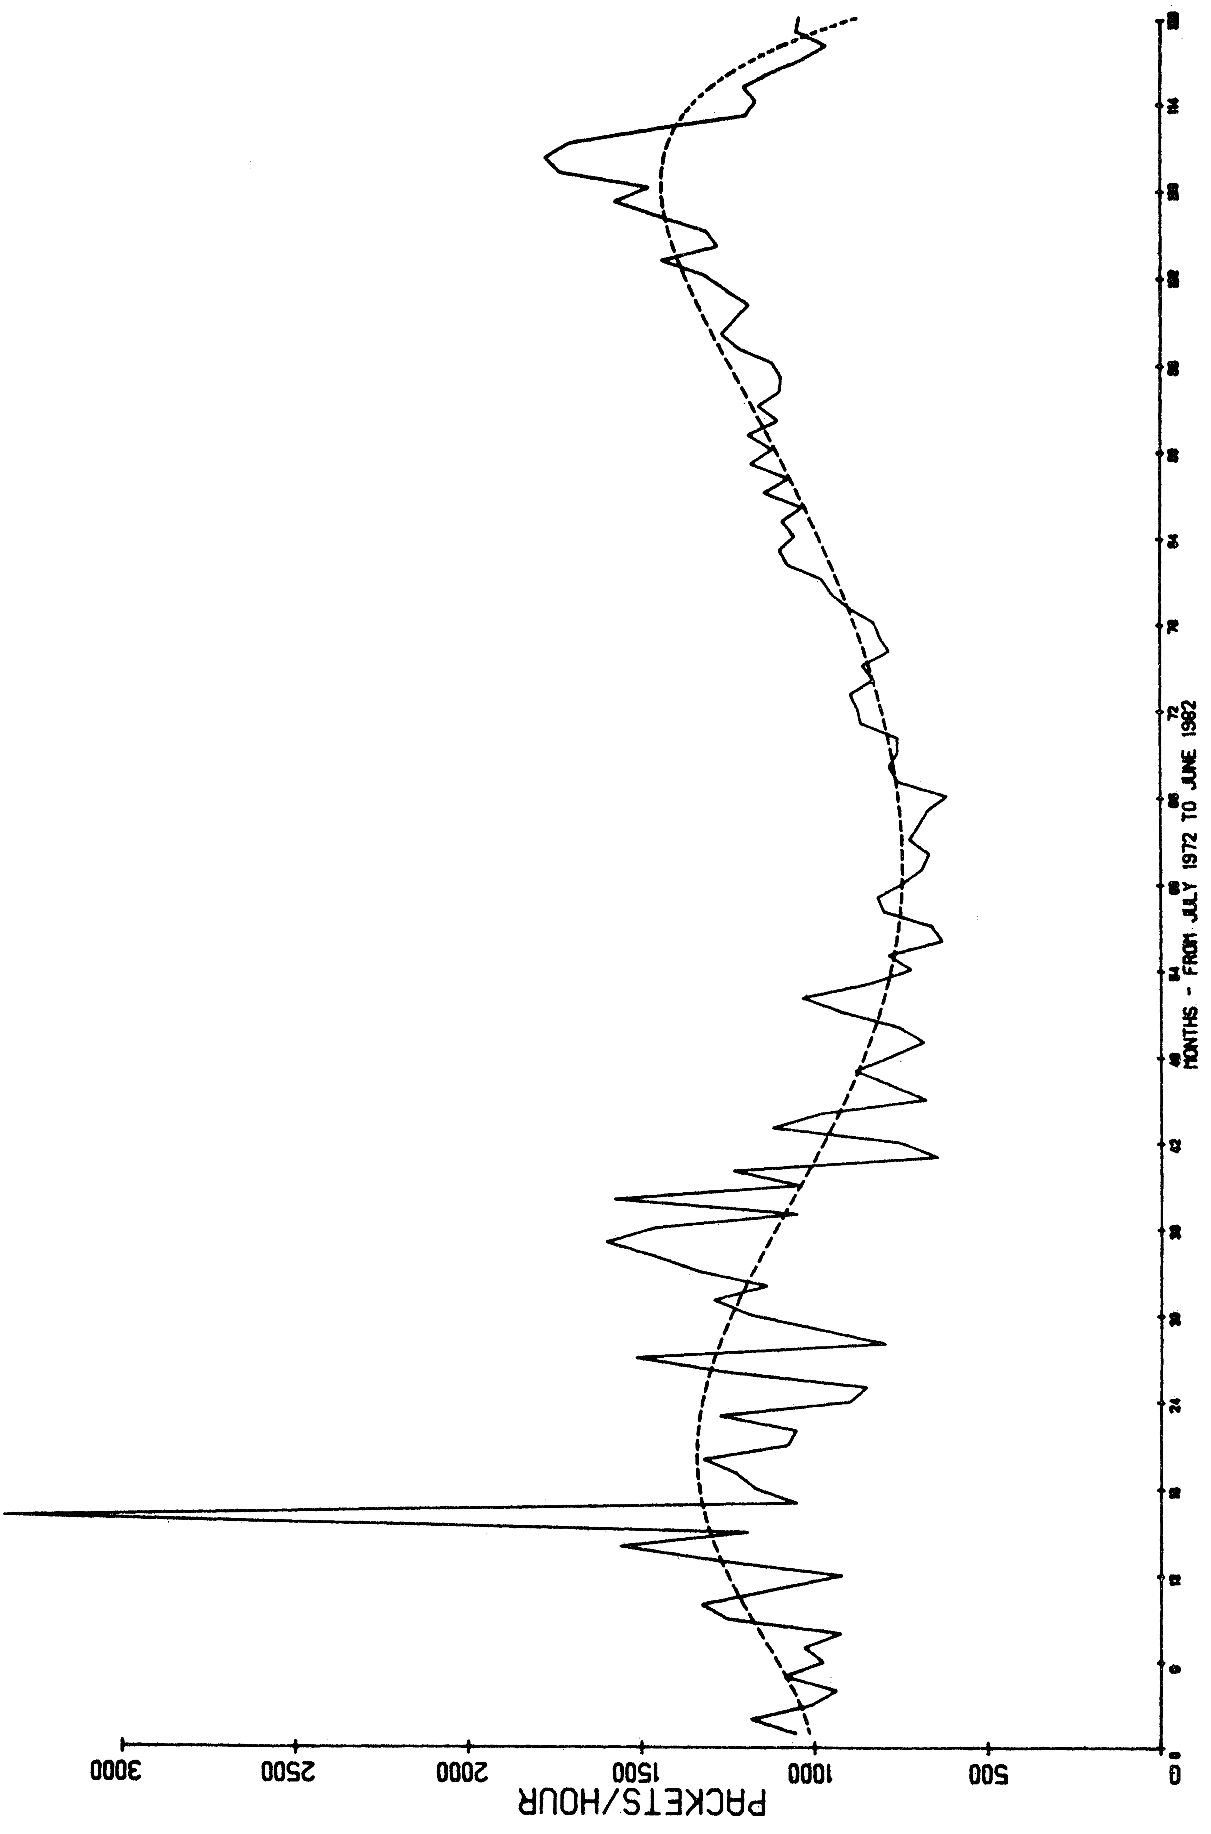

Fig. 20. Average number of packets per hour.

