Report Number: WUCS-91-12

1991-02-01

Performance Analysis of the Ethernet under Conditions of Bursty Traffic

Tony Y. Mazraani and Gurudatta M. Parulkar

In this paper we present a simulation study of the Ethernet performance under conditions of bursty traffic. This study is motivated by two observations: Ethernet will continue to be a widely used Local Area Network (LAN), especially as an access LAN for future high speed internet (or Broadband ISDN); and future high speed applications can best be modeled as bursty sources. Bursty traffic in this study is specified using three parameters: peak bandwidth, average bandwidth, and burst factor. The simulation study shows that the inherent behavior of the Ethernet does not change with bursty traffic. That is, as long... Read complete abstract on page 2.

Performance Analysis of the Ethernet under Conditions of Bursty Traffic

Tony Y. Mazraani and Gurudatta M. Parulkar

Complete Abstract:

In this paper we present a simulation study of the Ethernet performance under conditions of bursty traffic. This study is motivated by two observations: Ethernet will continue to be a widely used Local Area Network (LAN), especially as an access LAN for future high speed internet (or Broadband ISDN); and future high speed applications can best be modeled as bursty sources. Bursty traffic in this study is specified using three parameters: peak bandwidth, average bandwidth, and burst factor. The simulation study shows that the inherent behavior of the Ethernet does not change with bursty traffic. That is, as long as the utilization is less than a threshold value, packet delay, is almost equal to transmission time, queue lengths are minimal, and packet delay, queue lengths, and packet loss rate increase very quickly. Although the basic trend of the Ethernet performance is the same, performance metrics deteriorate faster with bursty traffic. For example, packet loss due to collision, packet delay, and buffer sizes increase with burstiness of traffic sources. The ratio of peak to average bandwidth of traffic sources has an unexpected effect on the packet loss rate and queue lengths. At high utilization, packet loss and queue lengths are less for higher peak-to-average ratio of bursty sources.

Performance Analysis of the Ethernet under Conditions of Bursty Traffic

Tony Y. Mazraani
Gurudatta M. Parulkar

WUCS-91-12

February 1991

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

Performance Analysis of the Ethernet under Conditions of Bursty Traffic

Tony Y. Mazraani[1]
tony.y.mazraani@sprintintl.sprint.com
+1 703 689 6619

US Sprint
Mail Stop: VARESB123B
12490 Sunrise Valley Drive
Reston, Virginia 22096

Gurudatta M. Parulkar
guru@flora.wustl.edu
+1 314 889 4621

Computer & Comm. Research Ctr.
Department of Computer Science
Washington University
Saint Louis, Missouri 63130

ABSTRACT

In this paper we present a simulation study of the Ethernet performance under conditions of bursty traffic. This study is motivated by two observations: Ethernet will continue to be a widely used Local Area Network (LAN), especially as an access LAN for future high speed internet (or Broadband ISDN); and future high speed applications can best be modeled as bursty sources. Bursty traffic in this study is specified using three parameters: peak bandwidth, average bandwidth, and burst factor. The simulation study shows that the inherent behavior of the Ethernet does not change with bursty traffic. That is, as long as the utilization is less than a threshold value, packet delay is almost equal to transmission time, queue lengths are minimal, and packet loss due to collisions is zero. However, as the utilization increases beyond the threshold, packet delay, queue lengths, and packet loss rate increase very quickly.

Although the basic trend of the Ethernet performance is the same, performance metrics deteriorate faster with bursty traffic. For example, packet loss due to collision, packet delay, and buffer sizes increase with burstiness of traffic sources. The ratio of peak to average bandwidths of traffic sources has an unexpected effect on the packet loss rate and queue lengths. At high utilization, packet loss and queue lengths are less for higher peak-to-average ratio of bursty sources.

1. INTRODUCTION

In this paper we present a performance analysis of the Ethernet under conditions of bursty traffic. This analysis is motivated by the following two step reasoning.

