CITI-TR-90-3 June 15, 1990

User Profile and Workload Analysis for Local Area Networks

Toby J. Teorey
David W. Bachmann
Thomas A. Makowski
Mandyam M. Srinivasan

Center for Information Technology Integration (CITI)
The University of Michigan, Ann Arbor, MI 48105-2016

Center for Information Technology Integration

The University of Michigan
Information Technology Division
535 W. William
Ann Arbor, MI 48103-4943

Center for Information Technology Integration (CITI)
The University of Michigan
535 West William Street
Ann Arbor, MI 48103-4943
(313) 763-5216

User Profile and Workload Analysis for Interconnected Local Area Networks

Toby J. Teorey
David W. Bachmann
Thomas A. Makowski
Mandyam M. Srinivasan

June 15, 1990

Abstract

Performance analysis tools for computer networks need accurate and comprehensive estimates of user workload. An approach is presented that estimates network impact for a wide variety of end user types and applications that are typical on local area networks. Fourteen user types and nine generic application types are defined, and data is collected to determine the average network bandwidth needed to accommodate the output of individual and aggregate user/application combinations. Workload is estimated using a combination of data obtained from live test experiments, and data collected from the literature. Finally, the implementation of this data in a highly interactive network modeling tool (NetMod) is illustrated with screen images generated during tool execution.

1. Introduction

Local area computer networks are now well established, and highly interconnected heterogeneous local networks are coming into prominence [Tane88, Stal88, Keis89]. The large-scale (tens or hundreds of thousands of sites) associated with interconnected local networks makes performance modeling difficult, and both analytical and simulation techniques are often employed to handle both large-scale and layered protocol details. In either case, the user workload has been traditionally specified in terms of the packet generation rate and packet size distributions at the host computers and/or individual microcomputers or engineering workstations.

Packet size distributions are usually known from manufacturer specifications and collected data on network traffic. The packet generation rates can also be collected from network traffic data. However, these rates do not reflect the real workload generated by the different workstations/computers. For instance, many subnetworks have users that fit a single profile, e.g. secretaries using Macintoshes connected with a LocalTalk network; or engineers developing CAD systems using an Apollo token ring, Sun Ethernet, or Digital DECNet. Thus, this method of data collection (based on network traffic) emphasizes "effect" rather than "cause" in terms of real workload generation, and thus is not at all flexible to predict the more detailed needs of specialized function networks or asymmetic network configurations. We want to be able to capture the exact details of potential network usage by studying the types of users and their most used applications.

At the University of Michigan, a large-scale heterogeneous environment has been developing rapidly, from a centralized computing system, the Michigan Terminal System (MTS) [ArBo69], in the mid-1960s, along with the MERIT wide area network [Aupp89] in the early 1970s, to a fully interconnected local network of 25,000 sites today, and estimated to expand to 50,000 sites withing the next five years. An analytical tool for modeling such an environment has been developed using NetMod [CST88, BSST90]. NetMod is a HyperCard based tool which requries the definition of icons representing approximately 50 different LANs, 50 interconnect devices, and 100 different collections of computer sites for individual departments and research centers on campus.

With such a wide variety of users (students, faculty, research staff, clerical staff, and adminstrators) and goals (term papers, programming assignments, research work, personnel records, grades, and financial management), there is a strong need for detailed user profiles and application type characteristics [FSZ83]. Thus an extensive focus study was conducted, in 1986-87 [Macr87], on computer usage by students, faculty, and administrative staff to determine how their computer time was broken down by application type, and what use of campus networks they were making or planning to make. Subsequently, network growth has exploded and the network usage has expanded into several new areas such as image processing and video. As a consequence the focus study data had to be updated considerably. This was accomplished with new interviews, live test data collection from the campus network, and in-house controlled experiments.

Section 2 defines the various user types and generic applications used in the study. Section 3 describes how data was collected for the user types and applications, and Section 4 shows how this information is used in the analytical tool, NetMod. Finally, Section 5 draws conclusions from the user profile study and modeling tool experience.

2. User Profiles and Generic Application Types

Table 2.1 illustrates fourteen generic user types we have identified for desktop computers and personal workstations as the basic components of interconnected local networks. They cover the spectrum from clerical people to administrators, technical categories of engineers and scientiests, and students and faculty in a university environment. In the second column we have estimated the expected number of hours each type of person would spend using a computer, and the succeeding columns estimate the percentage of that time spent with each of nine generic application types. The computation of these estimates are explained in Section 3. The far right column shows (in parentheses) the expected number of kilobits per second (kbps) generated by the given user type specified in each row. This number is taken as an idealized weighted sum of the expected number of kilobits per second generated by each application type from a workstation that has no constraints on speed. The packet generation rate for a real computer with real hardware constraints, used in an analytical model of user workload and system configuration, would be restricted by those constraints. The expected number of kilobits generated per second can be converted to packets per second (pps), given the expected packet size for the system confuration.

	Hrs/	**<*****Percentage of time for each activity******>**									
User Type	**day**	CAD	DB	Email	FileT	Prog	Scien	SpSh	Video	WP	Kbps
Administrator	1.5	0	10	20	0	0	5	30	0	35	1.0
Applications prog	7.0	5	5	5	5	40	30	0	0	10	54.7
Application soft user	6.5	10	30	5	0	0	15	20	0	20	18.2
Business analyst	3.0	5	20	10	0	0	5	35	0	25	5.3
Data entry clerk	7.0	0	30	5	15	0	0	40	0	10	136.8
Engineer - design	6.0	55	3	5	2	10	10	0	5	10	157.3
Engineer - production	1.5	0	15	15	0	0	0	15	5	50	19.4
Engineer - research	3.0	20	5	5	0	20	30	5	5	10	51.4
Faculty -engr/science	3.0	6	5	15	2	5	20	5	2	40	27.7
Faculty -arts/literature	1.5	0	10	20	0	0	10	10	0	50	0.8
Student -engr/science	4.0	8	0	15	0	35	15	5	2	20	28.0
Student-arts/literature	1.5	0	5	20	0	5	10	20	0	40	0.9
Systems programmer	7.0	10	5	5	10	60	0	0	0	10	102.1
Word proc./secretary	6.5	0	10	0	0	0	0	5	0	85	1.7

Table 2.1 User type profile: hours of computer use per day and % of time per application type.

The nine generic application types are defined as follows:

1. Computer Aided Design (CAD) -- graphical manipulation of 3-dimentional objects, including image processing.
2. Database applications (DB) -- database queries and update transactions in a commercial system, including real-time transactions.
3. Electronic mail (Email) -- sending and receiving messages from other sites.
4. File transfer (FileT) -- using a standard file transfer protocol to transfer large files across sites, possibly across wide-area networks.
5. Program development (Prog) -- software implementation, debugging, and testing.
6. Scientific programming (Scien) -- large scale number crunching (cpu-bound) tasks.
7. Spreadsheet (SpSh) -- Lotus or Excel like applications.
8. Video -- sending and receiving video images in real time.
9. Word processing (WP) -- use of a standard word processing or electronic publishing package with "save" capability.

