Real-Time Traffic Measurements on MAGNET II

AUREL A. LAZAR, MEMBER, IEEE, GIOVANNI PACIFICI, MEMBER, IEEE, AND JOHN S. WHITE, MEMBER, IEEE

Abstract—**Real-time traffic measurements on MAGNET II, an integrated network testbed based on asynchronous time sharing, are reported. The quality of service is evaluated by monitoring the buffer occupancy distribution, the packet time delay distribution, the packet loss, and the gap distribution of the consecutively lost packets. Our experiments show that both time delay and buffer occupancy distributions of multiplexed video sources display a marked bimodal behavior, which does not seem to depend on the buffer size. The reliance of the network designer on traffic sources that do not exhibit substantial correlations can lead to implementations with serious congestion problems. For ATS-based networks with different traffic classes, the impact of a traffic class on the performance of the other classes tends to be diminished when compared to single class-based ATM networks.**

I. Introduction

In the past years, we have designed and implemented a new generation of integrated networks that covers local and metropolitan areas (see [15], [16] and the references therein). These networks transport services such as video, voice, data, graphics, and facsimile. The design of these networks was mostly based on simulation results that were in part reported in [21]. No measurements on real-time networks have been reported. Real-time measurements lead, however, to a better understanding of the actual behavior of integrated networks. This understanding can be used for future designs and implementations as well as for supporting the network management and control task.

Currently, the engineering community has to approve a series of standards for B-ISDN that require making fundamental decisions [8]. Unfortunately, except for data networks, no experimental data on real-time networks is available in the literature. This paper is a first attempt to fill this gap. It gives the first results of measuring the behavior of a real-time network, called MAGNET II, that was designed based on the principle of asynchronous time sharing (ATS) [14]. ATS-based networks guarantee the quality of service as negotiated at call set up.

In order to understand the behavior of MAGNET II and to validate the ATS concept, a group of experiments was designed. The investigations reported here are limited to a single switching node. A set of representative results was chosen for presentation in this paper. Investigations of traffic behavior and performance of a distributed networking environment will be published elsewhere.

Measuring the traffic behavior of a network in real-time is a large undertaking. This is due for several reasons. First, in order to test high speed networks, a distributed experimental environment has to be created that is capable of loading the network. Second, the traffic generators distributed throughout the system have to be representative for a wide class of sources that might arise in practice. Third, a distributed monitoring system has to be built in hardware and software that can follow the fast events on the network [18], [19]. Finally fourth, the evaluation of distributed measurement data has to be undertaken and presented to a monitor.

All experimentation reported here was done within a network monitoring and control environment called WIENER. In order to be able to monitor the activity of the network, a hardware observation unit was built. This unit is supported by a parallel programming environment that allows to control the resources of the network. The final measurement results are presented by a Console.

In the context of data networks measurements results of operational networks were reported in [1], [2], [3], [12]. We measured performance parameters using real-time traffic sources as reported in the literature in [5], [17]. Theoretical work also guided our investigations [20], [9], [6]. The measurement results obtained are quantitative and reveal a great deal of information about the traffic behavior of MAGNET II. These results were qualitatively compared to measurements obtained using a class of traffic sources that exhibit both correlation between arriving packets transmitted by the same source as well as different sources. In particular, we sought to find traffic models that approximate the range of behavior exhibited by multiplexed video sources. To this end, we employed in our experiments stochastic models for composite video sources.

Although this work was conducted on MAGNET II, the results of our experiments have significance for ATM networks in general. Based on our measurements, we made the following observations.

- The *a priori* knowledge of the individual traffic sources (characterized by first-order statistics) is not adequate for estimating the behavior of a composite source. Different scenarios for multiplexed sources can lead to a substantial degradation of the quality of service (QOS).
- The reliance of the network designer on traffic sources that do not have substantial intersource correla-

Manuscript received May 13, 1989; revised November 15, 1989. This work was supported by the National Science Foundation under Grant CDR-84-21402.

The authors are with the Department of Electrical Engineering and Center for Telecommunications Research, Columbia University, New York, NY 10027-6699.

IEEE Log Number 8934095.

tions can lead to implementations that will exhibit serious congestion problems.

- For ATS-based networks with different traffic classes, the impact of a traffic class on the performance of the other classes tends to be diminished when compared to single class-based ATM networks.

This paper is organized as follows. In Section II, the MAGNET II architecture is presented. The foundations are first described, the resource allocation methodology explained and the functionality of a network node outlined. In Section III, the WIENER experimental environment for traffic monitoring and control is explained in detail. The hardware as well as the software monitoring devices are discussed. Section IV starts with a description of the real-time video traffic source and the Poisson-based traffic source model that are used in the first four experiments. Experimental results are compared and discussed. In the second part of Section IV, three correlated traffic source models are presented and discussed. The results of the last four experiments are compared to those of the first four.

II. The Integrated Network Testbed

The integrated network used in our experiments, called MAGNET II, has a mesh topology and transports services such as video, voice, data, graphics, and facsimile. The core of the network does not make a distinction between these services. It recognizes instead a set of well-established traffic classes.

MAGNET II supports four classes of traffic. The class is an abstract concept that is specified through delay and loss characteristics. The user, prior to negotiating the quality of service, maps his application into one of the four classes. One class of traffic supports network management and control traffic. This traffic class will not be further elaborated in this paper. The other three classes support user traffic. Characterization of these traffic classes is given in Section II-A below.

For this multiclass network model, asynchronous time sharing refers to the manner in which scheduling and buffer management resolves contention among the different traffic classes. The general ATS framework is presented in Section II-A below. Switching and communication bandwidth allocation, and buffer management are briefly described. Finally, the implementation in hardware of the asynchronous time sharing concept is presented in Section II-B.

A. Asynchronous Time Sharing

The multiclass network model considered in this paper supports three classes of traffic. These classes transport user traffic and are defined by a set of performance constraints.

Class I traffic is characterized by 0% *contention packet loss* and an end-to-end time delay distribution with a narrow support. The maximum end-to-end time delay between the source and destination stations is denoted by S^{I}. Class II traffic is characterized by ϵ% contention packet loss and an upper bound on the average number of consecutively lost packets η. It is also characterized by an end-to-end time delay distribution with a larger support than Class I. The maximum end-to-end time delay is S^{II}. Here, ϵ and η are arbitrarily small numbers and $S^{I} \leq S^{II}$. *Contention packet loss* represents packets that, because of network congestion, had an end-to-end delay greater than the maximum limit (S^{I} or S^{II}) or were blocked by a buffer management system. For Class I and II traffic, there is no retransmission policy for lost packets. Class III traffic is characterized by 0% *end-to-end packet loss* with an end-to-end retransmission policy for error correction. If requested, it is also characterized by a minimum average user throughput Γ and a maximum average user time delay T.

For the multiclass network model described above, the concept of *asynchronous time sharing* (ATS) refers to the manner in which scheduling and buffer management resolves contention between the different traffic classes. *Scheduling consists of switching and communication bandwidth allocation, while buffer management refers to buffer space partitioning.* The essential requirement on these resource sharing mechanisms is to guarantee the appropriate quality of service for each traffic class. The quality of service is monitored and controlled by the traffic control architecture (TCA) of the network [14].

The general concept of the resource allocation strategies is shown in Fig. 1. Switching and communication bandwidth allocation is implemented using time sharing. Buffer management is achieved using space partitioning. The resulting resource allocation policy is, therefore, based on a space/time allocation mechanism. This basic mechanism allows the TCA to provide guaranteed quality of service.

In order to achieve switching (communication) bandwidth allocation, the switching (communication) bandwidth is divided into time periods called cycles (see Fig. 1). Each cycle is divided into three subcycles. During each subcycle (I, II, III), the switch fabric is allocated to the corresponding traffic class (I, II, III). For example, during subcycle I, Class I packets enter the switch fabric. The length of a subcycle is measured in cells. A cell represents the time required to serve (switch) one packet. The boundaries between subcycles are determined by a maximum length movable boundary scheme.

The TCA uses three variables (MAX I, MAX II and MAX III) to determine the maximum boundary positions between subcycles. MAX I represents the maximum length (in cells) of subcycle I. MAX II represents the maximum length of subcycles I, and II combined. MAX III represents the maximum length of the entire cycle. These variables are controlled by the TCA of the switch (link) and will dynamically change according to the traffic load and mix.

In addition to the maximum length constraint, a moveable boundary scheme is used. This method switches subcycles when no more packets of the current traffic class are available. Thus, at the beginning of a cycle, the switch

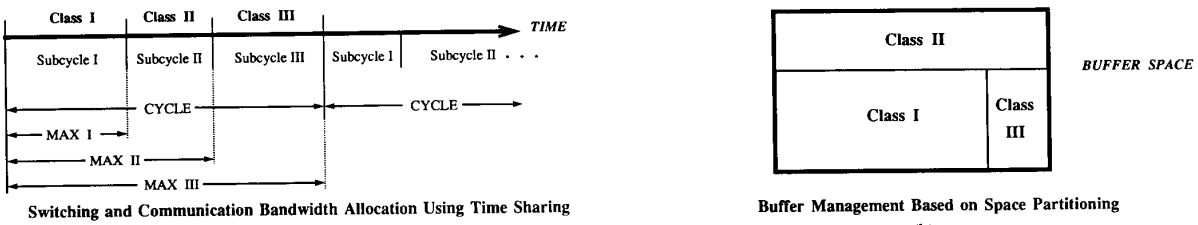

Switching and Communication Bandwidth Allocation Using Time Sharing

(a)

Buffer Management Based on Space Partitioning

(b)

Fig. 1. Resource allocation mechanisms.

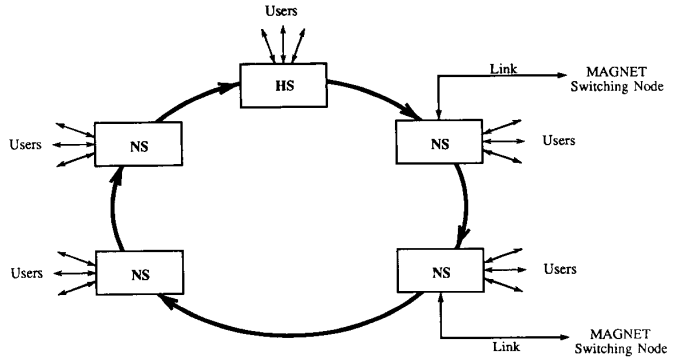

(a)

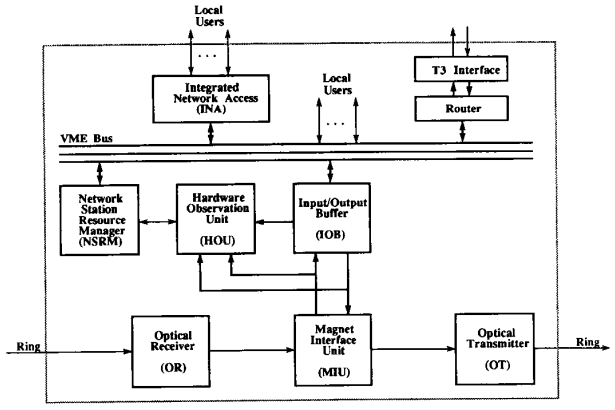

(b)

Fig. 2. MAGNET II switching node and network station architecture.

is allocated to Class I. The switch will serve Class I until either MAX I is reached or there are no more Class I packets available. At this point, the switch will change to subcycle II and serve Class II traffic. Class II traffic will be served until MAX II is reached or there are no more Class II packets available. When either condition occurs, the switch will change to subcycle III. Finally, a new cycle begins when MAX III is reached or there are no more Class III packets available.