First, high speed networking or Broadband ISDN (BISDN) technology has significantly matured, and its commercial viability and feasibility has become increasingly evident over the past few years. In fact, a number of vendors are planning to deploy pilot broadband networks in the near future. We believe that at least in the first phase of their deployment, broadband networks will be used as backbone networks at campus, regional, and national level, and the access networks will continue in most cases to be LANs such as Ethernet, token ring, and FDDI.

Second, BISDN deployment will also bring with it a whole set of new applications, such as scientific imaging and visualization, multi-media conferencing, video distribution, and others. These

[1] This work was done while the author was with the Computer and Communications Research Center at Washington University in St. Louis. This work was supported by the National Science Foundation, Bellcore, Bell Northern Research, Italtel, and NEC.

applications have very different traffic output from that of traditional applications, and require performance guarantees in terms of throughput, delay, and packet loss. We claim that most of the broadband applications need to be modeled as bursty sources, and therefore, previous analyses of the Ethernet with Poisson sources cannot be used to accurately characterize the performance obtained by these new applications. In this paper we report results of our simulation study aimed at characterizing the Ethernet performance for broadband applications.

The simulation study shows that the inherent behavior of the Ethernet does not change with bursty traffic, which is an expected result. That is, as long as the utilization is less than a threshold value, packet delay is almost the same as the transmission time, queue lengths are minimal, and packet loss due to collisions is zero. However, as the utilization increases beyond the threshold, packet delay, queue lengths, and packet loss rate increase very quickly.

However, the important conclusions of this simulation study are that the Ethernet performance deteriorates faster with the *burstiness* of traffic sources, and that bursty sources are more demanding than Poisson. For example, packet loss due to collisions, packet delay, and buffer sizes increase with the *burstiness* of traffic sources. The ratio of peak to average bandwidths of traffic sources has an unexpected effect on the packet loss rate and queue lengths. At high utilization, packet loss and queue lengths are less for higher peak to average ratio of bursty sources.

This paper is organized as follows. Section 2 presents an overview of the Ethernet network topology, frame format, and access protocol. Section 3 summarizes previous measurements and simulation studies. The Ethernet simulation model is presented in Section 4. This model uses traffic sources that are capable of generating both Poisson and bursty traffic. A description of these sources is presented in Section 5. The simulation results are discussed in Sections 6 and 7. As a validation step, Section 6 presents results obtained under conditions of Poisson traffic, and a comparison of these results with related work from the literature. Section 7 presents results obtained under conditions of bursty traffic, and examines the effect of varying various traffic parameters on the Ethernet performance. Finally, the conclusions of this paper are presented in Section 8.

2. BACKGROUND

Ethernet is a broadcast local area network with a branching-bus topology and uses an unslotted, 1-persistent, carrier sense multiple access protocol with collision detection (CSMA/CD) and binary exponential backoff. In this section we present an overview of the Ethernet network topology, frame format, and access protocol.

2.1. Ethernet Topology and Frame Format

The current generation Ethernet consists of a coaxial cable, or multiple segments of coaxial cable connected by repeaters, with a data rate of 10 Mbps. An Ethernet can have a maximum length of 2.5 km connected in 500-m segments, and up to 1024 connected stations. Stations attach to the cable using a tap, with the distance between any two taps being a multiple of 2.5 m. Each tap includes a transceiver, which transmits and receives signals on the cable and also does carrier sensing and collision detection.

The frame format for an Ethernet packet is shown in Figure 1. Each frame starts with a 7-byte *preamble*, a synchronization pattern consisting of alternating ones and zeros which allow the receiver's clock to synchronize with the transmitter's clock. Next, the start of the frame is indicated by a 1-byte *start-frame delimiter*. The *destination* and *source addresses* are then specified; each consisting of a 6-byte field. The source address is an unique address of the transmitter. The destination address, however, corresponds to either a unique station or a group of stations (in this case the destination address is often called a *multicast* address). The *length* field is two bytes long, and indicates the number of information bytes in the frame. The *information field* has a length between 46 and 1500 bytes. If the information supplied by the upper layer is less than 46 bytes, the pad field is used to bring the length to the minimum of 46 bytes. The final field is the *frame-check sequence* which consists of a 4-byte cyclic-redundancy-check code (CRC) for error detection. Therefore, the maximum and minimum packet sizes of the Ethernet are 1526 bytes (which include 1500 information bytes) and 72 bytes (which include 46 information bytes), respectively.