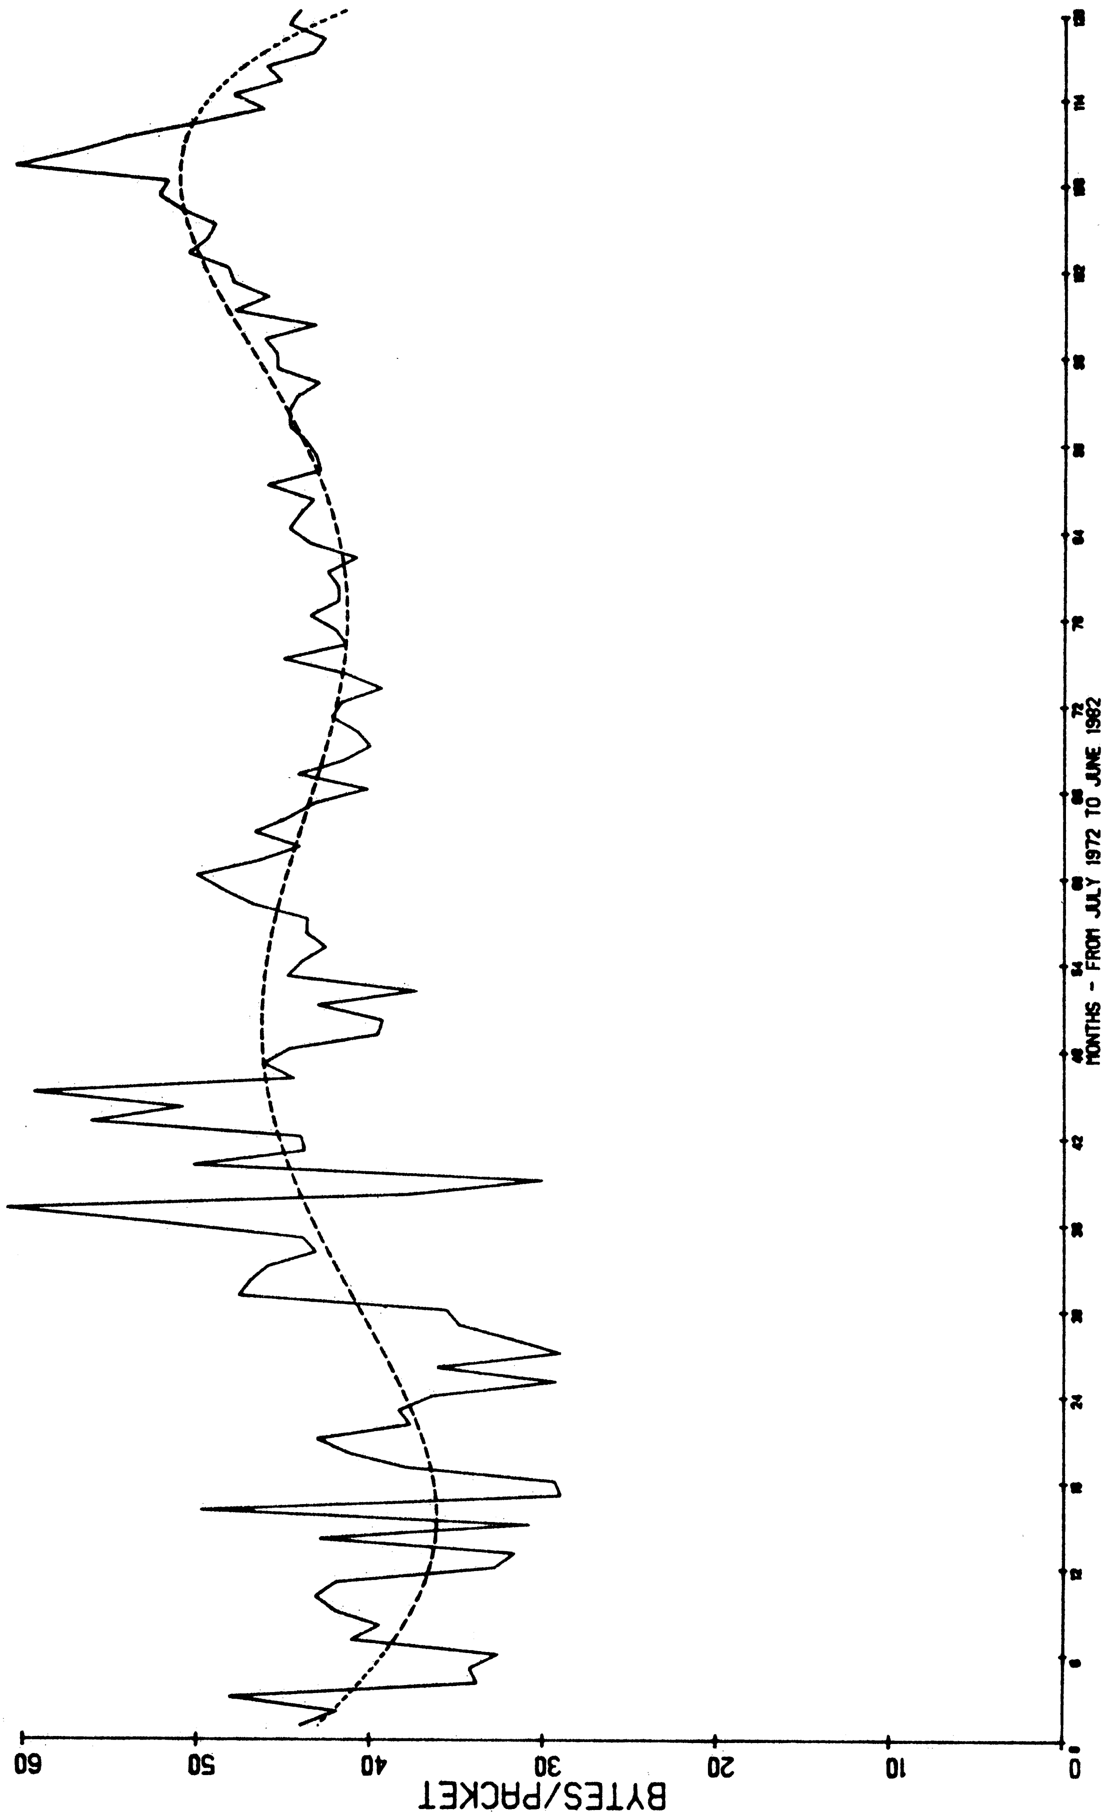

Fig. 21. Average number of bytes per packet.