In the next section we look at the ways in which the expected kilobit rate for each of these application types can be estimated.

3. User Profile Experimentation and Data Collection

3.1 Transformation of Focus Study Data into Network Traffic Model Input

This section analyzes the focus study data from [Macr87] for the University of Michigan user community workload. The sample workload is then transformed into our standard measure of the workload (in kilobits per second) imposed on the network. The relative user workload for active users (those who agreed to participate in the focus study) and average users is then estimated, and the ratio of peak workload per hour versus the average workload per hour is also estimated.

The focus study determined the connect hours per week each study participant estimated to use the computer, categorized by the various application types listed. In the original report [Macr87], there were 22 application types defined, and these have been mapped into the nine generic types listed above, excluding system overhead. There were 112 students, faculty, and administrative personnel participating in the study. The goal of this part of the analysis is to determine the total connect hours for the entire campus based on this sample.

The first step is to find the connect hours for the application categories by source of personnel:

A=administration (sample of 41, pop. of 9984*)
F=faculty (sample of 36, pop. of 2510)
S=students (sample of 35, pop. of 34,340)

*% using computers: professional and administrative (100%), clerical (100%), technical (25%), service and maintenance (10%).

The individual responses are contained in the original report. The cumulative connect hours for all personnel in each of the three basic personnel categories are listed in Table 3.1. Intra-off defines activity within the same LAN, inter-off defines activity across LANs within the campus, and external defines activities on a wide area network which is off-campus.

Applic.	**Pers**	**Intra-off**	**Inter-off**	**External**	**Total hrs/wk**
1. CAD	A	15	1	0	16
	F	14.5	5	0	19.5
	S	17.5	23	0	40.5
2. DB	A	120.3	27.5	8.8	156.6
	F	97.1	9	8	114.1
	S	4	13.75	0	17.75
3. Email	A	33.5	25	6.5	65
	F	52	77.2	21.2	150.4
	S	8.5	73.5	0	82
6. Scien	A	19.25	2.25	0	21.5
	F	106.2	5.5	7	118.7
	S	36.75	44.5	0	81.25
9. WP	A	81.25	16.2	17.25	114.75
	F	199.5	31	64.5	295
	S	83.5	23	0	106.5

Table 3.1. Connect hours per week estimated by focus study participants, summarized by category of personnel.

The cumulative connect hours per week resulting for each application on the network is computed by dividing each value above (hours/week for the sample group) by the sample size, and then multiplying the result by the total estimated population for that source. Combining over all sources we get the following total connect hours per week impact on the network (Table 3.2):

Applic.	**Intra-off**	**Inter-off**	**External**	**Total hrs**	**% of total**
1. CAD	21,834	23,159	0	44,993	9.7 %
2. DB/file	39,989	20,816	2,701	63,506	13.8 %
3. E-mail	20,123	83,584	3,061	106,768	23.1 %
6. Scien	48,149	44,592	488	93,229	20.2 %
9. WP	115,619	28,685	8,697	153,001	33.2 %
TOTAL	245,714	200,836	14,947	461,497	
	(53.2 %)	(43.5%)	(3.2%)		

Table 3.2. Connect hours per basic application type.

3.2 Experimental Data on User Workloads

The goal of this experiment was to estimate the current total user workload from the connect hour data collected in the focus study and summarized in Table 3.2 above. Several small experiments were run to collect Michigan Terminal System (MTS) statistics on four application types to see what the network traffic might be. The following statistics were collected:

Applic. (% of total)	**Diskio's /conn.hr**	**Adjusted diskio's**	**Bytes/ conn.hr**	**Connect hrs/wk**	**Bytes/ second**
3. DB/file (13.8 %)	17,498	13,561	55.5 MB	59,391	22.9 MBS
4. Email (23.1%)	6,099	152	.6 MB	106,767	.4 MBS
6. Scien (20.2 %)	19,328	14,979	61.4 MB	93,230	39.8 MBS
9. WP (33.2 %)	1,107	858	3.5 MB	141,968	3.5 MBS

					66.6 MBS

Table 3.3. Sample applications with accounting statistics shown for diskio's.

Diskio's per connect hour were obtained from the MTS accounting statistics at the end of each sample job. Three test cases were run for each application type with connect time varying from 2 to 26 minutes.

Adjusted expected diskio's per connect hour were obtained from UM Computing Center statistics showing that 25% of diskio's are for the system and 75% are for the user. System diskio's experience 90% cache hits (only 10% physical io's) while user diskio's require almost no cache hits (100% physical io's). The exception is E-mail which has almost 100% cache hits (0% physical io's) in the user data part.

Adjusted diskio's = .25 x diskio's x .1 + .75 x diskio's x 1.0
for all applications except E-mail.

= .25 x diskio's x .1 + .75 x diskio's x 0.0
for E-mail.

Connect hours per week are taken from Table 3.1. We determine bytes per connect hour and bytes per second by:

Bytes/connect.hr = Adjusted diskio's/connect hr. * 4096 bytes/diskio

Bytes/second = (Bytes/connect.hr) x (Connect.hrs/week) / (seconds/week)

Note: seconds/week = 40 hrs/week x 60 min/hr x 60 sec/min = 144,000

Now that we have bytes per second we only need to convert to packets per sec (pps). The model assumes that packet size is bimodal [NTI89]:

50% message packet: 512-1024 bytes uniformly distributed
50% acknowledgement: 10 bytes

The average packet size is therefore .5[10 + .5(512+1024)] = 389 bytes, and the projected total user workload is:

User workload = 66.6 MBS/(389 bytes/packet) = 171,208 packets/sec

Our experimental applications only represented 90.3% of the total workload, however, so we need to divide our result by .903 to get an estimate of the total workload.

Adjusted user workload = 171,208/.903 = 189,599 pps

Since there are approximately 25,000 workstations in the network today, the workload on each workstation is estimated at 189,599 pps/25,000 = 7.6 pps.

3.3 Refining the Accuracy of the Focus Study: Active Users vs. Average Users

We hypothesized that the participants in the focus study represented a class of very active computer users. If we extrapolate the user workload for these people across the whole campus we could be grossly overestimating the real workload, at least for the near future. To test this hypothesis we used the focus study data which categorized the inter-office application usage by MTS and LAN resources. Only the data for the students was chosen. Since the students dominate the user population we assumed they represented the entire population on this issue.

The focus study showed that students used MTS for 60.9% of their inter-office applications and LANs for the other 39.1%. Thus the weekly connect hours used for MTS would be 60.9% of 189,599 namely 115,466 hours. At the time the focus survey was conducted on the students, the total MTS usage was 141,357 connect hours. Assuming 4.3 weeks per month, the weekly average connect hours for one month would be 141,357/4.3 = 32,874 hours. Thus the ratio of the workload for the heavy users to the workload for an average user was estimated as 115,466/32,874 = 3.5.