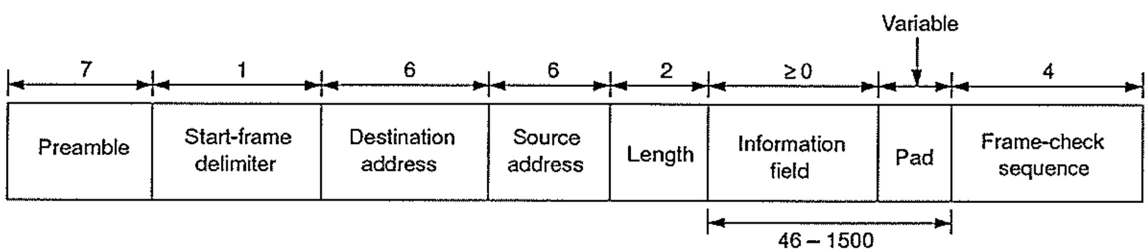

Figure 1: Frame format for an Ethernet packet

2.2. Ethernet Access Protocol

Ethernet uses an unslotted, 1-persistent, CSMA/CD protocol with binary exponential backoff. When a station is ready to transmit, it first senses the channel. If the channel is idle the station transmits its packet. If the channel is sensed busy, which indicates that another transmission is in progress, the station continues the sensing function until the channel is idle; then it transmits immediately. This access scheme, however, does not prevent packet collisions. Due to channel propagation delays, multiple stations may sense the channel idle and then transmit at nearly the same time causing a collision among all transmissions. Collision detection is carried out by the transmitting station which compares the transmitted signal with the received signal. When a mismatch (*i.e.* collision) occurs, the station involved immediately stops transmitting its packet, sends a jamming signal, and then waits for a random amount of *backoff* time before it attempts to retransmit the packet. The jamming signal ensures that all stations on the channel know of the collision and backoff.

Ethernet uses a truncated binary exponential backoff algorithm. In this algorithm, if a packet has been transmitted unsuccessfully n times, the next transmission is delayed by a random amount of time equal to the base backoff time (51.2 μs) times a random integer i. The integer i is generated from a uniform distribution in the range $0 \leq i \leq 2^k$ where $k = \min(n, 10)$. In the Ethernet implementation, each packet is allowed a maximum of 16 attempts. If all 16 attempts fail the packet is discarded, and the Ethernet retransmission algorithm reports an error.

For further details on the Ethernet, the reader is referred to the literature [1, 2].

3. PREVIOUS WORK

Performance studies of the Ethernet fall under two categories: measurements and simulation. Analytical models of the actual Ethernet are not available mainly because of the difficulties in modeling both the backoff algorithm and the physical distribution of stations. However, the unslotted and slotted 1-persistent CSMA/CD models from the literature have been used to approximate the performance of the Ethernet in parts of its operating region.

The commonly-used performance measures are *average delay, throughput*, and *channel capacity.* Average delay indicates the time to send a packet, and is measured from the time the station generates the packet until the packet is completely received. Throughput is equal to the fraction of the maximum channel bandwidth that is used for carrying useful data, which usually includes packet headers as well as the information field. The last measure is the channel capacity which is the maximum achievable throughput given a set of parameters, such as packet length and propagation delay.