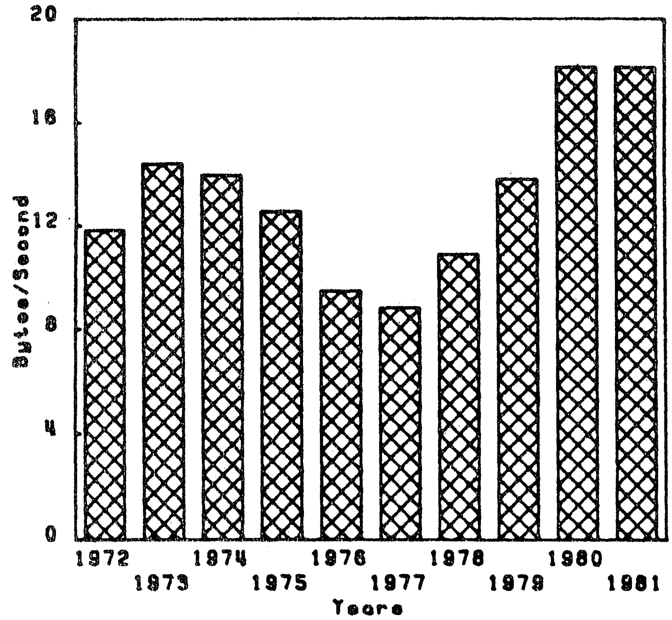

Fig. 22. Yearly average number of bytes per second.

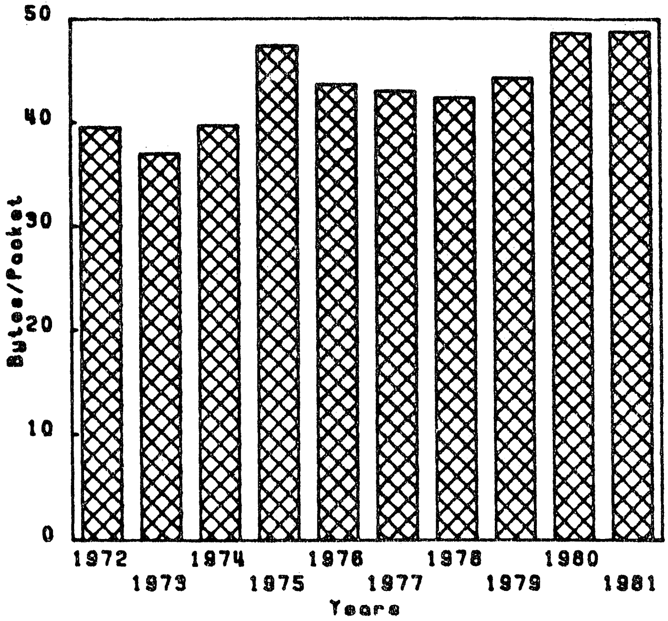

Fig. 24. Yearly average number of bytes per packet.

offer remote access to their clients through such a network. This allows them to operate a single central facility and yet serve clients throughout the country or even world. These companies find remote access to be a cost-effective alternative when compared with maintaining a distributed set of computing facilities.

Load sharing implies distributing a given load among the several computers on a network. This seemingly obvious and desirable objective is surprisingly difficult to implement. While effective load sharing is more easily accomplished when all of a network's computers use common hardware and software, it was rarely attempted until the last year or two with some LAN networks. Load sharing has never been successfully employed in heterogeneous networks.

Resource sharing means making available to a dispersed base of users all or many of the network's collective resources. Resource-sharing networks typically allow users access to unique computer hardware, data bases, programs, and programming languages that are not provided by the user's local system. Merit is an example of a resource-sharing network, and there are many others.

Process sharing combines and extends the features of load- and resource-sharing. In a process-sharing network all of its resources are automatically managed on a global scale so that data files and processes (programs) can migrate throughout the network in response to user requirements and the current system loading. Individual users may not necessarily care about or be aware of the physical location of their resources in this type of network. Some practical examples of local process-sharing networks are now being introduced.

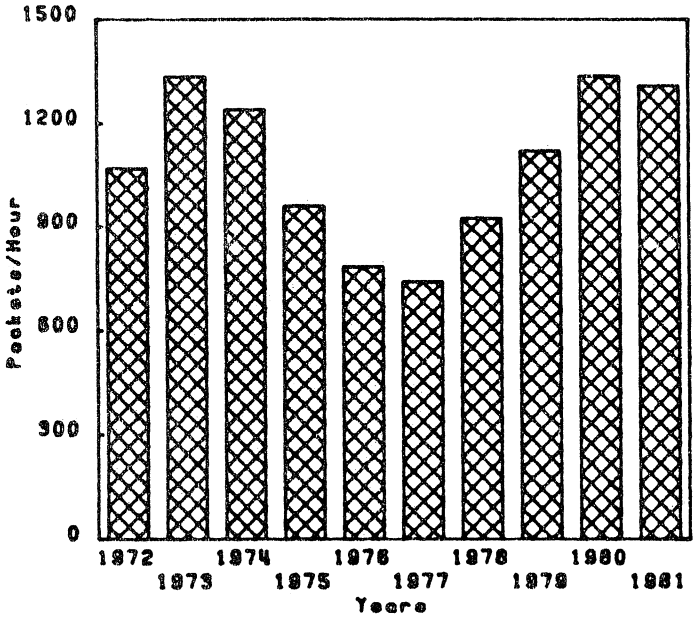

Fig. 23. Yearly average number of packets per hour.