3.4 Estimating the ratio of peak/average user workload over one hour.

We again obtained data from the University of Michigan Information Technology Division (ITD) that showed maximum (peak) and average workload statistics over various time periods. The available data showed:

1. Number of terminal users active at 15 minute sampling intervals over the entire day for a full month.

2. Number of terminal jobs between 9-5 each day for an entire month.
3. Number of terminal jobs over an hour's time for an entire month.

The peak/average ratio for the number of terminal users at 15 minute intervals ranged from 1.64 to 2.38 for three month's data. The terminal jobs between 9-5 measured over that entire time period had a peak/average ratio from 1.39 to 3.49 for the same three month's data. Both of these measures were seen as inferior to the third measure because the number of terminal users is not as good a measuere as the number of terminal jobs (a one-to-many relationship), and the time period of 8 hours (9-5 totals) was not a good sampling interval for peak usage. The best measure collectable on these systems was the number of terminal jobs submitted per hour over a full month's period. The global average of all the monthly peak/average ratios was 3.36. From this we conclude that a 3.5-to-1 peak/average ratio is a reasonable estimate, although a 7-to-1 ratio was seen at least once over the 7 month period.

3.5 Default Values for User Workloads in NetMod

The focus study results and live test experiments described above provided most of the needed data to estimate application workloads as well as user profiles for students, faculty, and administrators at a large research universtiy. The remaining data was estimated from our experience and from the technical literature. Fortunately, in NetMod (see Section 4.3) the analyst has the option to either use the default values or customize the tool to his or her own network environment. This flexibility is critical to the effective use of such analytical tools in network performance modeling and evaluation.

We now describe each of the nine generic application types in terms of their default network traffic impact (workload), and how each value was derived.

1. CAD/graphics/image processing: 160 kbps

The experiment for CAD applications (Suns, VAXstation and Ethernet in Fig. 3.1) used a screen of 1024x864 pixels (.8 MB) to generate an average of 34.6 packets per second at 578 Bytes per packet over the Ethernet. The experiment consisted of an extended period of 15 minutes of 3-dimensional object rotation, animation and parameter manipulation. The viewing program ran on the Sun 3/280 (citi), producing PHIGS (a graphics standard) commands which were translated by the device driver program on the Sun 3/75 (tahoe) into X Window System protocol requests which displayed the results on the VAXstation 2000 (chicago). The tcpdump monitoring program was run on the Sun 3/60 (hbomb) to record the traffic on the Ethernet. 34.6 x 578 x 8 = 159.99 kbps.

2. DB access, entry, on-line research: 7 kbps

Three sets of experimental data were collected at the University computing center for different database management systems: MTS database experiment in Sec. 3.2 (7.6 kbps), Oracle (1.8 kbps), and a library card catalog system, MIRLYN(24.1 kbps). The selected value, 7 kbps, represents the middle ground among the three disparate values.

3. Email/conferencing: 3 kbps

The experiment for email (Sec. 3.2) resulted in 3.7 kbps averaged over a 15-20 minute period of constant work. Another experiment conducted in the configuration of Fig. 3.1 returned values ranging from .6 to 1.2 kbps on a LocalTalk network. The selected value represents something closer to the higher rate that appeared for remote email access.

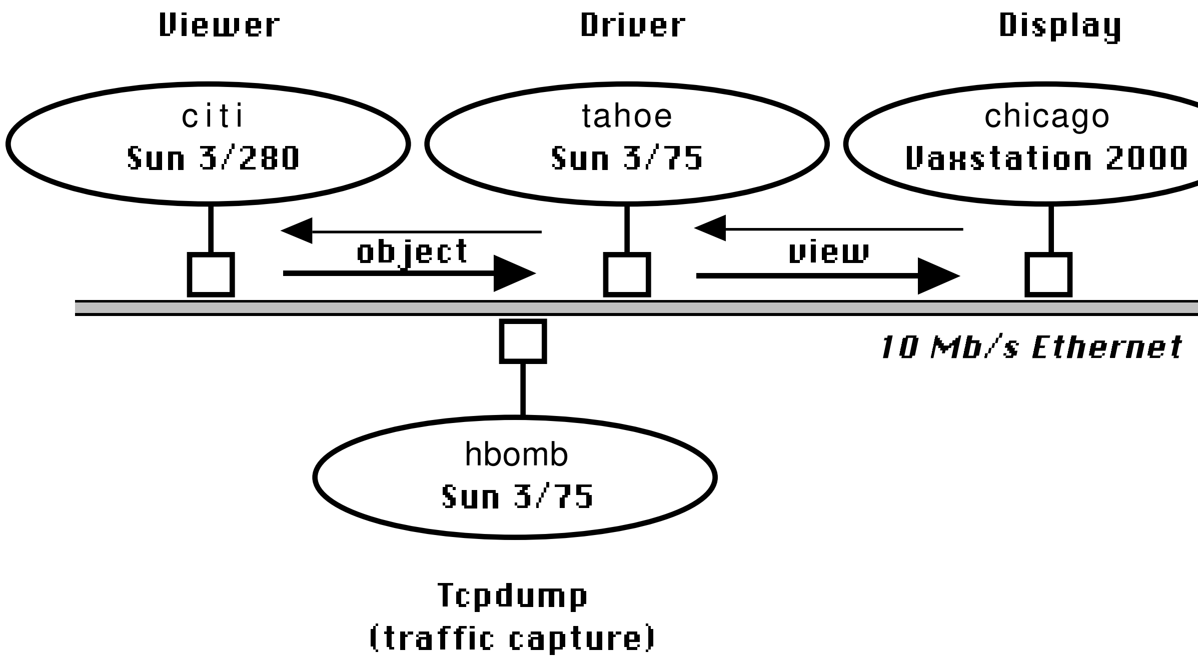

Figure 3.1 Experimental setup for CAD data collection

4. File transfer: 1000 kbps

The experiment for file transfer (MicroVAX and Ethernet in Fig. 3.2) used a file transfer between two MicroVAX (Ultrix) file servers across an empty and a busy Ethernet using the Network File System protocol (NFS) and the File Transfer Protocol (FTP). For light to medium loads the transfer rates were consistently the same, about 60 pps for NFS and 100 pps for FTP, resulting in bit rates ranging from 711-1160 kbps.

5. Program development: 0.1 kbps

This was esimated as one-tenth the rate used with a commercial word processing package, given the time a programmer needs to think about the next entry relative to a work processing user.

6. Scientific computation: 13 kbps

The experiment in Sec. 3.2 on scientific computation resulted in an average of 13.2 kbps used on MTS.

7. Spreadsheet: 10 kbps

This was estimated as halfway between database and scientific computation, which are the two basic features of spreadsheet progams (macros).