3.1. Measurement Studies

Gonsalves, 1987

Gonsalves [3] measured both throughput and mean delay as a function of offered load on an operational 10 Mbps Ethernet. The experimental environment consisted of three 500 meter segments connected in series by two repeaters, and 30--38 stations were used to generate the traffic. Each station had one packet buffer, and generated packets with random interarrival times. In each station, a new packet was generated only after the previous packet was successfully transmitted. Fixed length packets were used, and packets discarded due to excessive collisions were not included in the mean delay computations. The one-way propagation delay was estimated to be 30 ms. Gonsalves demonstrated that the Ethernet protocol achieves high throughput of 80 percent with long (1500 bytes) packets, and low throughput of 25 percent with short (64 bytes) packets. Furthermore, he showed that for an offered load below the channel capacity, the mean delay for all packet lengths was roughly equal to the packet transmission time. However, at high loads the mean delay increased rapidly to several times the packet transmission time.

Boggs, Mogul, and Kent, 1988

Boggs *et al* [4] measured data rate, transmission delay, and fairness all as a function of the number of stations of an Ethernet implementation. Fairness in this study was based on the standard deviation of throughput, computed from the throughput measures of individual stations. The measurement environment consisted of an equal number of stations connected to one of four multiport repeaters, whose transceivers were connected to a 910-meter long Ethernet with an inter-transceiver distance of roughly 303 meters. All stations but one were configured as packet generators and collectors of local statistics. One station acted as the test controller, broadcasting control packets announcing the beginning of a test and the length of the packets to be generated. The test controller also collected all statistics from individual stations at the end of a test to compute the final performance measures. When a test was started, each station waited for a short period of time, and then transmitted packets of the length specified by the control station for a predetermined period of time.

Using this measurement environment, Boggs *et al* conducted three experiments. In the first experiment, they used 24 stations to generate fixed length (in the range of 64 bytes – 4000 bytes) packets for each run on the experimental Ethernet. They showed that throughput increases with increasing packet length, and that the channel's fairness increases as the number of stations (*i.e.* the offered load) increases. The second experiment consisted of 23 stations sending fixed length packets on a 6 meter Ethernet which can be assumed to have zero propagation delay. As expected, the Ethernet performance improved significantly because the collision resolution time was negligible. The most interesting result in this experiment was that a throughput of 85 percent was achieved using 64 bytes packets. In the third experiment, all previous measurements were repeated using several bimodal packet length distributions, which were composed of 64 bytes and 1536 bytes packets. The results showed that the Ethernet performance improves as the ratio of the number of long packets to the number of short packets increases.

3.2. Simulation

O'Reilly, 1983

O'Reilly [5] analyzed the performance of the 2.94 Mbps Ethernet using simulation. Because the analysis results are normalized to the Ethernet data rate, they can be extrapolated to now standard 10 Mbps Ethernet without much error. The simulation assumes that the stations are equally spaced along the Ethernet and that the bus length is 550 meters. The jam time and base backoff time are 3 μs and 38 μs, respectively. The number of stations on the bus was varied from 5 to 15 while each station generated 10% of the total Ethernet data rate. The packet arrival process was Poisson and packet lengths were 512, 1024, or 4096 bytes. The simulation results show that under conditions of 200% offered load, the throughput varied from 0.851 for 512 byte packets to 0.976 for 4096 byte packets. The corresponding mean delay values for successfully transmitted packets were 2.055 msec and 9.597 msec, respectively.

Hughes and Li, 1982

Hughes and Li [6] used discrete-event simulation to determine the performance of a typical Ethernet. The bus bandwidth was 10 Mbps and the worst case end-to-end delay was 22.5 μs. Values for jam time, interframe gap, and base backoff time were 3.2 μs, 9.6 μs, and 51.2 μs, respectively. The packet arrival process was Poisson, and packet lengths were distributed as follows: 20% were 4000 bit packets and 80% were 250 bit packets. Simulation results show that the channel capacity in this case is only 63% of the maximum bus bandwidth, and that the average packet delay increases quickly when the offered load increases beyond 50%.