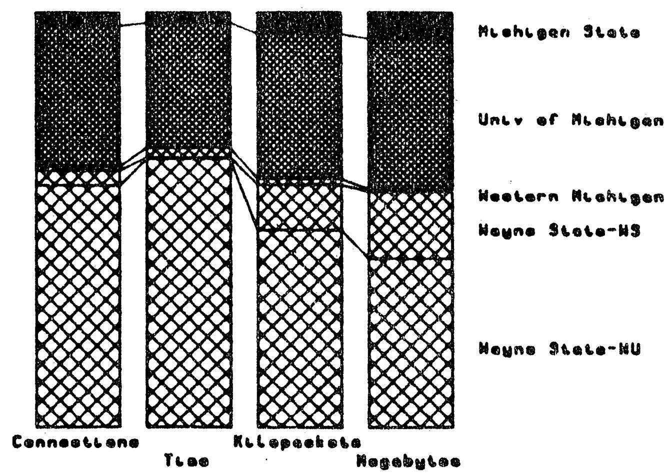

Fig. 25. Relative traffic to each host.

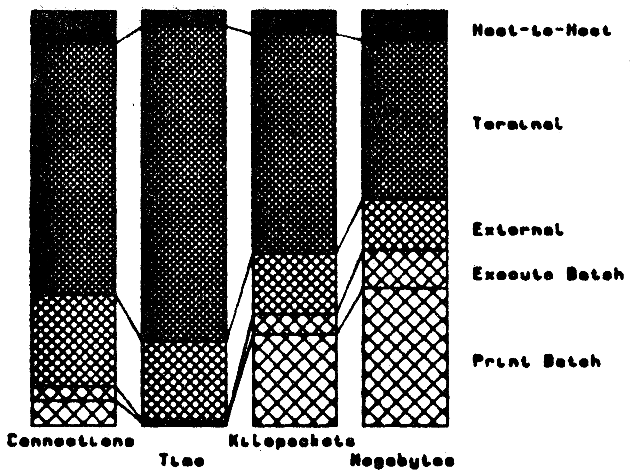

Fig. 26. Relative connection type distribution.

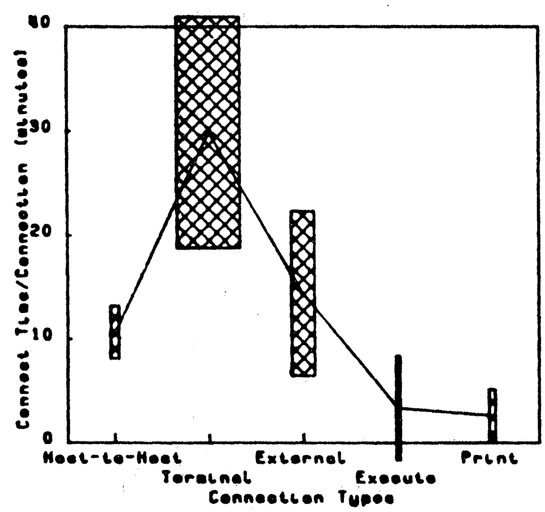

Fig. 27. Connect time per connection by connection type.

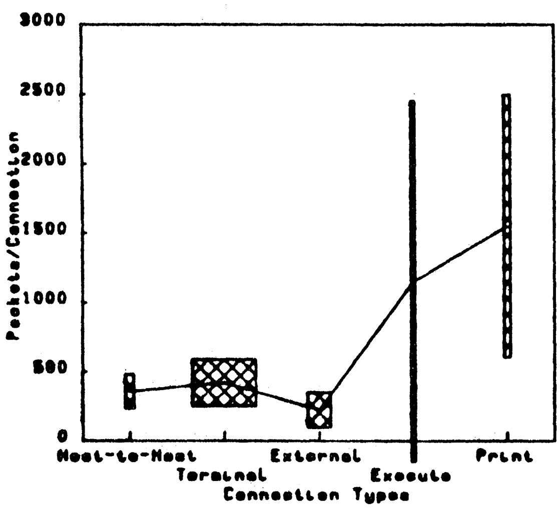

Fig. 28. Packets per connection by connection type.

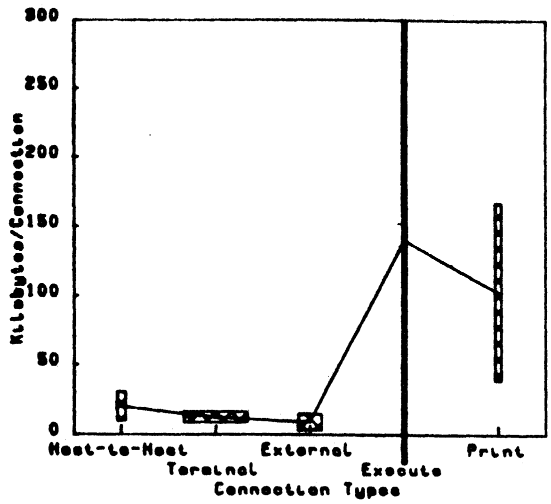

Fig. 29. Kilobytes by connection per connection type.