For a given traffic class, the available bandwidth must be allocated fairly among the multiple access points. A method to limit access in order to guarantee users the appropriate bandwidth is used. Each access point is assigned three LIMIT variables (L^{I}, L^{II}, L^{III}) by the TCA. These variables are defined as the maximum number of packets of each class that the access point can transmit during one cycle. The TCA dynamically controls these variables according to the traffic load and profile.

Each access point also requires a buffer organization that supports the quality of service of the three classes of packets. Thus, buffer management is an integral part of the resource sharing strategy. Buffer management is achieved using space partitioning. The buffer space is divided into three areas using three counters B^{I}, B^{II}, and B^{III}. The values of these counters is again set by the TCA.

B. Switching Architecture

MAGNET II is a network testbed designed based on the ATS principle described above. The network has a mesh topology in which switching nodes are embedded. Communication between switching nodes is achieved via T3/

DS3 links. More details of the MAGNET II system architecture can be found in [15] and [16].

The MAGNET II switching node (see Fig. 2(a) consists of a set of network stations (NS's) and a headend station (HS). The stations are connected in an active slotted ring topology (called ring switch fabric) via 120 Mbit optical links. Both HS and NS stations provide network access to the local user. The HS differs from the NS in that it has an additional hardware block, called the cell generator (CG). As indicated in Fig. 2(b), the NS consists of the following functional blocks: optical receiver (OR), optical transmitter (OT), MAGNET interface unit (MIU) input/output buffer (IOB), the hardware observation unit (HOU), the network station resource manager (NSRM), and the local users.

For the purposes of this paper, a brief summary of the data path through each station is in order. For the interested reader, more detailed information of the ring switch fabric (RSF) can be found in [15]. The transmission time on the RSF is divided into fixed length (1024 bits), contiguous *cells*. The unit of information transfer on the RSF is called a *packet*. Cells are either empty if they are unoccupied, or busy if they contain a packet. Cells are generated by the CG at the HS. The continuous stream of cells generated by the CG is divided into cycles. Each cycle is divided into three subcycles (I, II, III). Each cell is marked by the CG to indicate the subcycle. Cells belonging to subcycle I can only be used by Class I packets; those belonging to subcycle II can only be used by Class II packets; and those belonging to subcycle III can only be used by Class III packets. This process of setting the length of the switching subcycles is called ring scheduling.

An optically encoded cell arrives at the optical receiver (OR) at a station. There, the clocking information is extracted and a $C = 100$ Mbit/s serial data stream is fed to the MAGNET interface unit (MIU).The MIU converts the serial data stream to 16-bit parallel words. The output goes to the input/output buffer (IOB). The IOB interconnects the VME bus with the RSF. It contains buffers for transmitting and receiving packets, controls local access to the RSF and supports various traffic control functions. The input buffer stores packets going from the RSF to users on the VME bus. Three output buffers (one for each traffic class) store packets going from the user on the VME bus to the RSF. A busy cell with a destination address that matches the station address is removed by the IOB and placed in the input buffer (if there is room). If the cell is empty, and there is a packet of the same class waiting in the appropriate output buffer, the packet will be transferred to the cell. In most other cases, the cell just passes through.

The IOB clocks the cell out to the MIU, which converts the parallel words back to the serial bitstream. The cell duration in seconds is given by the expression $D = 1024/C$. This bitstream is fed into the optical transmitter (OT), which encodes the clock and data into a 120 Mbit/s optical signal and transmits it to the next station.

III. The WIENER Experimental Environment

The experimental environment is shown in Fig. 3. It consists of a MAGNET II switching node, three independent traffic generators, and a console system. The switching node is configured with a NS and a HS.

The NS is controlled via a transputer-based single board computer [11] called the NSRM. The IOB module in each station on the ring contains three variable counters (LIMIT I, LIMIT II, LIMIT III), which determine the maximum number of packets of each class that a station can transmit during the current RSF subcycle. In addition, the size of the three buffers can be controlled by the NSRM via three counters, B^{I}, B^{II}, and B^{III}. Thus three counters control the size of the three output buffers, respectively. The NSRM sets these variable counters in accordance with the directives of the WIENER subsystem.

The HS is controlled by a transputer-based single board computer called the HSRM. The HSRM sets the cycle lengths via the set of variables MAX I, MAX II and MAX III. Therefore, using the above three variables, the switching bandwidth allocation on the RSF can be controlled.

The three traffic generators are implemented on three separate VME-based transputer boards. A TG generates one class of traffic only. Since each output buffer on the IOB has a unique VME address, each TG will write to one address only. This reduces the contention on the VME bus from access contention and buffer sharing to access contention only. Thus, the VME operations of writing packets to different output buffers can be interleaved, without fear of data corruption.

Network activity at the NS is passively monitored by the HOU. This unit captures incoming and outgoing packet headers, IOB buffer status information, and ring timing information. Details of the HOU will be described in Section III-A below. The software architecture of the experimental environment is described in Section III-B.

A. The Hardware Observation Unit

The HOU packs the monitored information into 128 bit records (I-Records). The records are stored and analyzed by an onboard Transputer. Fig. 4 shows the functional block diagram of the HOU. There are 4 blocks: acquisition unit, processing unit, storage unit, and system unit (the latter unit is not depicted in Fig. 4).

The acquisition unit (AU) captures input signals from the MIU and IOB and packs them into 128 bit record formats (I-Records). It outputs this record to the storage unit.

The storage unit (SU) provides RAM storage for the HOU. Memory is split into two areas: one for the transputer program code and data (transputer RAM), and the other, called the history buffer, for storage of I-Records. This buffer can be read by the processing unit, and written to by the AU. It has the capacity for $M = 33K$ I-Records. The noteworthy aspect of the SU is that the data bus between the AU and the SU is 128 bits wide. Its width allows I-Records to be stored in the history buffer in one

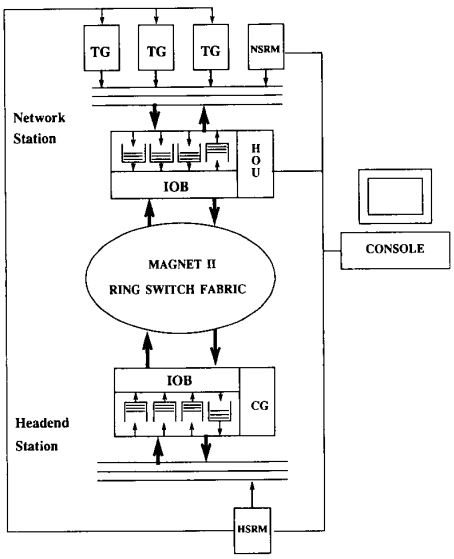

Fig. 3. WIENER experimental environment.

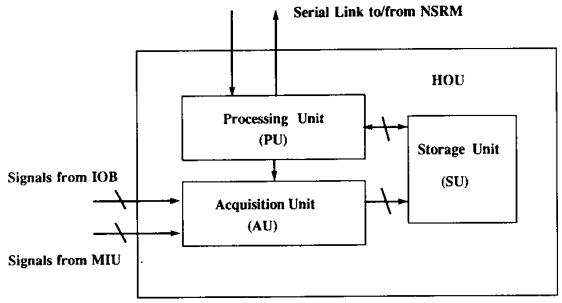

Fig. 4. HOU block diagram.

write cycle. This allows us to scale our design to higher base network speeds when necessary.

The *processing unit (PU)* consists of a 32 bit transputer (T414) that can control the AU, and read the history buffer for purposes of data reduction and analysis. Like all transputers, it communicates with the outside world via a 10 Mbit/s transputer serial link. Code is downloaded over the link when system reset is asserted. Note that the PU can only read the history buffer, it cannot write it.

The *system unit* (SYS) provides miscellaneous system functions to the HOU, including a 10 MHz system clock, and a reset signal.

The HOU is designed to operate in two acquisition modes.

- *Continuous Mode (CM):* The HOU records an I-Record for every cell that comes through the MIU. Each I-Record is numbered with a cell number, ranging from 0–64K. A special record is stored when the cell number wraps around 0. Using this mode, all network activity at the local node can be observed. This makes it easy to determine, for example, the bandwidth available on the RSF to a Network Station.
- *Event Mode (EM):* The HOU stores I-Records only when there is a local event, e.g., the arrival of a packet addressed to the local station, the arrival of a packet from a user to an output buffer, transmission of a packet from the local station, etc. In this way, infrequent events can be recorded over long periods of time.

In both acquisition modes, acquisition can be started either by a direct command from the PU or by the arrival of a control packet solely used for this purpose. This latter method is employed to synchronize data acquisition throughout the network. In both modes, the AU is stopped by a direct command issued by the PU.

To monitor network activity, the HOU is operated in the following way. First the history buffer is reset and the AU is activated. After a user selectable period of time of T seconds (HOU acquisition interval), the AU is stopped and the data is statistically evaluated. It is worth noting that T must be selected carefully to avoid overflow of the history buffer. This selection depends upon the acquisition mode and the offered load. The maximum value $T_{\max}$ for T can be obtained from the following relation:

$$\sum_{j=0}^{\lfloor T_{\max}/D \rfloor} X_j = M$$

where $\lfloor x \rfloor$ denotes the greatest integral value less than or equal to x and X_j is the number of cells per second acquired during the time interval $jD \leq t < (j+1)D$. If the HOU operates in CM, $X_j = 1\ \forall_j, t$, while in EM, $X_j = 1$ only when at least one of the following events occurs in the time interval $jD \leq t < (j+1)D$:

- packet clocked into the IOB input or output buffer,
- packet clocked out of the IOB input or output buffer.

While in all other cases X_j is equal to zero.

B. The Software Monitor

The loading and observation of the network requires a considerable amount of software. There are currently 6 transputers running in a loosely coupled parallel fashion. There are five main software objects which run in parallel: traffic generator (TG), headend station resource manager (HSRM), network station research manager (NSRM), HOU server, and console/experiment manager (console). These programs have been written in OCCAM [10].