It is important to note that none of these previous studies have considered the Ethernet performance under conditions of bursty traffic resulting from real-time applications such as video, facsimile, voice, and data. Therefore, it is difficult to use these studies to predict the performance that can be provided to bursty applications over the Ethernet in an emerging broadband communication environment. In this paper, we present a simulation model for the standard Ethernet with sources capable of generating Poisson as well as bursty traffic. As a verification of the simulator, we first present simulation results of the Ethernet using Poisson traffic, and show that these results match the results reported in the literature.

4. SIMULATION MODEL

The simulation model simulates the operation of the standard Ethernet at the level of the data link layer. Some features of the physical layer, such as cable propagation and repeater delay are also included. The network consists of a single channel and a set of stations. Each station has an independent packet generator and a buffer of infinite capacity as shown in Figure 2. Hence, packet loss in this model occurs only when the Ethernet protocol discards a packets after 16 unsuccessful retransmission attempts[2]. In each station, the first packet in the queue is completely transmitted before a second packet, if it exists, is dequeued. A *balanced star* topology is assumed, *i.e.* the propagation delay between any two stations on the channel is constant. The simulator uses discrete-event simulation techniques, and is written in C and runs under the Unix operating system.

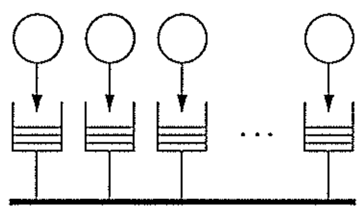

Figure 2: Network Configuration

4.1. Simulation Operation

A high-level block diagram of the simulation model is shown in Figure 3. A simulation run starts with the user specifying the input parameters (Section 4.4) which describe the network topology and the traffic characteristics. The simulator runs for a specified amount of time, and then prints out the output statistics (Section 4.5) which consist of performance measurements and debug information (if requested).

The simulator is divided into four modules, the *input parser, packet generator, network simulator,* and *output analyzer.* The input parser reads input parameters, and checks for errors. If no errors are detected, it computes various parameters for packet generator, does appropriate unit transformations, sets up basic data structures and event queue, and then initializes variables used in the simulation. All events in the event queue are ordered by time stamp with the lowest at the head of the

[2] However, we do monitor the packet queue of each station, and have found that the average and maximum queue lengths remain reasonable. See Section 7 for details.

queue. The packet generator can be considered as a module which is running in the background (Section 5) and generating packets with a specified arrival process. The packet generator gets the traffic characteristics from the input parser and generates packets with distribution specified by these characteristics. The network simulator, the major component of the simulation model, simulates the Ethernet access protocol and collects various statistics. The general operation of the network simulator follows the classic *next-event time advance discrete-event* simulation. In this approach, the network simulator enters the following loop: it removes the first event (*i.e.* the event with the lowest time stamp) from the event queue, advances the simulation clock to the time of occurrence of that event, performs the required operation for this event, updates the state of the system, and then inserts new events in the event queue. This process is continued until a prespecified stopping condition is satisfied. At this point, the output analyzer takes control. It computes the final statistics from the statistics collected by the network simulator, and then prints out a summary of the results.