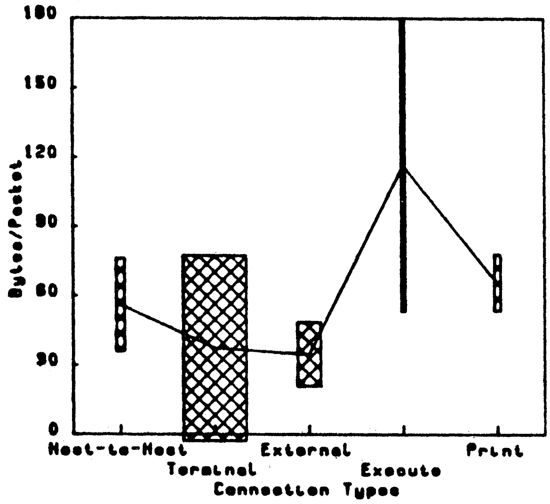

Fig. 30. Bytes per packet by connection type.

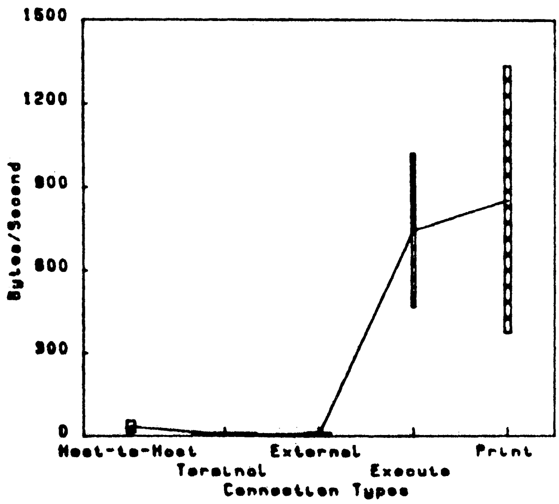

Fig. 31. Bytes per second by connection type.

A common thread in the purposes just listed is the concept of sharing. In most networks, resources available on any computing systems are available for use at other computers in the network. The following kinds of resource sharing are usually involved:

1) Device Sharing: The ability to connect to and use the resources of a remotely located computing device as if it were local; for example, to produce a listing on a remote line printer.

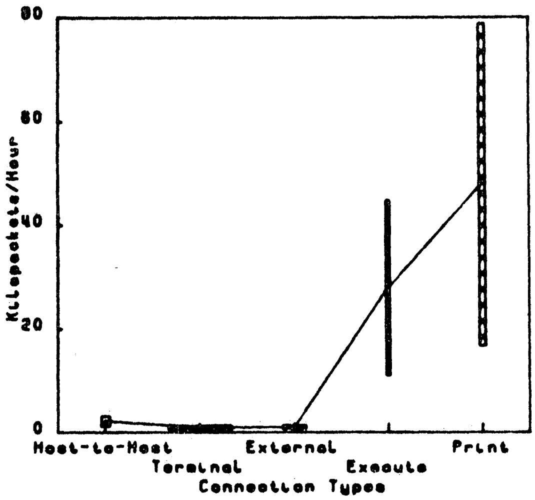

Fig. 32. Kilopackets per hour by connection type.

2) *File Sharing:* The ability to read from, write to, or update files on a remotely located computing system as if they were local.

3) *Program Sharing:* The ability either to send a loadable program to a remotely located computing system to be loaded and executed by that system, or to execute a remotely located program with local data.

4) *Interprocess Sharing:* The provision of data paths between programs running on separate hosts, so that large tasks may be divided into smaller units for execution at different computing sites in the network.

Merit is primarily a resource-sharing network; it is not, nor for a long time is it likely to become, either a load- or process-sharing network. In varying degrees Merit has implemented all four of the sharing modes described above. But the onus is still on users to know where resources are located, how they are used, and how these resources are accessed through the network. A user cannot enter the network and simply request "run my Pascal program" with the expectation that the network will handle all of the necessary details. Instead, users must specify which node they wish to access, sign on or log in to that host, and issue the necessary commands, in the selected host's own command language, to achieve their computing task. This differs from IBM's SNA approach to networking, for example.

While users have the responsibility for managing their own resource-sharing, they are by the same mechanism afforded the opportunity to load-share. They may, from time to time, direct some of their work to less busy hosts or they may regularly use a remote host's resources for load- rather than resource-sharing reasons. Such user-scheduled load sharing does not account for much network traffic.