The TG is currently capable of generating five different types of traffic sources: L_1, *subband video coder source*; L_2, *Poisson process source*; L_3, *Markov modulated Poisson process source*; L_4, *on-off source with constant arrivals*; L_5, *simulated video coder source*. These are described in more detail in Section IV.

The general method of traffic generation is the same for all five types of traffic sources. A large table of interarrival times ($I = 400\ 000$) is created before the experiment is run. To create traffic, a process just fetches the next interarrival time from the table, and waits on a timer running at a resolution of one microsecond. When the timer expires, a 128 bit packet containing the MAGNET II packet header is sent to the IOB. The IOB hardware gen-

erates a full packet. Thus, the TG is implemented in part in the MAGNET II hardware. This design feature was necessary in order to be able to load the switch. Naturally, the method of building the interarrival time table differs for each type of traffic.

The case of the L_1 source differs from the others in that a schedule is created from the output of a set of simulated subband coders processing 120 s of real-time video sequences. This is based on the work of [5]. The subband coder is based on a constant quality, varying bit rate video coding technique that was proposed for ATM networks. For a video frame of each source, an output summary of each subband is constructed (e.g., the total number of packets per frame). These summaries are kept in a file. To produce a schedule, the user selects the output summary files, the coding rate (e.g., the maximum bit-rate of the coder), the number of traffic sources, and the temporal correlation between these sources. An absolute schedule of packet generation is constructed for each source. These are merged to produce a single interarrival time table for each traffic class. This table is loaded into the appropriate TG. Note that the schedule accurately reflects the composite traffic pattern at the multiplexer of real-time video sources which use the subband coding technique.

In the case of the stochastic sources L_2, L_3, L_4, and L_5, the table is built by computing successive interarrival times (τ_i) drawn from a preassigned distribution. This distribution might be state dependent and/or deterministic.

All of the work of computing the interarrival times at these different rates is performed before the table is filled. By using this method of building the interarrival time table offline, the computation required during real-time traffic generation is kept at a minimum. This allows us to generate pseudorandom traffic with a minimum interarrival time $\tau_{\min} = 10\ \mu$s. When $\tau_i < \tau_{\min}$, the computation loop overhead in the TG dominates the interarrival times. Note that the size of the table is such that the sum of the interarrival times is always much greater than the HOU maximum acquisition interval $T_{\max}$, e.g.,:

$$\sum_{i=0}^{I} \tau_i > \sum_{i=0}^{I} \tau_{\min} > T_{\max} = 327.68 \text{ ms}$$

where $T_{\max}$ has been evaluated, from the previous relation, under the worst case condition given by

$$X_i = 1\ \forall j, t,\quad jD \le t < (j+1)D.$$

We would like to emphasize once more that, although the traffic schedule is computed off-line, the schedule is always played out in real-time, allowing us to perform real-time traffic measurements. A new table is computed with new seeds between each successive run of the experiment.

The TG is also responsible for evaluating the packet gap distribution that it transfers to the console on demand. For each packet that is transmitted, the buffer-full status of the appropriate IOB buffer is checked. If it is full, the packet is blocked. An array of counters is maintained in which the jth counter represents the number of times $j+1$ packets were successively clipped. Once the TG is initialized, it is always actively generating traffic. Note that the network station is not observed until the transient behavior at start-up ceases and the system is in steady state.

The HSRM controls the resource allocation on the RSF. From the console, the user can chose between a multiplicity of variations of fixed and movable boundary schemes in the form of cycle modes. For each scheme there are boundaries which are set by the user at start up. The functionality of the HSRM includes the routing of messages between the TG's and the console.

The NSRM controls the local resources of the network station. This is accomplished by means of setting the limit counters, LIMIT I, II, III, and the buffer size counters B^{I}, B^{II}, and B^{III}. The NSRM supports additional services such as remote configuration of Routers, pinging services, echo services, etc. As shown in Fig. 3, for the measurements presented in this paper only one MAGNET II node was used.

The HOU server controls the mode and sample period of data acquisition. In addition, it computes buffer occupancy, packet delay distributions, packet loss and gap distribution of consecutively lost packets as well as the free bandwidth on the RSF that is available to the NS. Each of these statistical performance evaluations are executed for each class of traffic upon demand. In the future, the HOU server will do traffic estimation, prediction, and classification.

The console provides the human interface to the software monitor. The operator can select programmable parameters for each of the other objects, (although technically speaking, the human interface for each object is part of that object). It also controls the experiment, by sending commands to each object over the serial links. The console is also responsible for some reduction, presentation and storage of experimental data.

IV. Real-Time Traffic Measurements

Using the experimental environment described in Section III, we will first present measurement results of a MAGNET II Network Station loaded with real-time traffic sources. In Section IV-A we introduce, in detail, the traffic sources utilized in the first set of experiments, and the real-time measurements results obtained with this system configuration. In particular, we show the bimodal distribution of the buffer occupancy and time delay distributions that is caused by the intersource correlation of the real-time video sources. The *quantitative* behavior of the different buffer classes is discussed in detail. In Section IV-B, we load the network with other model sources in an attempt to repeat this behavior. The sources used and the corresponding measurement results are presented therein. The emphasis here is on the characterization of the *qualitative* behavior of the buffer classes.

A word about notation. The intensity of the arrival process (if it exists) is denoted by $\lambda(t, i)$ where t stands for

time and i is the traffic source number. Sources L_1 and L_5 do not have an intensity in the classical sense. By abuse of notation, the average rate for these sources will be denoted by $\lambda(t, 1)$ and $\lambda(t, 5)$ (see also [6]). A roman superscript (I, II, or III) denotes the traffic class. Thus, $\lambda^{I}(t, 4)$ is the intensity rate of an L_4 source that carries Class I traffic.

A. Real-Time Traffic Sources and Measurements

The first two types of traffic sources used are as follows.

L_1: Subband Video Coder Source (VC): The simulated video coding subsystem is depicted in Fig. 5. 10 NTSC video sequences containing the head and shoulders of a talking person, captured by a single stationary camera without zoom or pan, were acquired in real time. For each sequence, the luminance signal was extracted, digitized and stored on a disk. Each sequence was processed by a simulated subband video coder. This coder is of the constant quality, varying bit-rate type. Using a combination of subsampling and filtering, a subband video coder splits the digitized three-dimensional video input signal (two-space dimensions, one-time dimension) into eleven subbands. Each subband is coded independently. The separation of image information in subbands results in a vastly different behavior of the output channels. For the class of video sequences considered here, only four of these eleven subbands are important to image quality. These four subbands (with index numbers 1, 5, 8, 9) are coded to produce a varying number of bits for each successive frame. Thus, for each frame, the coder will transmit a varying number of packets for each subband. Due to temporal subsampling, the frame duration time is equal to $1/15$ s [5].

Each subband is treated as an independent traffic source on the network. In terms of picture quality, subband 1 contains the most important visual information, and has the most rigid QOS requirements. Thus, subband 1 requires Class I service. Subbands 5, 8, and 9 have less rigid QOS requirements and, therefore, are transported as Class II traffic.

The packet output rate of the coder was fixed. For example, if a coder operating at rate c were to output N_i packets during the ith frame, those packets would be sent at times $0, d, 2d, \cdots, (N_i - 1)d$ where $d = 1024/c$ is the packet duration time.

Of critical importance is the manner of multiplexing the video sources. As will be shown, the temporal correlation between the frames of different traffic sources results in markedly different traffic statistics. To highlight this fact, four modes of multiplexing were used. In *mode* 0, the beginning time of the frames of all participating sources is deterministically (evenly) distributed over the frame duration time. From the network's point of view, this provides the most regulated load minimizing the aggregate peak rate. In *mode* 1, all frames start in synchrony. This results in a short, high intensity burst at the beginning of each frame period. Since, the aggregate peak rate is maximized, the loading might reach a critical level. In *mode* 2, the beginning time of each frame is a random variable uniformly distributed over the frame interval duration. In *mode* 3, the sources were divided in two subsets. For the first subset, the frame starting time was a uniformly distributed random variable over the frame duration interval, while for the second subset all starting times were in synchrony. Thus, *mode* 3 is a "combination" of *mode* 1 and *mode* 2.

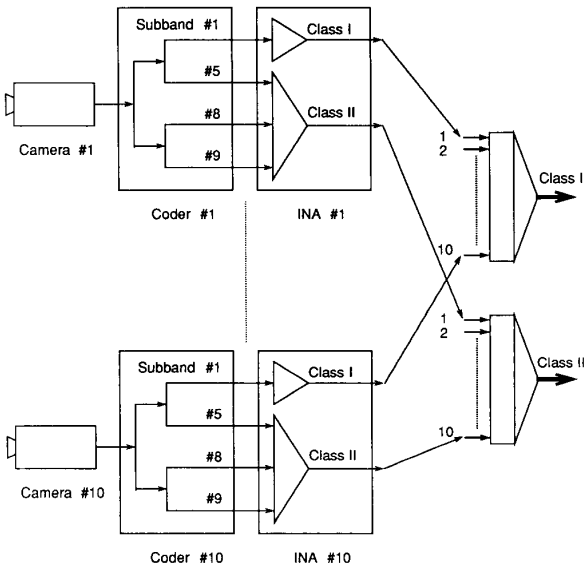

Fig. 5. Subband video coder sources.

L_2: Poisson Process (PP): Each source generates cells according to a Poisson process with average rate $\lambda(t, 2)$.

In all of the experiments, the maximum amount of resources allocated at the RSF to each traffic class was kept constant; in particular an output buffer size of $B^{I}, B^{II}, B^{III} = 15$ cells and a maximum capacity of MAX I = 5, MAX II = 8, and MAX III = 10 cells was adopted. Finally, since only one network station is active, the value of the limit counters for each traffic class were set equal to the maximum cycle length MAX III [14], [15].

The results shown here concern: i) the buffer occupancy distribution $p^{i}(n)$, $0 \leq n \leq 15$ where $p^{i}(n)$ is the probability that at a given time t the Class i buffer, $i = \mathrm{I}, \mathrm{II}, \mathrm{III}$, has n cells waiting for transmission; ii) the time delay distribution $f^{i}(t)$, $0 \leq t < \infty$, experienced by cells inside the Class i buffer; iii) the gap distribution $q^{i}(l)$, $1 \leq l \leq \infty$ where $q^{i}(l)$ is the probability that l cells are consecutively lost due to blocking (Class i buffer full). Along with the above performance parameters, the percentage of totally blocked cells ϵ^{i} experienced by Class i were measured. Finally, with W^{i} we indicate the maximum buffer time delay, i.e., the maximum time spent by a Class i cell inside the corresponding output buffer.