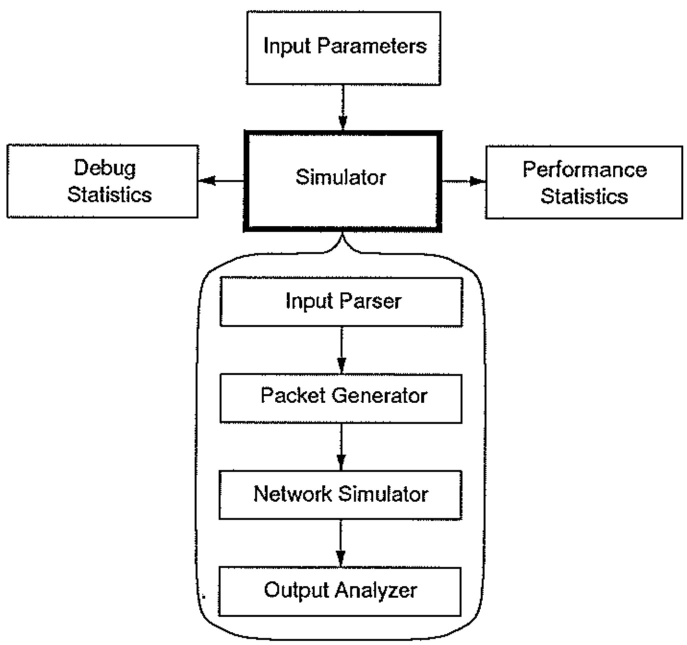

Figure 3: Block diagram of the simulation model

4.2. Input Parameters

The simulation input parameters are shown in Table 1. They are chosen to simulate the 10 Mbps Ethernet implementation whose characteristics are given in [1]. The channel propagation delay is equivalent to about 5 km of coaxial cable (the maximum allowed channel length) and is used to examine the worst-case performance of the Ethernet. All remaining parameters are specified by the user, and are as follows:

- number of stations on the channel,
- offered load (in Mbps),

- packet length (in bytes),
- simulation length (in seconds),
- and packet generator parameters: peak bandwidth (in Mbps), average bandwidth (in Mbps), and burst factor (in packets).

Parameter Description	Notation	Value
base backoff time	t_{bbt}	51.2 μs
collision detection time	t_{cd}	4.8 μs
interpacket gap	t_{gap}	9.6 μs
length of jamming signal	tjam	4.8 μs
channel propagation delay	τ_{pd}	23.2 μs

Table 1: Simulation input parameters

4.3. Output Statistics

The output of the simulator consists of two parts: *simulation trace* and *output statistics*. When the debug option is asserted, every time the simulation state changes the simulator prints out a trace showing the up-to-date state of the system and various performance statistics. This option is very useful in debugging the simulator and verifying the output statistics, however they are beyond the scope of this paper. The following is a description of the output statistics of the simulation.

- *Offered load* is the sum of all individual offered loads of stations on the channel normalized to 10 Mbps, the maximum capacity of the Ethernet. The offered load of an individual station is defined as the number of bits per second that the station attempts to transmit on the channel.
- *Channel throughput* is the total number of useful bits per second transmitted on the channel, normalized to 10 Mbps. Useful bits include all headers, trailers, and data fields of successfully transmitted packets.
- *Mean packet delay* in milliseconds is computed from the individual delays of all packets transmitted (including those discarded). The delay experienced by a given packet is computed from the time the packet is assembled until the time the packet is completely transmitted or discarded due to repeated collisions.
- *Maximum packet delay* in milliseconds is the mean of all maximum packet delays computed per station. The per-station maximum packet delay is computed from transmitted as well as discarded packet.
- *Mean queue length* in kbytes is computed using a trajectory of the total number of packets in the system as a function of time. This trajectory changes every time there is a packet arrival or a packet departure from the system. At the end of the simulation, the area under the trajectory is computed, and then divided by the simulation time times the number of stations on the channel.

- *Maximum queue length* in kbytes is the mean of the per station maximum queue lengths.
- *Packet loss* is computed from the number of packets discarded due to excessive collisions. Packet loss is expressed as the ratio of the number of packets discarded divided by the sum of the number of packets transmitted and the number of packets discarded.

5. SOURCE MODEL

In a broadband communication environment, an application can be best modeled as a source of bursty traffic [7, 8], which consists of bursts of data, possibly at maximum application data rate, followed by idle periods. In such environment, an application has to specify its bandwidth requirements using more parameters than just its peak bandwidth. Reserving peak bandwidth for each application is obviously inefficient, and thus other parameters such as average bandwidth and burst size are needed to give a more accurate estimate of bandwidth and buffer requirements.

In this paper, a bursty source is modeled as a *Markov chain* consisting of two states, *active* and *idle*, as shown in Figure 4(a). When the source is in the active state, it generates packets at a rate λ_p. In the idle state the source is shut off. The holding times in both the active and the idle states are exponentially distributed with means $1/\beta$ and $1/\alpha$, respectively.