The purpose in further elaborating on the typical Merit user's viewpoint is to provide a better basis for assessing the statistics contained herein. Certainly the technology, functional capabilities, and user environment of existing networks vary greatly. These factors must influence the statistical measures observed. But it is nearly impossible to comment intelligently on how typical or atypical Merit's data are relative to the experiences of others since little similar information has been published. Perhaps this paper will stimulate others to provide corresponding usage data about their networks.

TABLE I
CONNECT TIME PER CONNECTION (HOURS)

Connection Type	Number of Connections	Mean Value	Standard Deviation	Maximum Value
All	47 218	0.379	0.733	16.624
IT	3607	0.163	0.372	12.910
TL	30 291	0.463	0.819	16.624
X3	9642	0.331	0.608	12.900
EX	1597	0.017	0.103	2.450
PR	1766	0.032	0.265	10.847

TABLE II
PACKETS TRANSMITTED PER CONNECTION

Connection Type	Number of Connections	Mean Value	Standard Deviation	Maximum Value
All	47 218	394	1163	55 845
IT	3607	293	1364	55 845
TL	30 291	386	914	34 997
X3	9642	315	815	22 581
EX	1597	453	2108	48 112
PR	1766	1173	3111	49 535

TABLE III
KILOBYTES TRANSMITTED PER CONNECTION

Connection Type	Number of Connections	Mean Value	Standard Deviation	Maximum Value
All	47 218	17	81	6132
IT	3607	15	64	1709
TL	30 291	14	35	2091
X3	9642	13	37	1115
EX	1597	39	262	6132
PR	1766	85	270	4211

Throughout this paper comments amplifying the data have been made where it seemed appropriate to do so. A few additional, more general, ones, have been left for this section. The first of these addresses how individuals make use of the network.

Without question the most helpful network services are the Hermes (direct terminal-to-host) and external (X.25 link with Telenet) services. The introduction of these services, more than any other single factor, has boosted the network's level of usage. The obvious inference is that most of the networking needs of our users are very simple. They want to access a remotely-located resource interactively. The network provides them with an alternative to placing a long-distance telephone call, which, perhaps, they would not make. A further indication of the importance of these services is given in Table IV. This table lists the hours of Hermes and Telenet use per month during the July 1, 1981 to June 30, 1982 fiscal year. The times reported are total usage hours, including access through Hermes to a node's local host.

While the host-to-host interactive connections are far less frequently used, they provide an important service to both novice and advanced network users. Novices will usually use this service to transfer a data file between hosts. For some, a

TABLE IV
FISCAL 1981 HERMES USAGE IN HOURS[3]

Month	From Ann Arbor	From Detroit	From E. Lansing	From Kalamazoo	From Telenet	Total
July	2069	5503	1267	396	2726	11 961
Aug.	1811	4961	1144	334	2579	10 829
Sept.	1999	5198	1089	294	2562	11 142
Oct.	2544	6844	1304	226	2422	13 340
Nov.	2202	6808	1244	175	2507	12 936
Dec.	2328	6185	1240	192	2278	12 223
Jan.	2392	6056	1183	384	2548	12 563
Feb.	2863	6946	1139	592	3045	14 585
Mar.	3602	8164	1173	827	3191	16 957
Apr.	3307	8484	1097	871	3192	16 951
May	2703	7067	1154	787	2870	14 581
June	2849	7195	966	870	2813	14 693

single transfer may even be their only use of the network. At the other extreme are individuals who make very clever uses of the IT connection type to perform such tasks as distributed multiprocessing.

The batch services were not initially available and saw only limited use when they were introduced, but the number of these connections continues to increase. To some users they are very important. For example, the resources on Wayne State University's MVS system are accessible only in batch mode. Another reason for interest in batch is that it is easier to use a foreign system in this fashion than it is to become an effective interactive user. This is particularly true for a classroom application, where a course instructor can easily set up a specific set of commands to compile and execute a source program for his class on a remote host. In cases like this, the students need very little, if any, understanding of another computer's command language, editing system, and file system.

A final note about batch services is that, in addition to the execute and print services described and reported on earlier, there is a third service that carries data to be punched. Although this service has been available just as long as the other batch services, only a handful of punch connections are opened each month. Evidently our users have little use for cards; this is not a great surprise.

Another set of observations to make about Merit's statistical experiences pertains to the relationships between packets, bytes, and the length of network connections. The average number of bytes per packet is about 50 when all connection classes are aggregated. When the various connection types are considered individually, the averages vary from a low of about 35 for the interactive TL and X3 types to over 115 for the batch EX type, as shown in Fig. 30. But all of these averages are much smaller than Merit's maximum packet size of 240 bytes.