Four different classes of experiments were conducted (see Table I [Fig. 6] for an overview). The plots in Figs. 7(a)–9(a) show the buffer occupancy distribution $p^{i}(n)$ for

TABLE I
TRAFFIC LOAD CHARACTERISTICS FOR EXPERIMENTS 1, 2, 3, AND 4

Exp. #	Class I	Class II	Class III
1	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{I}(t,1) = 29.0Mbit/s$ Subband 1 Multiplexing mode 2	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{II}(t,1) = 10.0Mbit/s$ Subbands 5, 8, 9 Multiplexing mode 2	L_2 Source $\mathbf{E}\lambda^{III}(t,2) = 20Mbit/s$
2	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{I}(t,1) = 29.0Mbit/s$ Subband 1 Multiplexing mode 0	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{II}(t,1) = 10.0Mbit/s$ Subbands 5, 8, 9 Multiplexing mode 0	L_2 Source $\mathbf{E}\lambda^{III}(t,2) = 20Mbit/s$
3	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{I}(t,1) = 29.0Mbit/s$ Subband 1 Multiplexing mode 1	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{II}(t,1) = 10.0Mbit/s$ Subbands 5, 8, 9 Multiplexing mode 1	L_2 Source $\mathbf{E}\lambda^{III}(t,2) = 20Mbit/s$
4	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{I}(t,1) = 29.0Mbit/s$ Subband 1 Multiplexing mode 3	L_1 Source 10 Video Sources $c = 6.0Mbit/s$ $\mathbf{E}\lambda^{II}(t,1) = 10.0Mbit/s$ Subbands 5, 8, 9 Multiplexing mode 3	L_2 Source $\mathbf{E}\lambda^{III}(t,2) = 20Mbit/s$

Fig. 6. Traffic load characteristics for experiments 1, 2, 3, and 4.

Class I, II and III, respectively. In Figs. 7(b)–9(b) the time delay distribution $f^i(t)$ is shown, while Figs. 7(c)–9(c) plot the gap distribution $q^i(l)$ for all three classes, respectively.

In *experiments 1, 2, 3, and 4*, the network was loaded with 10 L_1 video sources. In each experiment, a different multiplexing mode was used. Subbands 1 of each source were multiplexed together resulting in an average aggregate rate of $\mathbb{E}\lambda^{I}(t, 1) = 29$ Mbits/s. Subbands 5, 8, and 9 of the same sources showed an aggregate rate of $\mathbb{E}\lambda^{II}(t, 1) = 10$ Mbits/s. The maximum coding rate c for each video source was fixed at 6 Mbits/s.

Finally, subbands 1 of each source were transported as Class I traffic, Subbands 5, 8, 9 of each source were transported as Class II traffic and a Poisson source L_2 characterized by $\mathbb{E}\lambda^{III}(t, 2) = 20$ Mbits/s was transported as Class III traffic.

Figs. 7–9 show that for *experiments 1 and 2* with multiplexing modes 2 and 0 enforced, respectively, the delay and buffer distributions exhibit desirable unimodal behavior with the $W^{I} = 100\ \mu s$ and $W^{II} = 200\ \mu s$, and $\epsilon^{I} = \epsilon^{II} = 0$.

By contrast, experiments 3 and 4 display markedly different performance, with both delay and buffer occupancy distributions exhibiting bimodal behavior and nonnegligible clipping. The most prominent result is that, although the ratio between the Class III traffic load and the network resources allocated for Class III traffic, i.e.,

$$R^{III} = \frac{\mathbb{E}\lambda^{III}(t, 2)}{C * \dfrac{\text{MAX III} - \text{MAX II}}{\text{MAX III}}} = 1.0$$

is considerably larger than the relative portion of network resources allocated to Class I and II, given by

$$R^{II} = \frac{\mathbb{E}\lambda^{II}(t, 1)}{C * \dfrac{\text{MAX II} - \text{MAX I}}{\text{MAX III}}} = 0.33$$

and

$$R^{I} = \frac{\mathbb{E}\lambda^{I}(t, 1)}{C * \dfrac{\text{MAX I}}{\text{MAX III}}} = 0.59,$$

the delay and buffer distribution for Class III *does not* exhibit bimodal behavior. This is a strong evidence for the argument that the Poisson traffic sources cannot be utilized alone to evaluate the performance of an integrated switching system. This is mainly due to the correlation between composite real-time sources such as video, voice, and facsimile, that cannot be reproduced with the Poisson model. By correlation *between* sources, it is meant that the arrival of a packet at a multiplexer from one source occurs within time ϕ of the arrival of a packet from another source. Note that, in addition, real-time sources ex-

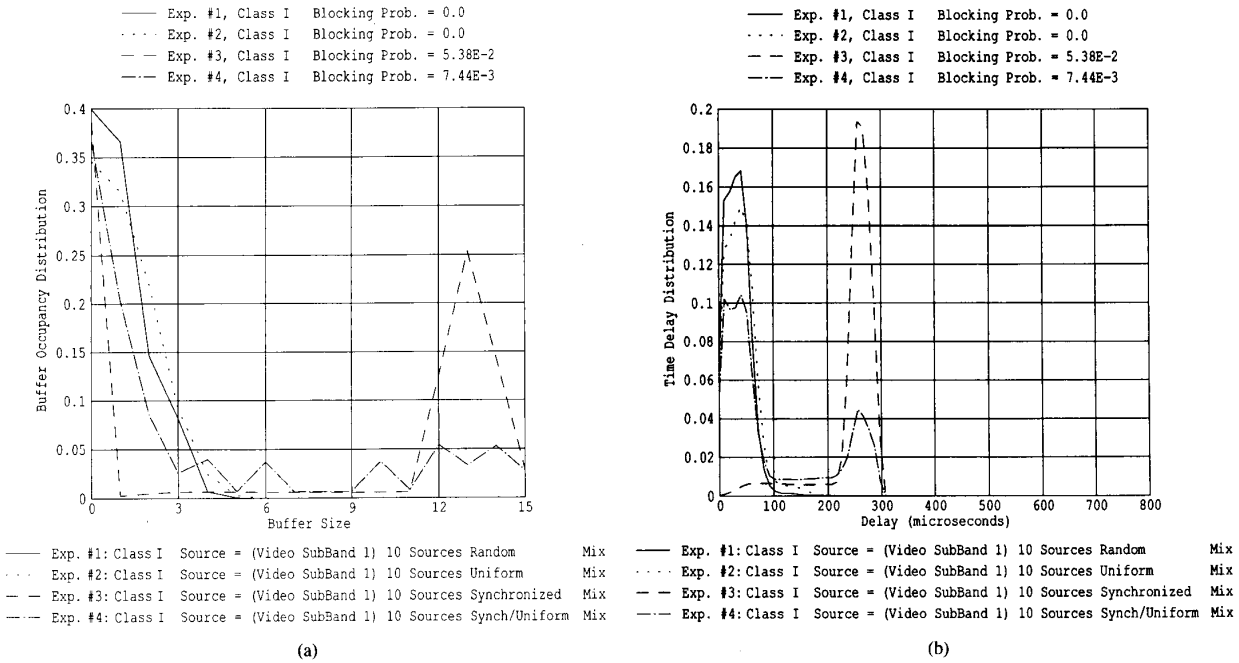

(a) (b)

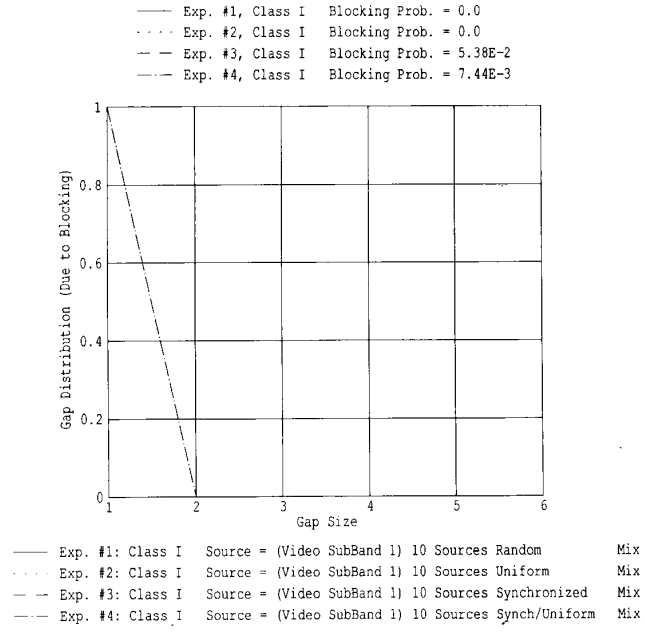

(c)

Fig. 7. (a) Buffer occupancy distribution for Class I traffic. (b) Time delay distribution for Class I traffic. (c) Gap distribution (due to blocking) for Class I traffic.

hibit correlated behavior, in which the arrival time of a packet is dependent on the arrival time of the previous packets; again, this is not reproduced by the Poisson source model.

In the following sections we introduce other traffic source types, with the intent of reproducing the qualitative network behavior shown in Figs. 7(c)–9(c). The results obtained are again compared to the Poisson sources.

B. Traffic Source Models and Measurements

In order to simulate the behavior of a real-time traffic source, the following three models were implemented.