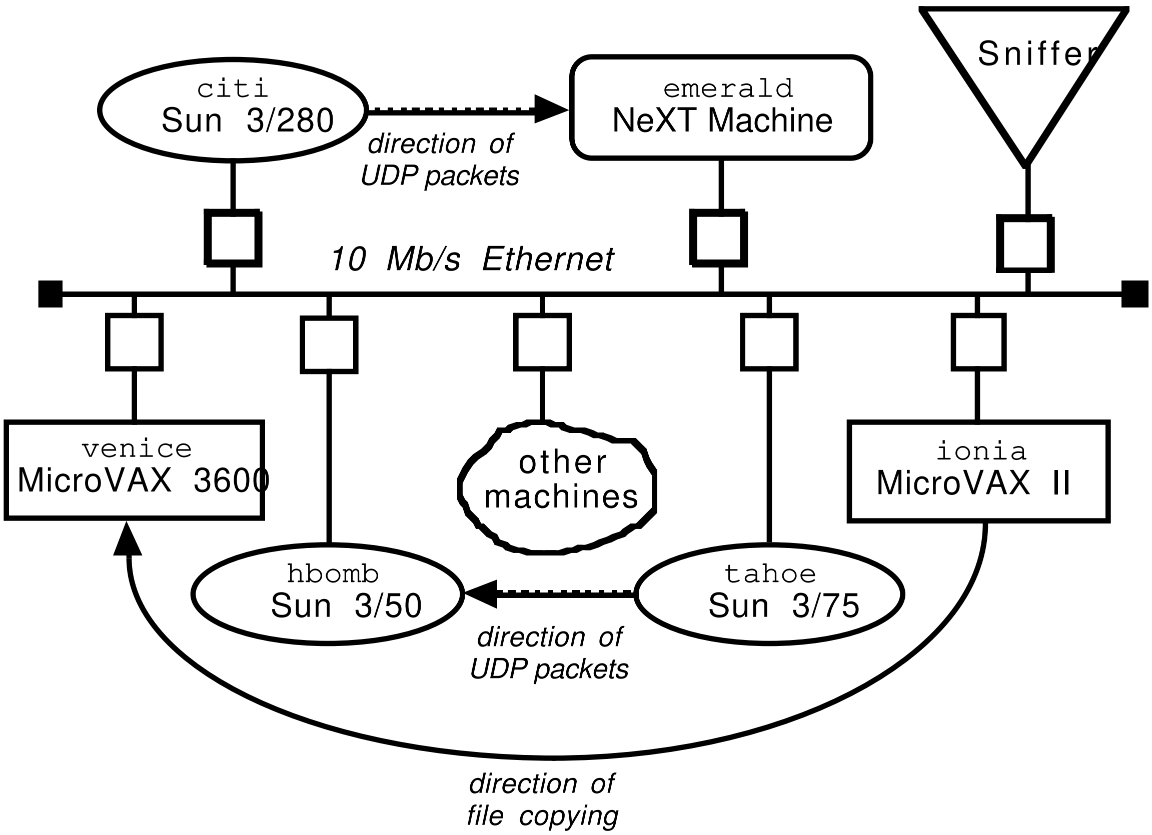

Figure 3.2 Experimental setup for file transfer data collection

8. Video (motion): 2000 kbps

This was estimated from discussions with video engineers who maintained that slow video was feasible for T1 links, but could easily use an order of magnitude increase in medium capacity. Thus we estimated a range of 1.544 - 15.44 Mbps for marginal video transmission, and selected a value near the lower end of this range to avoid total domination of the other values used in this study.

9. Word proc./electronic publishing:1 kbps

The experiment in Sec. 3.2 obtained an average rate of 1.15 kbps for word processing using the MTS file editor from a personal computer to a mainframe.

4. User Profiles in the NetMod analytical tool

4.1 NetMod

The Network Modeling Tool (NetMod) uses simple analytical models to provide the designers of large interconnected local area networks with an in-depth analysis of the potential performance of these systems. This tool can be used in either a university, industrial, or government campus networking environment consisting of thousands of computer sites. The principal application of the tool is to assist the campus network designer in configuring a potential user's data network using proposed network hardware and software components.

NetMod can analyze an existing or proposed network in terms of its basic performance characteristics, e.g. component utilization, throughput, and packet delay times. It provides a capability for sensitivity analysis of the performance based on changes in the workload parameters and either minor or major changes in the network topology and connectivity. Comparisons of alternate configurations will be possible as part of the network design process. NetMod models popular LAN technologies such as token ring and Ethernet and their variations, as well as special network components such as routers, bridges, gateways and adapter cards. Network operating system details are beyond the scope of the existing analytical models, but are the subject of ongoing research.

The tool enables the designer to display the current or proposed system topology graphically on the screen. This is currently done through a HyperCard interface on a Macintosh computer. The designer uses the given topology to evolve to a system design based on his or her experience or understanding of the user's requirements, and knowledge of the conceptual model. The tool then estimates the performance of the design and displays it to the designer.

The facilities provided by NetMod include ease of subnetwork definition, the flexibility to easily include/delete other existing subnetworks, the ability to provide substantial user interaction for either simple parameter changes or major reconfiguration of LANs and backbones, and the provision of a hierarchical modeling capability for extremely large models. For example, a model for the University of Michigan's current and proposed network includes over 160 submodels representing different types of LANs, interconnect devices, mainframes, and workstations (including personal computers); approximately 25,000 devices are represented in this model.

HyperCard is a software tool for developing applications for the Macintosh. What makes it unique is that it requires no understanding of the inner workings of the Macintosh. NetMod uses a HyperCard stack to allow the user to create their network graphically. Network submodels are represented by buttons which are interconnected to form the complete network. Calculations can also be made for the model within HyperCard and their results can be displayed with the Network submodels.

4.2 NetMod Features

The user inputs the network description through a Hypercard interface. Alternately, the user could choose to input the description using Excel dialog boxes. In the input process, the user decomposes the network into submodels (for example a departmental Ethernet might be one submodel). These submodels are analyzed using Excel spreadsheets. Parts of the spreadsheet calculate the level of traffic that passes through a particular submodel, while other parts use these traffic levels to calculate the output statistics for that submodel. Section 3 presents a detailed discussion on how NetMod is used. In this section the parameters input by the user, the output statistics generated, and the submodels used, are described.

User Inputs

The inputs to NetMod are i) the topology parameters, ii) the workload characteristics, and iii) the characteristics for the devices/transmission media in the network. These are input using either the HyperCard interface or the Excel dialogues. With the Hypercard interface, the user can choose a graphical (Hypercard) output. Alternately, the input can be passed to the Excel software, with the output appearing in the form of a spreadsheet. The output statistics for the entire model are generated.

In the following, the term node will frequently be used to refer to a device or a LAN.

Topology parameters

To describe the network topology, the different node types, namely host computer, workstation, LAN, or other network device such as gateway, router, bridge, etc. are input. In addition, their connectivity must be specified, namely which nodes are connected to each other.