What causes these particular values? Again the answer is fairly obvious. All of the hosts on the network except for WMU's DEC system and the newly added VAX use record-oriented input and output rather than character-stream I/O. This means that input records are card images or successive but individual lines of information from terminals. Likewise, output is typically a series of printer lines or terminal output lines. Therefore, most input lines contain fewer than 80 characters and most output lines fewer than 135 characters. For this reason, it is not at all surprising that the packets transmitted over the network average between 35 and 115 bytes.

It should be noted that while Merit's maximum packet size is 240 bytes, it is possible to send very long records as strings of maximum-length packets. For example, a single record of any length up to 32 767 bytes may be transmitted between the MTS hosts. Evidently few long records are sent, but this capability is exercised whenever object files are transferred between hosts and for some other purposes as well. This ability to ship long records partially account for the somewhat higher byte per packet ratio for IT connections as contrasted with TL connections.

The ratio of either bytes or packets to the duration of a connection demonstrates the dramatic superiority in transmission rate achieved by batch connections, both EX and PR types, over the much slower interactive connections. The relative throughput rate differs by more than an order of magnitude, as shown in Figs. 31 and 32. The primary reason for this difference is surely that I/O from an interactive terminal is much slower than interhost, machine-controlled data transfers. But there are also some technical factors related to Merit's host-to-host protocols and the use of data compression within the network that influence this difference in rates.

Merit's average packet per hour rate for TL connections is 890 and for X3 is 998. These rates compare closely with the kilopacket per hour rate experienced by some other packet-switching networks.

A factor of frequent concern to users is how responsive a computing system is. There are many ways to quantify response performance. For example, the National Bureau of Standards [3] has defined various performance measures and conducted some experiments related to this issue. Merit's experience has been that for typical interactive computing sessions the network does not introduce noticeable additional response delays. Any user-observable response sluggishness is attributable to a host system. On the other hand, a network file transfer between hosts takes much longer than would a comparable intrahost data transfer. This delay is due in part to the relatively low-speed data lines linking the Communications Computers, and in part to idiosyncrasies of the supporting interactive protocols in the hosts, and for this reason users are advised to transmit long files between hosts using the network's batch services. In any case, Merit has not attempted either to systematically record or to seriously study response performance.

[3] As of June 1981 there were 32 ports in Detroit; 17 ports in Ann Arbor; 10 ports in East Lansing; and 8 ports in Kalamazoo. Telenet is serviced through an X.25 synchronous link capable of supporting 32 concurrent Inbound connections.

REFERENCES

[1] W. D. Becher and E. M. Aupperle, "The communications hardware of the Merit computer network," *IEEE Trans. Commun.*, vol. COM-20, June 1972.

[2] E. M. Aupperle, "Merit computer network experiences," in *Proc. Int. Conf. Commun. Conf. Record*, vol. III, June 16-18, 1975, pp. 33-2-33-6.

[3] M. D. Abrams, S. Treu, and R. P. Blanc, "Measurement of computer communications networks," NBS Tech. Note 908, July, 1976.
[4] E. M. Aupperle, "Mini and microcomputers and computer communication," *The International Computer Technology Conference, ASME Century II—Emerging Technology Conferences, Advances in Computer Technology—1980*, vol. 1, pp. 265-270, Aug. 12-15, 1980.

Eric M. Aupperle (S'57-M'61-SM'65) was born in Batavia, NY, on April 14, 1935. He received the B.S.E. degree in electrical engineering, and the B.S.E. degree in mathematics in 1957, the M.S.E. degree in nuclear engineering in 1958, and the Instm. E. (the professional degree in engineering granted by the Program in Computer, Information, and Control Engineering, formerly the program of Instrumentation Engineering) in 1964, all from the University of Michigan, Ann Arbor.

From 1957 to 1963 he was an Assistant Research Engineer, from 1963 to 1967 an Associate Research Engineer, and from 1967 to 1969 a Research Engineer at Cooley Electronics Laboratory, the University of Michigan, working with fuze countermeasures, application of solid-state devices, and pseudorandom communication system countermeasures. He has been a Lecturer with the Department of Electrical and Computer Engineering since 1963 and a Research Scientist with the Institute of Science and Technology since 1969 at the University of Michigan. His present responsibilities include serving as Director for the Merit Computer Network, a joint interuniversity computer network activity of Michigan State University, East Lansing, the University of Michigan, Ann Arbor, Wayne State University, Detroit, and Western Michigan University, Kalamazoo. He is also an Associate Director for Communications at the University of Michigan Computing Center.

Mr. Aupperle is a member of Phi Eta Sigma, Eta Kappa Nu, Tau Beta Phi, Sigma Xi, Phi Kappa Phi, and the University of Michigan Science Research Club. He also has held several offices in the National Electronics Conference. In addition, he has acted as Chairman and faculty member of the University of Michigan Engineering Summer Conference's short course on minicomputers for several years. He is a Registered Professional Engineer in the State of Michigan.