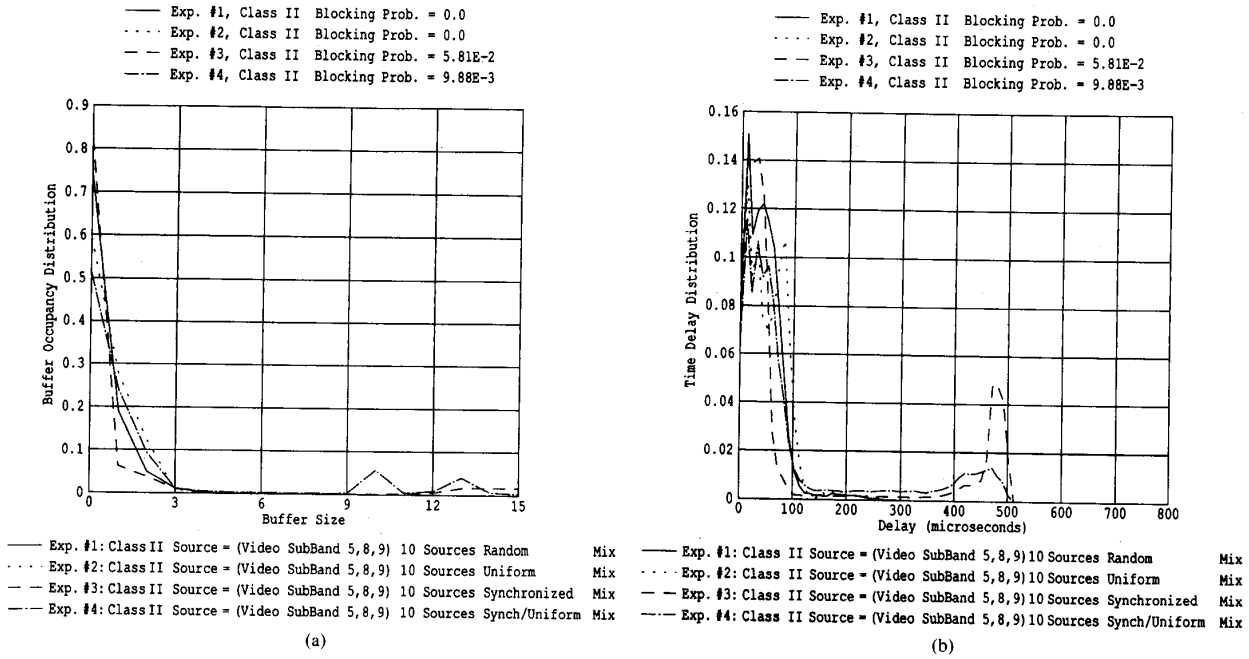

(a) (b)

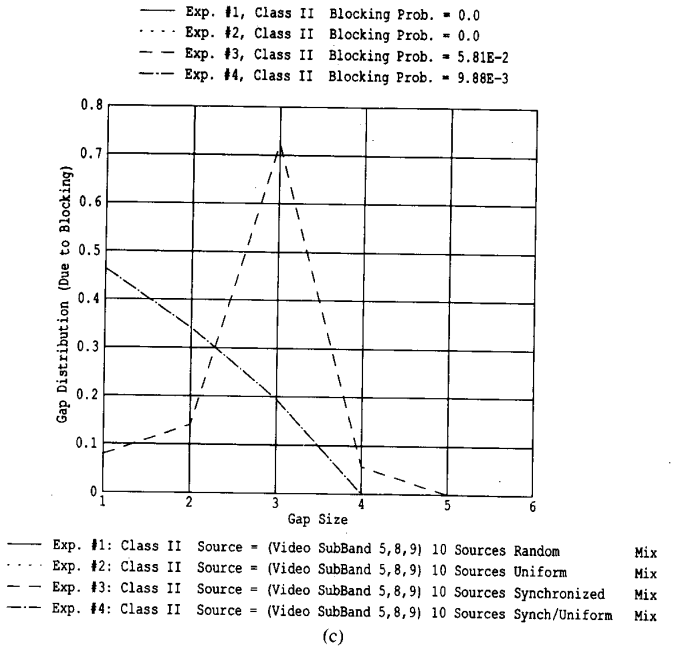

(c)

Fig. 8. (a) Buffer occupancy distribution for Class II traffic. (b) Time delay distribution for Class II traffic. (c) Gap distribution (due to blocking) for Class II traffic.

L_3: Markov Modulated Poisson Process (MMPP): Each source emits cells as a Poisson process with average rate

$$\lambda j = (j + 1) * c$$

where j is determined by the birth-death model with birth rate δ_j and death rate μ_j given by

$$\delta_N = 0, \quad \delta_j = \delta, \quad 0 \le j < N,$$
$$\mu_0 = 0, \quad \mu_j = \mu, \quad 0 < j \le N$$

where $(N + 1)$ denotes the maximum number of levels. The probability of emitting cells at rate λ_j is given by

$$\frac{\rho^j(1 - \rho)}{1 - \rho^{N+1}}$$

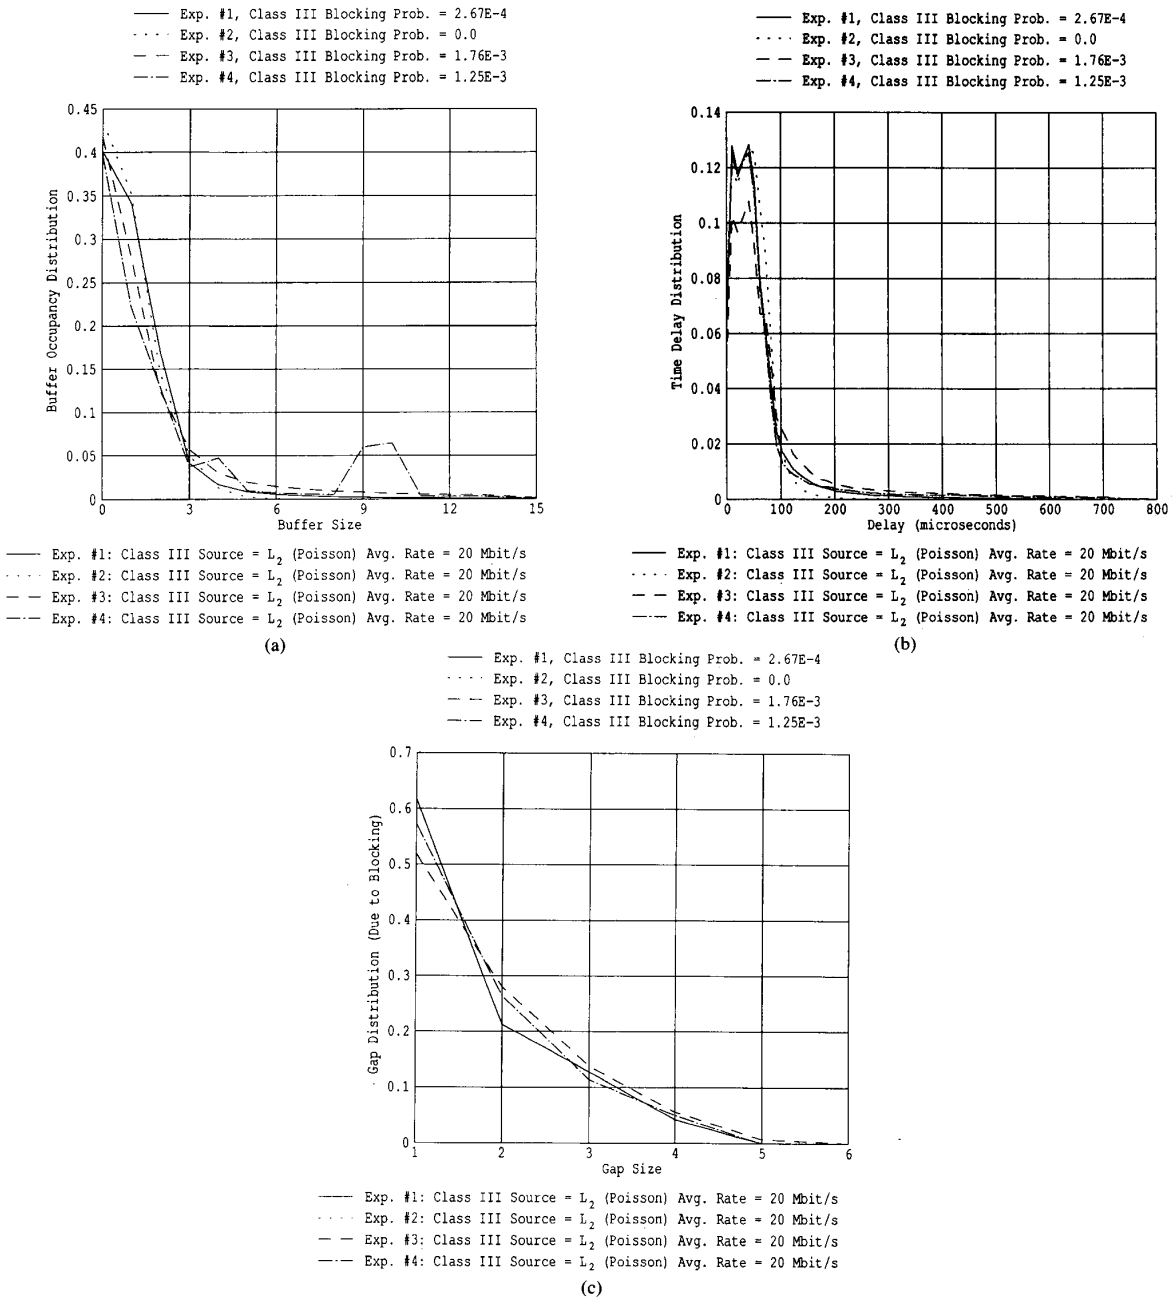

Fig. 9. (a) Buffer occupancy distribution for Class III traffic. (b) Time delay distribution for Class III traffic. (c) Gap distribution (due to blocking) for Class III traffic.

where $\rho = \delta/\mu$. The L_3 source global average rate is given by

$$\mathbb{E}\lambda(t, 3) = c * \left(\frac{\rho}{1 - \rho^{N+1}} * \frac{1 + N\rho^{N+1} - (N+1)\rho^N}{1 - \rho} + 1 \right).$$

L_4: On-Off Source with Constant Arrivals (OO): This source class is characterized by an active period, in which cells are generated with constant rate c, and a silence period, in which no cells are generated. Σ_{on} and Σ_{off} are two statistically independent random variables with a negative exponential distribution describing the on/off behavior of the source. $\mathbb{E}[\Sigma_{on}]$ and $\mathbb{E}[\Sigma_{off}]$ denote the average values of Σ_{on} and Σ_{off}, respectively.

TABLE II
TRAFFIC LOAD CHARACTERISTICS FOR EXPERIMENTS 5, 6, 7, AND 8

Exp #	Class I	Class II	Class III
5	L_2 Source $\mathbb{E}\lambda^I(t,2) = 40Mbit/s$	L_2 Source $\mathbb{E}\lambda^{II}(t,2) = 26Mbit/s$	L_2 Source $\mathbb{E}\lambda^{III}(t,2) = 20Mbit/s$
6	L_3 Source $\lambda^I = 21Mbit/s;\ N = 2$ $\delta^I = 10^3;\ \mu^I = 1.25 * 10^3$ $\mathbb{E}\lambda^I(t,3) = 39Mbit/s$	L_3 Source $\lambda^{II} = 13Mbit/s;\ N = 2$ $\delta^{II} = 10^3;\ \mu^{II} = 1.25 * 10^3$ $\mathbb{E}\lambda^{II}(t,3) = 25Mbit/s$	L_2 Source $\mathbb{E}\lambda^{III}(t,2) = 20Mbit/s$
7	L_4 Source $\mathbb{E}^I[\Sigma_{on}] = 2ms$ $\mathbb{E}^I[\Sigma_{off}] = 0.25ms$ $c^I = 45Mbit/s$ $\mathbb{E}\lambda^I(t,4) = 40Mbit/s$	L_4 Source $\mathbb{E}^{II}[\Sigma_{on}] = 2ms$ $\mathbb{E}^{II}[\Sigma_{off}] = 1ms$ $c^{II} = 39Mbits/s$ $\mathbb{E}\lambda^{II}(t,4) = 26Mbit/s$	L_2 Source $\lambda^{III}(t,2) = 20Mbit/s$
8	L_5 Source $F^I = 0.067s$ $\mathbb{E}^I[\Sigma_{active}] = 0.039s$ $c^I = 50Mbit/s$ $\mathbb{E}\lambda^I(t,5) = 29Mbit/s$	L_5 Source $F^{II} = 0.067s$ $\mathbb{E}^{II}[\Sigma_{active}] = 0.013s$ $c^{II} = 50Mbits/s$ $\mathbb{E}\lambda^{II}(t,5) = 10Mbit/s$	L_2 Source $\mathbb{E}\lambda^{III}(t,2) = 20Mbit/s$

Fig. 10. Traffic load characteristics for experiments 5, 6, 7, and 8.