Device/transmission medium characteristics
For each device in the network, information on the device throughput capacity such as packets per second (pps) is input. For each LAN, the type of LAN such as IBM Token Ring, Ethernet, etc., and the transmission medium capacity in Megabits per second (Mbps) is input. NetMod provides the known capacity of the device or LAN as a default value, which the user may replace if desired.

Workload characteristics
Included here are the packet generation rate for each end node type (typically a workstation or mainframe) in packets per second, the average packet size in bits (either averaged over messages and acknowledgements (ACKs), or obtained from some given distribution), and the routing probability for messages arriving at a LAN or network interconnect device. The routing probability specifies the probability that a packet arriving at this node has a destination beyond this node. The packet generation rate for each workstation is specified by the user. For all the other nodes it is internally generated, based on the routing probabilities. For each workstation type, a default packet generation rate is used; this can be changed by the user. The generation of packets at the workstations is assumed to follow a Poisson process.

The hierarchical model of the entire network is decomposed into a number of interrelated submodels each of which is analyzed in greater detail. Typically, four parameters are passed to each submodel:

λ: the mean number of packets per second,
size: the mean size of packets in bits,
size2: the second moment of packet size in bits,
speed: the transmission speed of the medium in bits per second.

In addition, for those submodels that implement bulk arrivals, the mean number of packets per message, c, and its second moment, $c^{(2)}$, is passed. Currently it is assumed (except for the interconnect devices mentioned later) that the buffer capacity for the nodes in the network is practically infinite, i.e. it is sufficiently large that it is not overflowed. Additional features targeted for future implementation include the modeling of finite buffer sizes and flow control mechanisms.

Output Statistics

The principal output statistics provided by NetMod are the mean delay experienced by a packet that is transmitted across the network, (obtained as the sum of mean packet delays in milliseconds across each submodel) and the effective utilization of the transmission media for each submodel. If the submodel is saturated, i.e. unstable, the packet delay function will pass back the word "saturation".

The utilization function passes back the utilization of the submodel as a percentage. If the submodel is saturated, this percentage will be a value greater than 100%. Although technically inaccurate, this feature allows the designer to assess the degree to which the submodel is overloaded (e.g. if the utilization of a token ring is shown as 200%, the traffic through that submodel must be cut by at least one half).

4.3 Using NetMod Visualization Techniques (HyperCard)

A new network model is initiated by selecting "new model" from the Options Menu (Fig. 4.1). New submodels are created by selecting one of the icons from the palette on the left side of the window. These icons can be thought of as a "library" of pre–defined submodel types. Each icon is, in fact, a pop-up menu of alternative submodel types within a class. As more submodel classes are added to the icon palette, the arrows at the top and bottom of the palette will scroll the entire set of icons up or down. The topmost icon in the palette shown in Fig. 4.1 is a ring icon. If this icon is selected with a mouse, it pops up a menu with entries for FDDI, IBM Token Ring, Apollo Ring, Proteon Ring, etc. If one of these menu items is selected the corresponding submodel button appears on the screen. Each submodel in Netmod represents a generic "class" of networks or nodes. These buttons have a set of attributes such as network type, speed, name, etc. each of which has a default value which may be changed by the user.

In Fig. 4.1 there is an Ethernet connected to the FDDI backbone ring via a router. The details of the router are specified by double-clicking the router icon and setting the appropriate attributes (router type and throughput capacity) in the router window. If there are other subnetworks that are not identical in terms of detailed connections and/or parameter data, they must be modeled as two separate buttons with different names.

Buttons with icons representing rings, buses, stars, gateways, and workstations may be linked and moved around on the screen. Buttons representing other subnetworks may be used to reduce the clutter on one card; double-clicking the button takes the user to the card for that subnetwork (Fig. 4.2). Double-clicking on any of the other buttons brings up a form displaying the current attributes of that submodel such as network type, speed, name. etc. which can be changed by the user. A submodel can be deleted by merely dragging it into the trashcan icon. Clicking on the question mark icon brings up an identical card whose buttons describe their function when clicked on.

Two submodels, e.g a network and a workstation, may be linked together by clicking on the "connect" button. Alternately, the "Connect Submodels" can be chosen from the options menu, then the higher-level submodel (the "parent") is clicked followed by the lower one (the "child"). A line will then be drawn between the two icons. Links are broken in a similar manner.

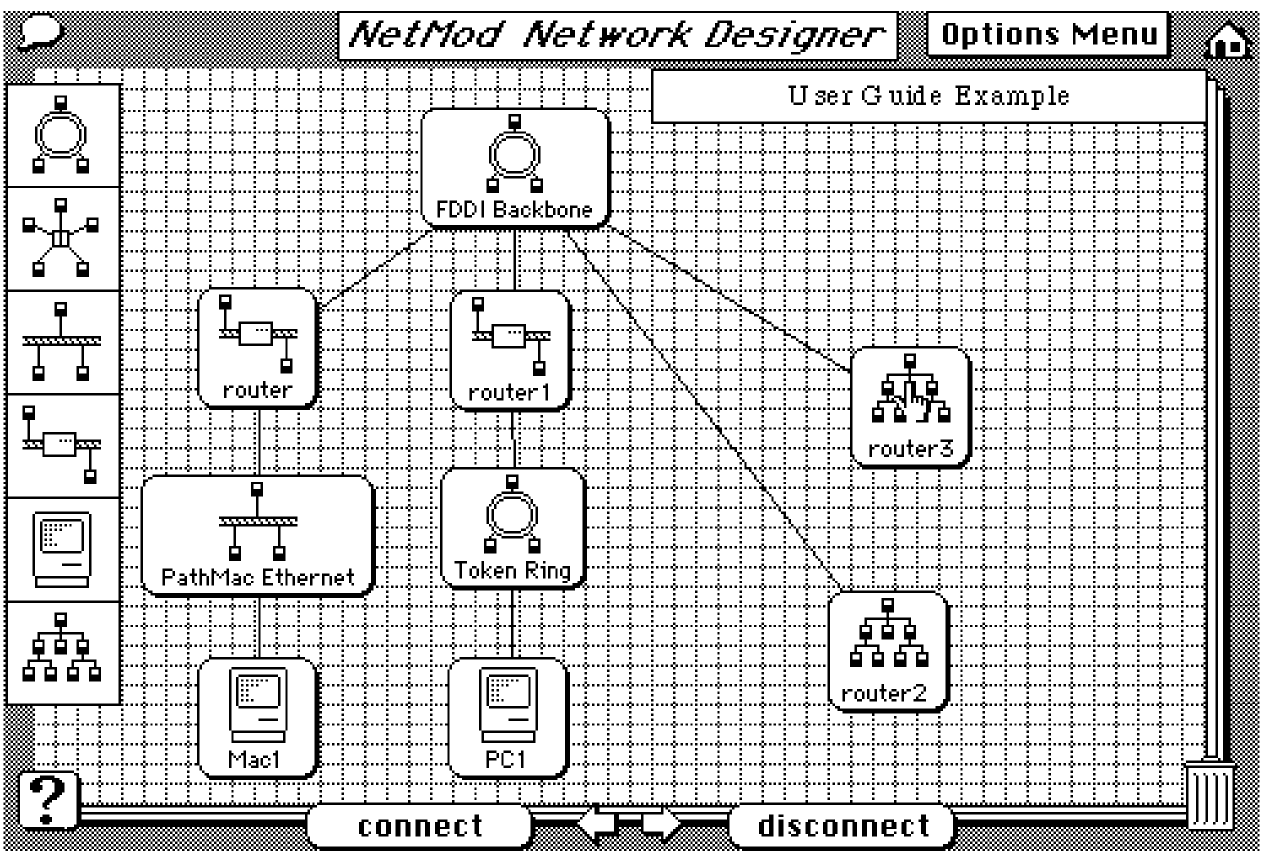

Figure 4.1 NetMod window for drawing a network topology.