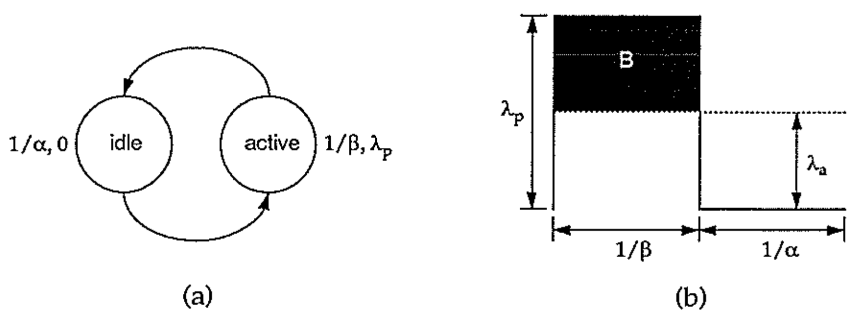

Figure 4: Model of a bursty source

In this section, we are interested in expressing α and β in terms of three parameters: the peak packet arrival rate (λ_p), the average packet arrival rate (λ_a), and the burst factor (B). In the simulator, we specify the traffic characteristics using *peak bandwidth, average bandwidth,* and *burst factor.* Therefore, we will also express α and β in terms of these three parameters.

We first define the burst factor as follows

$$B \equiv (\text{mean time in active state}) \times (\text{peak rate} - \text{average rate}),$$

which is equivalent to

$$B = \frac{1}{\beta}\,(\lambda_p - \lambda_a) \tag{1}$$

In other words, the burst factor represents the average number of packets which are generated at a rate higher than the average, as shown in Figure 4(b). The average packet arrival rate as a function of α, β, and λ_p is given by:

$$\lambda_a = \frac{\alpha\lambda_p}{\alpha+\beta} \tag{2}$$

Substituting for λ_a from (2) into (1), we obtain:

$$B = \frac{\lambda_p}{\alpha+\beta} \tag{3}$$

Using (2) and (3), we can now express α and β in terms of λ_p, λ_a, and B as follows:

$$\alpha = \frac{\lambda_a}{B} \tag{4}$$

$$\beta = \frac{\lambda_p - \lambda_a}{B} \tag{5}$$

Let L be the packet length in bytes, B_p be the peak bandwidth in Mbps, and B_a be the average bandwidth in Mbps. Equations (4) and (5) can be written in terms of B_p, B_a, and B as follows:

$$\alpha = \frac{B_a}{8LB} \tag{6}$$

$$\beta = \frac{B_p - B_a}{8LB} \tag{7}$$

Note that for Poisson traffic, the mean holding times in the active and idle states of the packet generator are not specified. The packet generator is configured to remain in the active state for the duration of the experiment (*i.e.* $1/\alpha = 0$ and $1/\beta = \infty$). The only remaining packet generation parameter is λ_p, the mean packet arrival rate per station. We compute λ_p by first computing the total packet arrival rate from the offered load and packet length, and then dividing the result by the number of stations on the channel.

6. SIMULATION VALIDATION

In order to validate our simulation, we compared results of our simulation with Poisson traffic with those reported in the literature. The comparison showed that our results match other results very well as shown in Figure 5. The results were obtained using 30 stations and fixed length packets. The interarrival times of packets at each station are independent and exponentially distributed.

A study of Figure 5 shows that our simulation results are consistent with the measurements conducted by Gonsalves [3] and simulation results presented by O'Reilly [5]. Given the assumption of balanced star topology, the results presented in this chapter evaluate the worst-case performance of the Ethernet because the channel propagation times are on the order of 20 µs – 30 µs, which correspond to relatively large cable lengths. The Ethernet is capable of performing better for smaller cable length, *i.e.* smaller channel propagation and collision resolution times. Boggs *et al* [4] have shown using an experiment on a 6 meter long Ethernet that a throughput of 0.85 can be achieved with 64 byte packets.