Finally,

$$\mathbb{E}\lambda(t,4) = \frac{\mathbb{E}[\Sigma_{on}]}{\mathbb{E}[\Sigma_{on}] + \mathbb{E}[\Sigma_{off}]} * c$$

is the L_4 source average rate.

L_5: Simulated Video Coder Source (VS): This source is intended to model the behavior of N subband video coders at the multiplexer. It is a periodic random process that is characterized by a frame duration F cell rate c and average active period, $\mathbb{E}[\Sigma_{active}]$. At time $t = jF$ the cycle begins, and the source emits cells at a rate c. During this time the source is active. The active period, Σ_{active}, is a random variable with a exponential distribution truncated at time $t = F$. At the end of the active period, the source stops emitting cells for the duration of F, e.g., $F - \Sigma_{active}$. The cycle then repeats. The values of F, c, and $\mathbb{E}[\Sigma_{active}]$ are specified at the console. The value of τ is computed at the console by solving the equation:

$$\mathbb{E}[\Sigma_{active}] = \tau[1 - \exp(-F/\tau)]$$

for τ. The L_5 source average rate is given by

$$\mathbb{E}\lambda(t,5) = \frac{\mathbb{E}[\Sigma_{active}]}{F} * c.$$

(The choice of a truncated exponential distribution for the duration of the active period is a simplification that is supported by the easy implementation.)

Four experiments using simulated traffic models (Experiments 5, 6, 7, and 8) were conducted (see Table II [Fig. 10] for an overview). The plots in Figs. 11(a)–13(a) show the buffer occupancy distribution $p^i(n)$ for Class I, II, and III, respectively. In Figs. 11(b)–13(b) the time delay distribution $f^i(t)$ is shown, while Figs. 11(c)–13(c) plot the gap distribution $q^i(l)$ for the three classes, respectively.

In the *fifth experiment*, all three traffic buffers were loaded with an L_2 source model with $\mathbb{E}\lambda^I(t,2) = 40$ Mbits/s, $\mathbb{E}\lambda^{II}(t,2) = 39$ Mbits/s, and $\mathbb{E}\lambda^{III}(t,2) = 20$ Mbits/s.

Figs. 11–13 show that the performance of Class I and Class II traffic are greatly improved with respect to those obtained in experiments 1, 2, 3, and 4. It is important to note that for the measurements presented in this section, the same scheduling policy as in the experiments 1, 2, 3, and 4 was used. Note that, the values of R^I and R^{II} are greater than in the previous experiments. Since for experiment 5 the bimodal behavior is neither present for Class I nor for Class II traffic, the difference in performance is neither due to the network scheduling policy nor to the average load, but solely depends on the different traffic sources utilized.

Fig. 11(a) shows that, under such loading conditions, the probability of finding more than 10 cells inside the Class I output buffer approaches zero. Although R^I is small ($R^I = 0.4$), the stochastic characteristics of the Poisson source together with a finite buffer size leads to a nonzero blocking probability ($\epsilon^I = 5.40 * 10^{-5}$). This behavior is also illustrated in Fig. 11(c) which shows that the maximum number of consecutively lost cells is not greater than 2.

The same results were obtained for Class II and III traffic [Figs. 12(a), 13(a) and Figs. 12(b), 13(b)]. Since $R^{II} = 0.66$ and $R^{III} = 1$ these traffic classes experience a greater buffer maximum time delay ($W^{III} >> W^{II} > W^I$). Furthermore, Figs. 12(c) and 13(c) show that the gap size for Class III cells is greater than the gap size for Class II cells.

In the *sixth experiment*, the Class I buffer was loaded with a L_3 process with $\lambda^I = 21$ Mbits/s, $N = 2$, $\delta^I = 10^3$, and $\mu^I = 1.25 * 10^3$. The Class II buffer was loaded with the same traffic source as Class I with parameter $\lambda^{II} = 13$ Mbits/s, $\delta^{II} = 10^3$, and $\mu^{II} = 1.25 * 10^3$. The Class III buffer was loaded with the same traffic source as in the first experiment. Figs. 11(a)–13(a) along with Figs. 11(b)–13(b) show that the performance for all three traffic classes are basically the same as in the first experiment. Since ρ is equal to 0.8 for both Class I and II traffic sources, it follows that for 75% of the time these sources are emitting at a lower rate than those in the first experiment, while the *average* rates are the same. Thus, the system performance in this case is slightly better than in the previous experiment.

In the *seventh experiment* the average rate for the three classes is again the same as in the previous experiments, while the traffic model for Class I and II has been changed into L_4, with $\mathbb{E}^I[\Sigma_{on}] = \mathbb{E}^{II}[\Sigma_{on}] = 2$ ms, $\mathbb{E}^I[\Sigma_{off}] = 0.25$ ms, $\mathbb{E}^{II}[\Sigma_{off}] = 1$ ms, $c^I = 45$ Mbits/s and $c^{II} = 39$ Mbits/s.

Since the bandwidth allocated to Class I traffic is much larger than c^I, it results that under these conditions block-

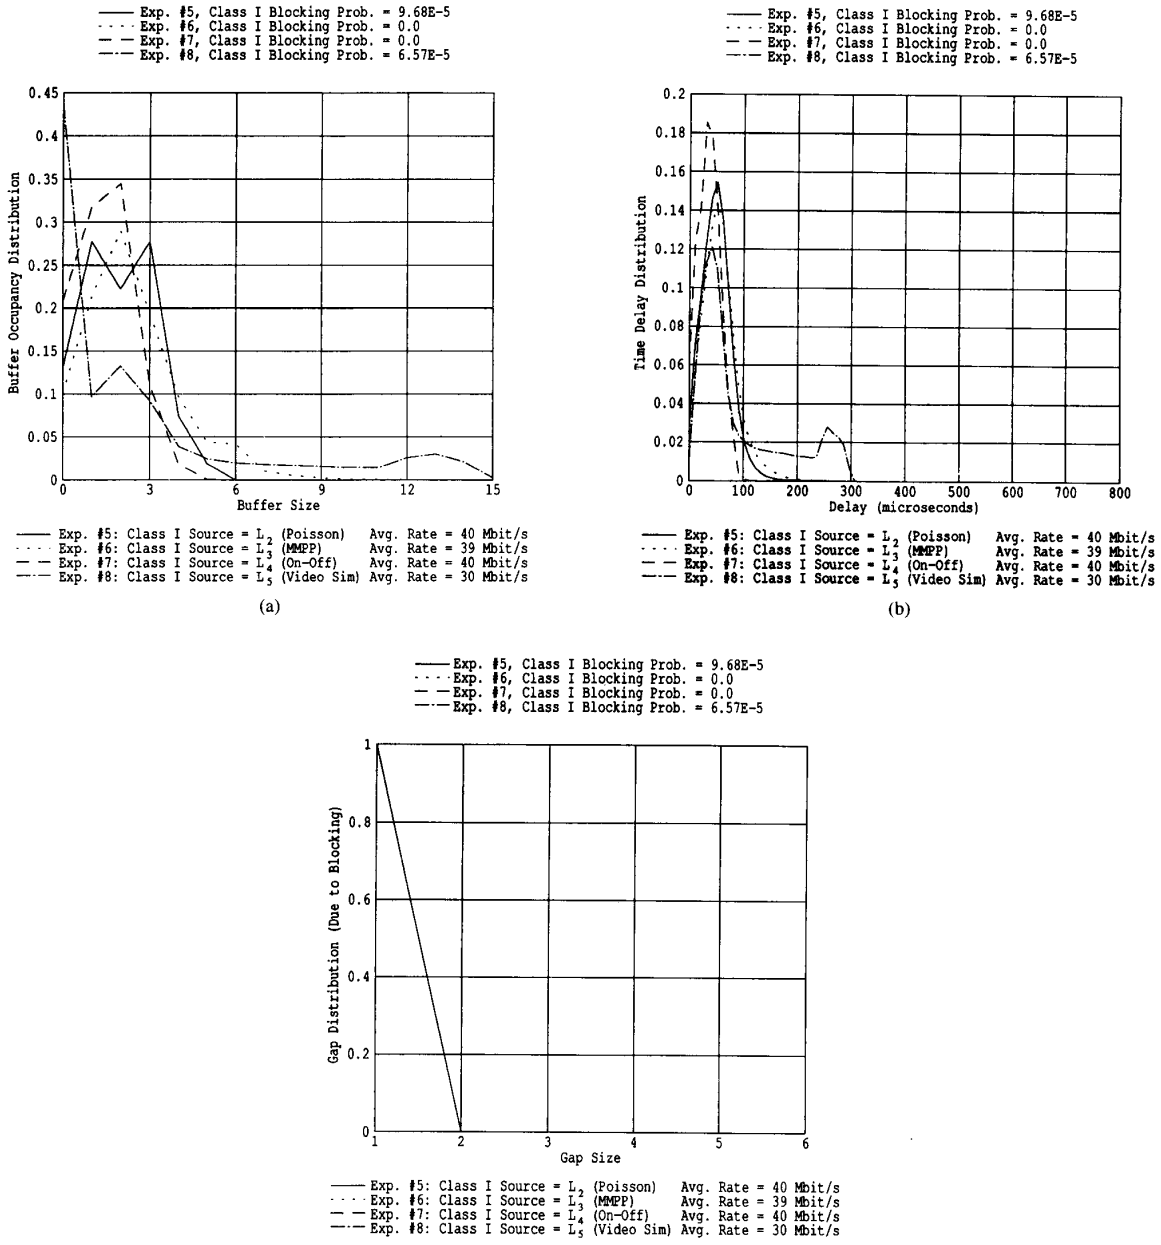

Fig. 11. (a) Buffer occupancy distribution for Class I traffic. (b) Time delay distribution for Class I traffic. (c) Gap distribution (due to blocking) for Class I traffic.

ing does not occur [Fig. 11(c)]. Fig. 11(b) shows that $W^{\mathrm{I}} = 100\ \mu$s. This value is justified by the fact that in the worst case, a cell has to wait only for one cycle before being transmitted (this is the case when the source goes active immediately after the end of subcycle I). Since $c^{\mathrm{I}} = 4.3$ cells per cycle, Class I will utilize either 4 or 5 cells each cycle. Since the subcycle switching overhead [14] is equal to one cell, it results that subcycle I will be always 5 cells long whenever the Class I traffic source is active.