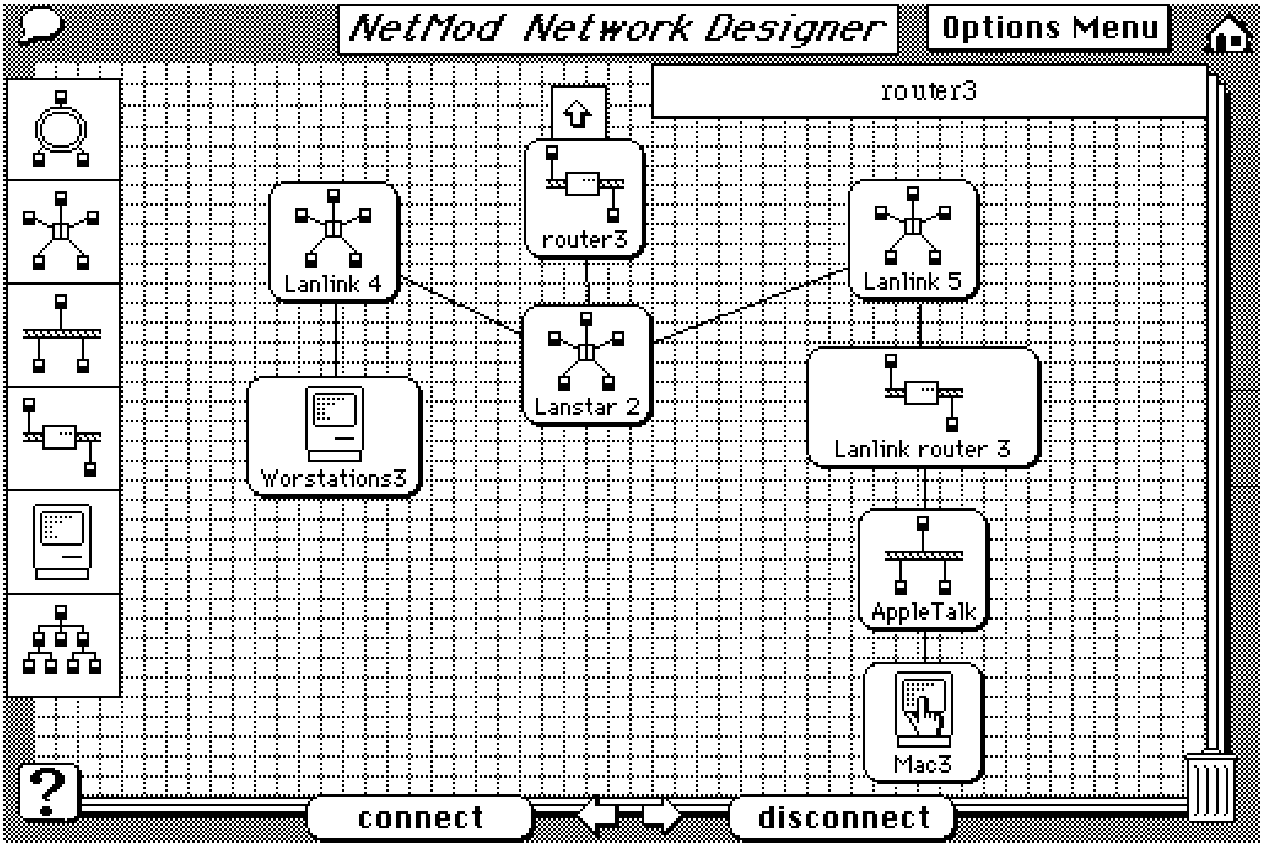

Figure 4.2 NetMod window for subnetworks (continuation of initial window).

4.4 User Profile Specification: User Category

The user profile of, for example, the Mac1 workstation in Fig. 4.1 is specified by double-clicking on the workstation icon. When the new window pops up (Fig. 4.3), an appropriate user category, e.g. "Application user" is selected from the user list by clicking on the name. The User Category field will automatically be filled in with this name. Also, the corresponding packet generation rate and packet size for this user type will be placed in the Packet Rate field and Packet Size field respectivel. The Packet Rate field and the Packet Size fields can be filled in manually if desired. On the other hand, the number of Mac1's must always be manually specified . Any of these values can be modified (customized) at any time to fit the particular network configuration being analyzed.

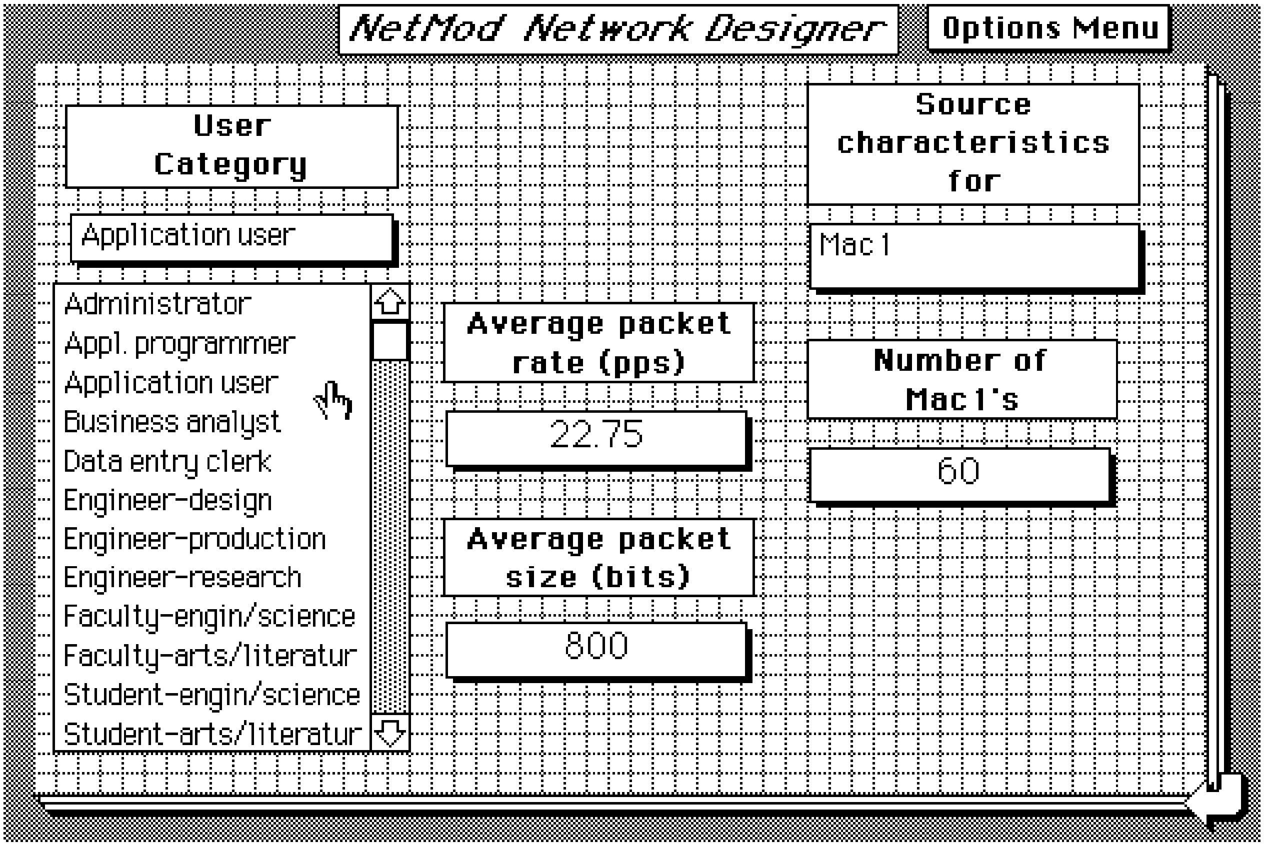

Figure 4.3 User Category window in NetMod.

Modifying, Adding, and Deleting User Categories

The user categories can be modified and deleted as desired. Also, new user types can also be defined. To modify a user, the option "modify user" is selected from the options menu. This produces the window shown in Fig. 4.4.

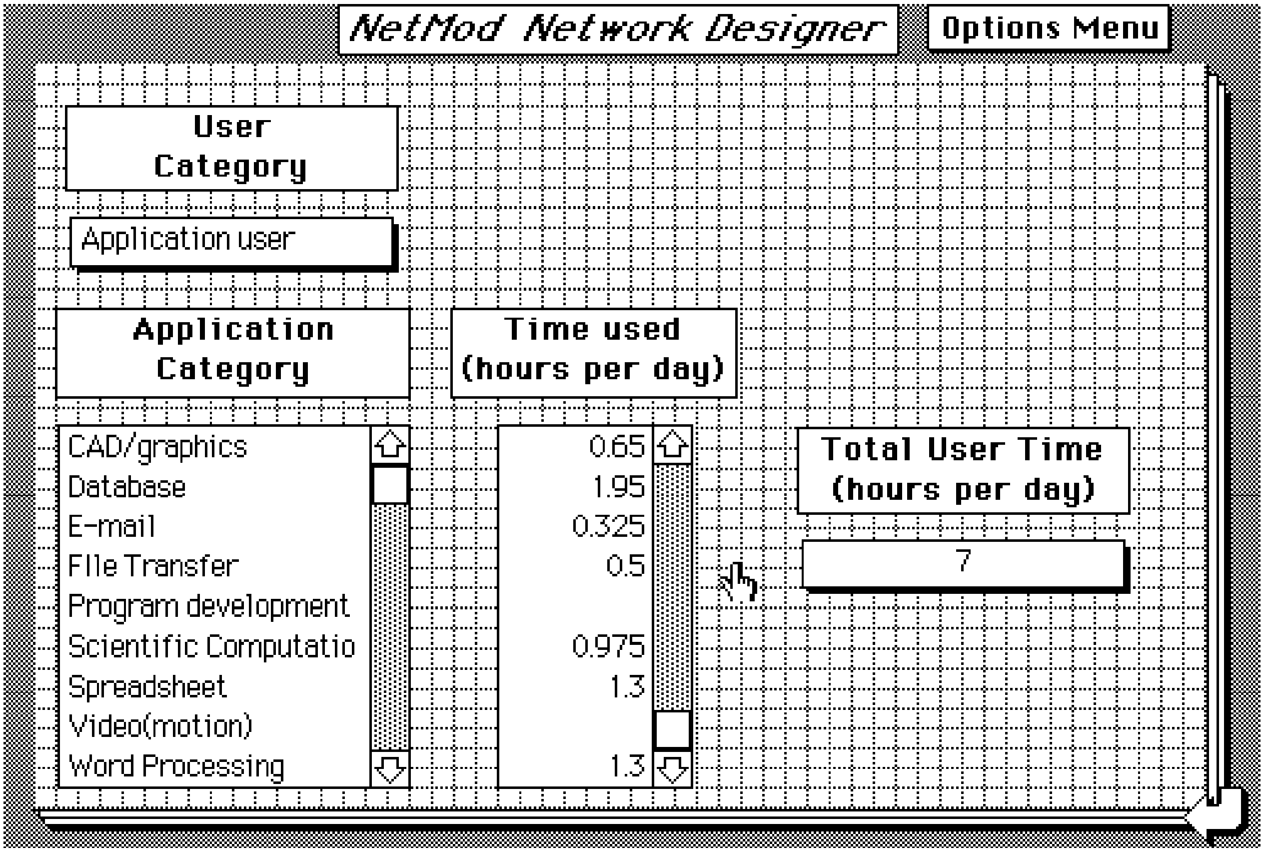

Figure 4.4 User category modification in NetMod.

A user category is defined by the amount of time per day the user will be actively using each type of application. To alter the user category, simply change the Time Used field for any desired application type. The total user time field will be updated according to the changes that are made. For example, we can add "0.5" in the line corresponding to file transfer for the "Application user" category as shown above, and click on the return arrow in the bottom right corner, confirming the change when prompted. The packet generation rate as well as the packet size for this user category will be automatically updated, and the newly updated "Application user" category will now be available for use anywhere in the network model.

To add a new user category, the option "add user" is selected from the options menu. This causes a prompt for a name. The category will be immediately added and the control will be transferred to the modify user screen, at which time the new category may be specified by filling in the Time Used field as appropriate and then saving the changes. To delete a user category, the option "delete user" is selected from the options menu.

Note that the changes to user categories will not alter the packet rate and packet size for other workstations in the network. In other words, the current packet rate and size to which they are set will remain unchanged until a new or updated user category is actively selected for them.

Modifying, Adding, and Deleting Application Categories

You may also modify, add, or delete the application categories which define the users, entirely similar to the way in which the corresponding operations are done for user categories. To modify an application category, the option "modify application" is selected from the options menu.