The Class II output buffer was loaded with 3.8 cells per cycle. As mentioned above, MAX II = 8, which means that the Class II traffic source will be able to emit either 3 or 4 cells per cycle whenever the Class I source is inactive, and only 3 cells per cycle whenever the Class I source is active. In the latter case, the Class II buffer will build up at the rate of 4 cells every 5 cycles. If $\Sigma^{\mathrm{I}}_{\mathrm{on}}$ is greater than 20 cycles, blocking will occur with subsequent cell loss [Fig. 12(c)]. Under such conditions, the last cell that entered the Class II buffer has to wait 5 cycles

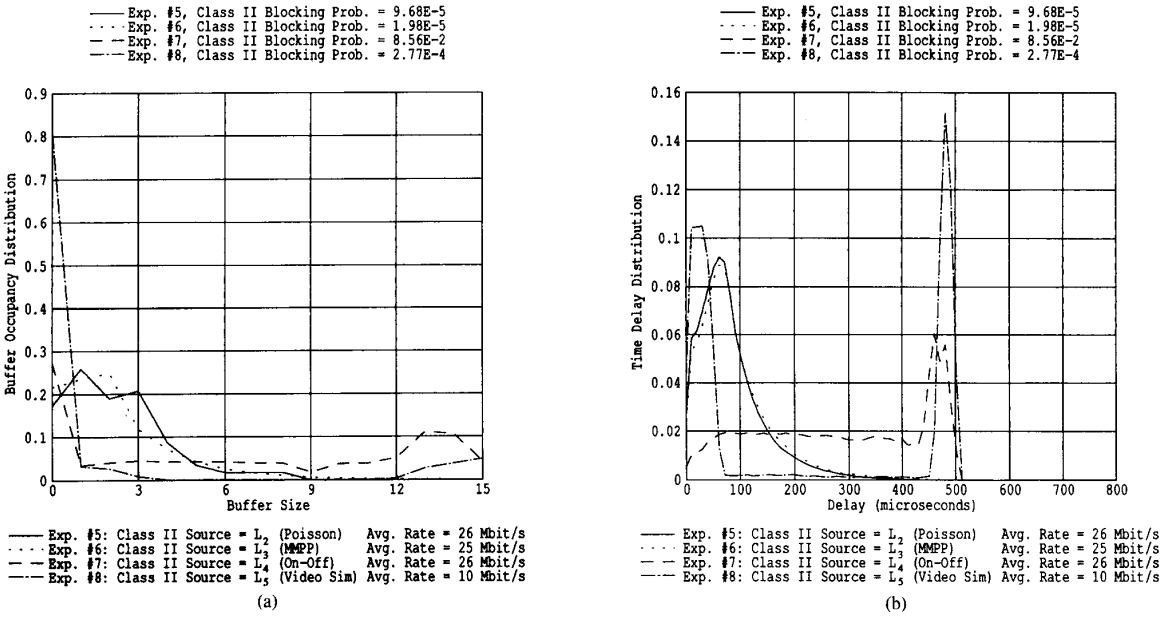

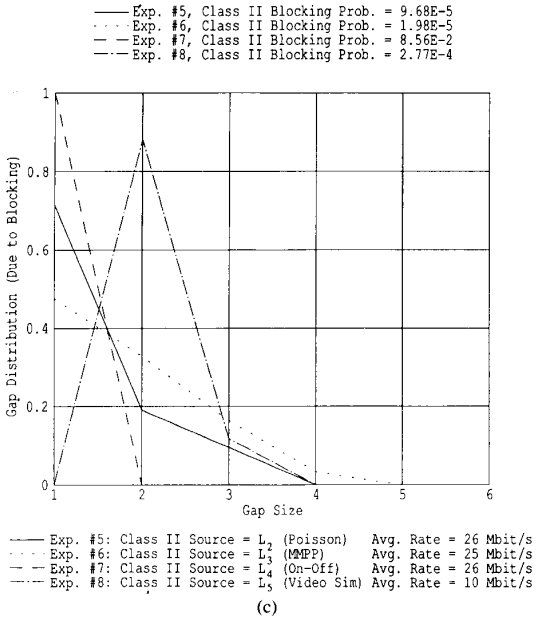

Fig. 12. (a) Buffer occupancy distribution for Class II traffic. (b) Time delay distribution for Class II traffic. (c) Gap distribution (due to blocking) for Class II traffic.

before being transmitted, leading to an W^{II} value of $\simeq 500\ \mu s$ [Fig. 4.8(b)].

Therefore, the Class II buffer will be either almost empty (Class I or II sources inactive) or almost full (Class I and II sources both active) as shown in the plots of Fig. 12(a)–(b). This reproduces the behavior found in the plot of Figs. 8(b) and 9(b). It is clear that the interarrival process of source L_4 is characterized by a high degree of correlation, which suggests that this is the source characteristic that determines the bimodal behavior in the delay and buffer distributions.

Fig. 13(a)–(c) show that Class III traffic appears to be less sensitive with respect to the Class I and II traffic model. *This supports one of the main tenets of the ATS scheme: the separation of the integrated traffic into different traffic classes tends to reduce the transient effects between traffic classes.*

Finally, in the *eighth experiment*, we evaluated the per-

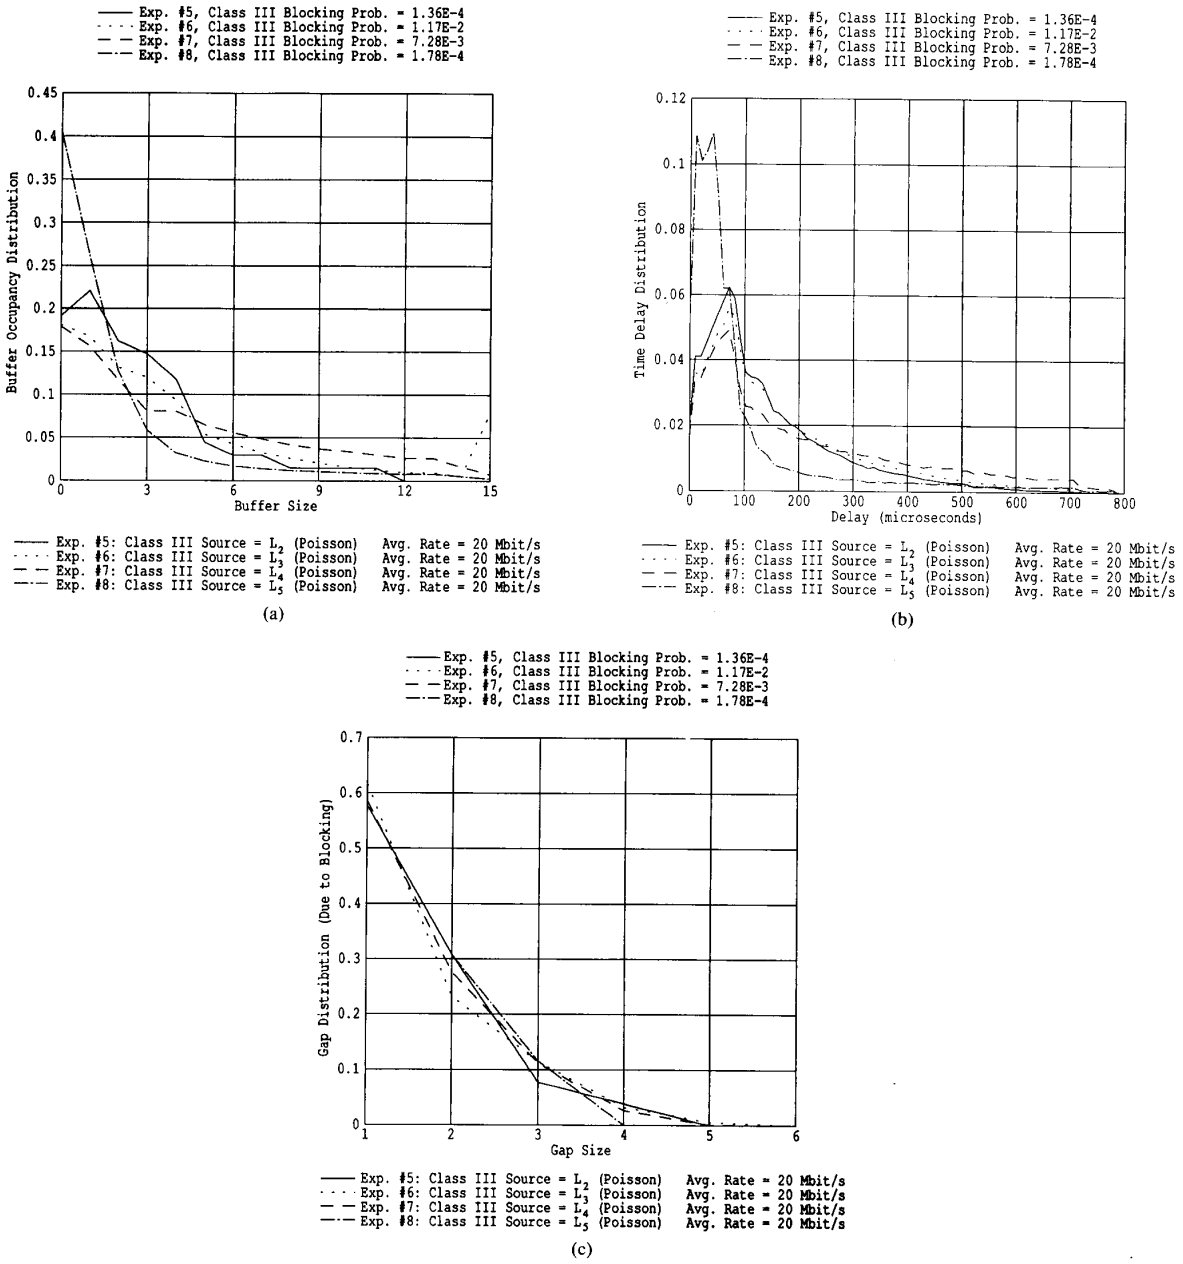

Fig. 13. (a) Buffer occupancy distribution for Class III traffic. (b) Time delay distribution for Class III traffic. (c) Gap distribution (due to blocking) for Class III traffic.

formance of the network when loaded with the L_5 source. The Class I buffer was loaded with an L_5 source operating at a peak rate $c = 45$ Mbits/s, for a frame duration $F = 1/15$ s. The average active period was chosen to be 58% of the frame duration. The Class II buffer was loaded with a source operating at a peak rate $c = 50$ Mbits/s, for the same frame duration $F = 1/15$ s. The average active period was chosen to be 20% of the frame duration and therefore, $\mathbb{E}\lambda^{I}(t, 5) = 29$ Mbits/s and $\mathbb{E}\lambda^{II}(t, 5) = 10$ Mbits/s. Thus, the same average load was imposed as the one generated by the real-time video sources in experiments 1, 2, 3, and 4.