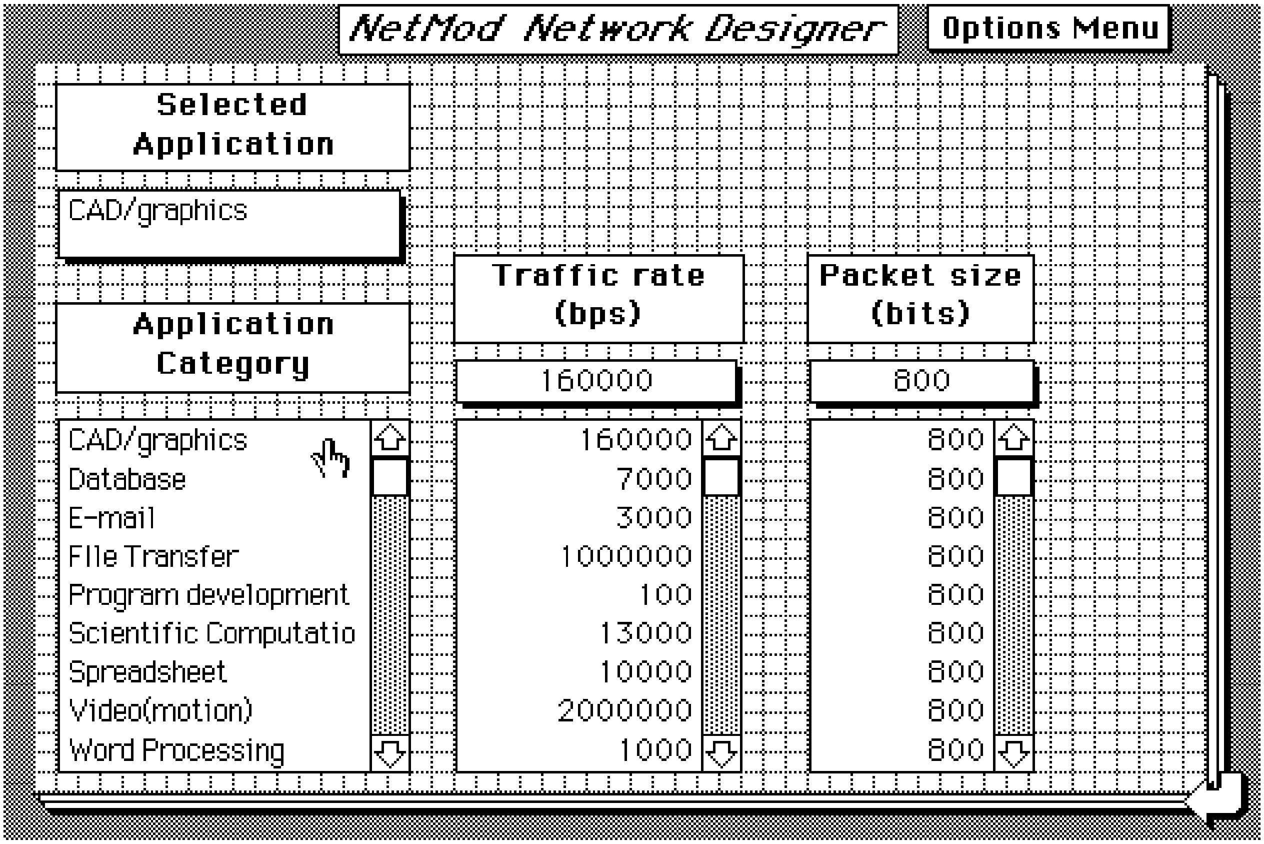

Figure 4.5 Application category modification in NetMod.

Applications are defined by their average traffic production rate, specified in bits per second, and their average packet size, specified in bits. To change the specifications, the category is chosen by clicking on the name in the Application Category field. The current rate and size specifications will automatically be placed in the Traffic Rate and Packet Size fields respectively. These values can be modified by altering the the Traffic Rate and Packet Size fields, and automatically saved. Applications can be added and deleted in the same manner as user categories.

5. Conclusions

The Network Modeling Tool (NetMod) has been shown to be a very pragmatic tool for the design and analysis of medium to large interconnected campus networks, involving up to tens of thousands of computer workstations. It helps the network system analyst to not only thoroughly understand the device interrelationships in the network, but also develop a detailed view of user profiles to provide realistic workloads in the models. Because of the volatility of meaningful data on user workloads, flexibility is built into NetMod to have both default values based on live-test experiments conducted on local networks as well as to allow NetMod users to customize the user categories to conform to their own user environment.

Acknowledgements. This work was entirely sponsored by Northern Telecom Inc.

References

[ArBo69] B.W.Arden, and D.Boether, "Measurement and Performance of a Multiprogramming System," *Proceedings 2nd ACM Symposium on Operating System Principles*, Princeton University, October 1969, pp. 130-146.

[Aupp89] E. M. Aupperle, "Building MERIT," *Harvard Info. Tech. Qtrly.*, Vol. 8, No. 2-3, Summer-Fall 1989, pp. 7-10.

[BST89] D.W. Bachmann, M. E. Segal, and T. J. Teorey "On the performance of copying large files across a contention-based network," submitted to *IEEE Network*..

[BSST90] D.W. Bachmann, M. E. Segal, M.M. Srinivasan, and T. J. Teorey "NetMod: A Design Tool for Large-Scale Heterogeneous Campus Networks," CITI-TR-90-1, Center for Information Technology Integration, The University of Michigan, February 15, 1990. To appear in *IEEE Jnl. on Selected Areas in Communication.*

[CST88] J. Chiarawongse, M.M. Srinivasan, and T.J. Teorey, "Performance Analysis of a Large Interconnected Network by Decomposition Techniques," *IEEE Network*, Vol. 2, No. 4, July 1988, pp. 19-27.

[FSZ83] D.Ferrari, G.Serazzi, A.Zeigner, *Measurement and Tuning of Computer Systems*, Prentice Hall, Inc., Englewood Cliffs, New Jersey, 1983.

[Keis89] G. E. Keiser, *Local Area Networks,* McGraw-Hill Inc., New York, 1989.

[Macr87] A.MacRae, R.Frost, D.Haggerty, S.Robinovitz, and T.J.Teorey, "A Wideband Campus Network Architecture at the University of Michigan," presented at the 1987 EDUCOM Conference, Los Angeles, CA., Oct. 27-30, 1987, CITI Tech. Rep. 87-5, Univ. of Michigan, Ann Arbor.

[NTI89] Northern Telecom, Inc., Private Communication.

[Stal88] W. Stallings, *Data and Computer Communications*, 2nd Edition, Macmillan Pub. Co., New York, 1988.

[Tane88] A. S. Tanenbaum, *Computer Networks*, 2nd Edition, Prentice Hall, Inc., Englewood Cliffs, New Jersey 1988.