The results are displayed in Figs. 11–13. As expected, both the buffer occupancy and the time delay distribution

of Class II traffic display the bimodal behavior shown by the L_1 source in the previous experiments. As expected, the bimodality is more pronounced than previously seen, since this source represents the worst case for multiplexed periodic sources.

Note that, Fig. 11(a)–(c) also show that, although the average traffic loads for Class I are equivalent for each experiment [i.e., $\mathbb{E}\lambda^{I}(t, 4) \simeq \mathbb{E}\lambda^{I}(t, 3) \simeq \mathbb{E}\lambda^{I}(t, 2)$], the value of the maximum time delay for Class I in the fifth experiment (W^{I}) is twice the value of the maximum delays experienced by Class I in experiments 6 and 7 (W^{I} and W^{II}). The same results apply to Class II [Fig. 12(a)–(c)] where the buffer occupancy and gap distributions obtained in the seventh experiment are drastically different from those obtained in the fifth and sixth experiments. Finally, comparing the results obtained in experiments 7 and 8, we can conclude that, as expected, the worst case is represented by multiplexed input streams with the maximum correlation of the relative phase. This confirms again that *the correlation in the traffic stream has a significant impact on the system performance and must be taken into account, along with the source average load, in developing robust control strategies for integrated networks.*

V. Conclusions

We measured the performance of a MAGNET II network station under different loading conditions. In order to reproduce loading conditions arising in practice, a subband video traffic source and four sources based on different stochastic models were employed. The goal of the experiments was to obtain the *quantitative* behavior of MAGNET II (and its ATS-based concept) as well as a basic *qualitative* understanding of some of the possible congestion problems that might arise in practice.

The *quantitative* results obtained with packet video traffic generated by multiplexed subband video coder sources show that the *a priori* knowledge of the individual traffic sources (characterized by first order statistics) is not adequate for estimating the behavior of a composite source. The phase between the frames of different traffic sources at the multiplexing or switching stage plays an important role in predicting the performance of the network. Our experiments showed that both time delay and buffer occupancy distributions of the multiplexed video sources display a marked bimodal behavior, which does not seem to depend on the buffer size employed. Thus, the stochastic behavior of traffic sources as well as their relative phase has a significant impact on the network performance and must be taken into account when developing robust control strategies for integrated networks.

In order to get a *qualitative* understanding of the bimodal behavior of the time delay and buffer occupancy distribution, four traffic source were implemented. These sources exhibit different degrees of correlation between successive arrivals. These are the *Poisson process source*, the *Markov modulated Poisson process source*, the *on-off source with constant arrivals*, and the *simulated video coder source*. Of the four, the *on-off source with constant arrivals* and the multiplexed *simulated subband video coder source* with suitable chosen parameters best approximated the bimodal behavior of the real-time *subband video coder source*.

In addition, our preliminary experiments provide strong evidence for a basic principle of ATS: *if the network resource is partitioned into separate traffic classes, the performance of one class reduces the effect on the performance of another class*. This is certainly true for the relatively simple case of a MAGNET II network station. Measurements of performance of two and three concurrently active stations are underway. What has yet to be conclusively shown is that the adjustment of resource scheduling and admission control parameters will allow the network to guarantee the QOS negotiated at call set up. Experiments to determine the correct resource scheduling parameters for a given set of traffic sources are being developed at the time of this writing.

We plan to implement more classes of traffic sources such as the embedded model of [7] and add more *real-time* video and voice traffic sources [13], [4]. Further experiments will be conducted in which a large number of such correlated sources are multiplexed in the different modes. In conjunction with this implementation, the problem of the number of moments for the characterization of real-time traffic and the intersource correlation will be raised.

References

[1] P. D. Amer, "A measurement center for the NBS local area computer network," *IEEE Trans. Comput.*, vol. C-31, pp. 723–729, Aug. 1982.

[2] R. S. Bloomfield and K. P. Spies, "User-oriented performance evaluation of data communication services: Measurements results," in *Proc. Int. Conf. Commun.*, Boston, MA, June 11–14, 1989, pp. 3.5.1.–3.5.7.

[3] D. R. Boggs, J. C. Mogul, and C. A. Kent, "Measured capacity of an Ethernet: Myths and reality," in *Proc. ACM SIGCOMM'88 Symp.*, Stanford, CA, Aug. 1988, pp. 222–234.

[4] T. O. Brunner and J. S. White, "Implementation of packet telephone and video services on a local area network," CTR Tech. Rep. 103-88-42, Cent. Telecommun. Res., Columbia Univ., NY, Aug. 1988.

[5] P. Douglas, G. Karlsson, and M. Vetterli, "Statistical analysis of the output rate of a sub-band video coder," in *Proc. SPIE Conf. Visual Commun. Image Processing*, vol. 1001, Cambridge, MA, Nov. 1988, pp. 1011–1025.

[6] J. M. Ferrandiz and A. A. Lazar, "Modeling and analysis of real-time packet traffic," CTR Tech. Rep. 119-88-47, Cent. Telecommun. Res., Columbia Univ., NY, Dec. 1988.

[7] J. Filipiak, "Multi-layer analysis of packet delay and blocking in statistical multiplexing," TRC Tech. Rep. 3/89, The Teletraffic Res. Cent. Univ. Adalaide, Adalaide, Australia, Jan. 1989.

[8] J. Gechter and P. O'Reilly, "Conceptual issues for ATM," *IEEE Network*, vol. 3, pp. 14–16, Jan. 1989.

[9] H. Heffes and D. M. Lucantoni, "A Markov modulated characterization of packet voice and data traffic and related statistical multiplexer performance," *IEEE J. Select. Areas Commun.*, vol. SAC-4, pp. 856–868, 1986.

[10] Inmos Corporation, *Occam Programming Manual.* Englewood Cliffs, NJ: Prentice-Hall, 1984.

[11] Inmos Corporation, *Transputer Development System*, 1988.

[12] L. Kleinrock and W. Naylor, "On measured behavior of the ARPA network," *AFIPS Conf. Proc. 43*, May 1974, pp. 767–780.

[13] A. A. Lazar and J. S. White, "Packetized video on MAGNET," *Opt. Eng.*, vol. 26, no. 7, pp. 596–602, July 1987.

[14] A. A. Lazar, "A traffic control architecture for integrated networks,"

CTR Tech. Rep., Cent. Telecommun. Res., Columbia Univ., NY, in preparation.
[15] A. A. Lazar, A. Temple, and R. Gidron, "A metropolitan area network based on asynchronous time sharing," in *Proc. IEEE Int. Conf. Commun.*, Boston, MA, June 11-14, 1989, pp. 20.3.1-20.3.5.
[16] A. A. Lazar, R. Gidron, and A. Temple, "A switching architecture for asynchronous time sharing," in *Proc. IEEE Global Telecommun. Conf.*, Dallas, TX, Nov. 27-30, 1989, pp. 1166-1172.
[17] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J. D. Robbins, "Performance models of statistical multiplexing in packet video communications," *IEEE Trans. Commun.*, vol. 36, pp. 834-844, July 1988.
[18] S. Mazumdar and A. A. Lazar, "Knowledge-based monitoring of integrated networks," in *Proc. First Int. Symp. Integrated Network Management*, Boston, MA, May 14-17, 1989, pp. 235-243.
[19] ——, "Monitoring integrated networks for performance management," in *Proc. IEEE Int. Conf. Commun.*, Atlanta, GA, Apr. 15-19, 1990.
[20] V. Ramaswami, "Traffic performance modeling for packet communication—whence, where, and whither," keynote address given at the *Third Australian Teletraffic Seminar*, Nov. 1988.
[21] M. El Zarki, A. A. Lazar, A. Patir, and T. Takahashi, "Performance evaluation of MAGNET protocols," in *Local Area & Multiple Access Networks*, R. L. Pickholtz, Ed. Rockville, MD: Computer Science, 1986, pp. 137-154.

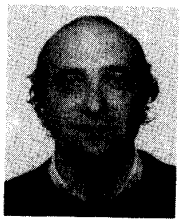

Aurel A. Lazar (S'77-M'80) was born in Zalau, Transylvania, Romania, on January 30, 1950. He received the Dipl.-Ing. degree in communications engineering (Nachrichtentechnik) from the Technische Hochschule Darmstadt, Darmstadt, Federal Republic of Germany, in 1976, and the Ph.D. degree in information sciences and systems from Princeton University, Princeton, NJ, in 1980.

In 1980 he joined the faculty in the Department of Electrical Engineering, Columbia University, New York, as an Assistant Professor. Since 1988 he has been a Professor and Director of the Telecommunication Networks Laboratory. His current areas of interest include control and management of telecommunication networks and, the mathematics of networks and intelligent systems.

Dr. Lazar is a founding member of the Center for Telecommunications Research at Columbia University and a member of ACM.

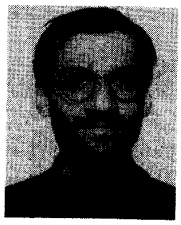

Giovanni Pacifici (S'80-M'85) was born in Rome, Italy, on September 27, 1957. He received the Dr.Eng. and Ph.D. degrees from Department of Information and Communication Technology at the University of Rome *La Sapienza* in 1984 and 1989, respectively.

As a student, his main activities were focused on the performance evaluation of local and metropolitan area networks, with an emphasis on the integration of voice and data. In the course of his studies, he was a Visiting Scholar at the Center for Telecommunications Research, Columbia University, from 1987-1988. In 1989, he joined the staff of the Center for Telecommunications Research as a Research Scientist. His current interests include monitoring and control of Broadband Integrated Networks, real-time traffic generation, parallel processing, and switching architectures.

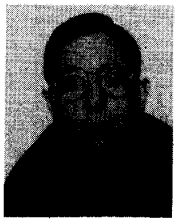

John S. White (M'84) was born in St. James, NY, on June 9, 1953. He received the B.S. and M.A. (mathematics) degrees from SUNY, Stony Brook, in 1976 and 1978, respectively. In 1987, he received the M.S.E.E. degree from Columbia University.

From 1979 to 1984 he has worked in industry. Since 1984, he has been a research staff member at the Center for Telecommunications Research, working in the areas of packet video, packet voice, observation, control, and management of broadband integrated networks. His current interests include network observation and control, and parallel processing.

Mr. White is a member of the Association for Computing Machinery.