A number of other results were obtained with Poisson traffic using our simulation model and compared with results in the literature to conclude that our results are consistent and simulation model valid. Refer to [9] for more details.

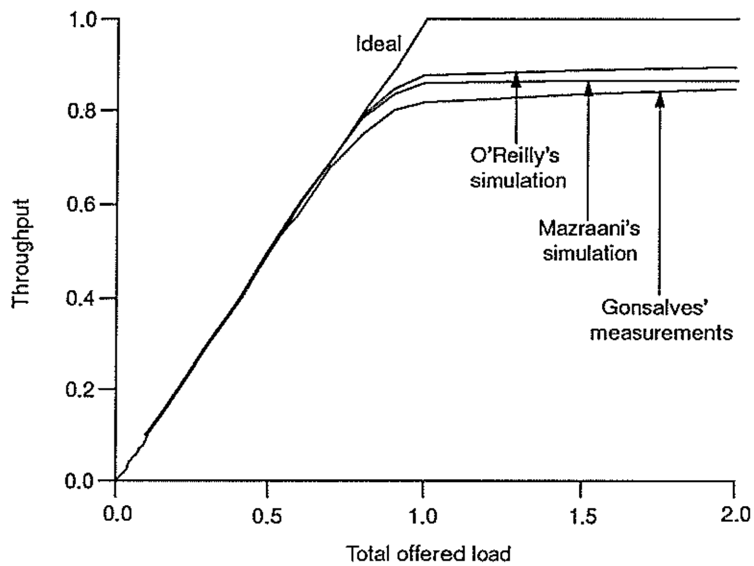

Figure 5: Throughput vs. offered load— measurements and simulation

7. SIMULATION RESULTS WITH BURSTY SOURCES

We now present simulation results for the Ethernet performance with bursty traffic. The input parameters and the output statistics are presented in Sections 4.4 and 4.5. The offered load in this case is defined in terms of three parameters: peak bandwidth in Mbps, average bandwidth in Mbps, and burst factor in packets.

We analyze the Ethernet performance by examining the effect of traffic parameters and number of traffic sources on the following performance measures: packet delay, packet loss, and buffer requirements. All results presented in this section were obtained using 30 stations, 1500 byte packets, and exponential interarrival times of packets when the generator is active, unless otherwise specified. Simulation results using short packets are not presented here because they exhibit similar behavior as results with long packets.

7.1. Effect of Varying Burst Factor on the Ethernet Performance

To examine the effect of burstiness on the Ethernet performance, we fix the peak (λ_p) and average (λ_a) packet arrival rates, and vary the burst factor. As a result, the packet generator stays in both active and idle states for longer time. This is depicted in Figure 6 in which we assume deterministic holding times in active and idle states. In Figure 6(a) we show a source with $\lambda_p = 2\lambda_a$. We fix λ_p and λ_a, and we double the burst factor as shown in Figure 6(b). As a result, the holding times in active and idle states are doubled.

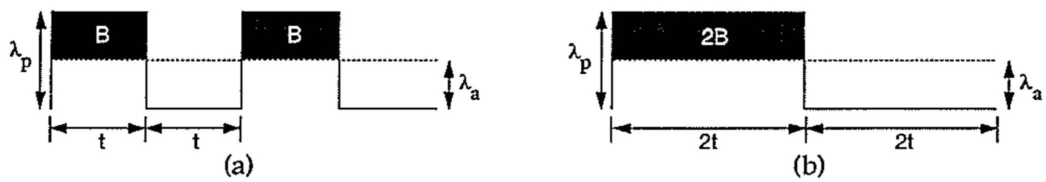

Figure 6: Traffic characteristics when the burst factor is doubled

Effect of Varying Burst Factor on Packet Delay

The mean packet delay as function of throughput and burst factor is shown in Figure 7. The peak-to-average ratio is equal to 4. Note that bursty traffic with zero burst factor is equivalent to Poisson traffic. For low throughput values, the contention on the channel is low for various burst factors. As a result, the mean packet delay is almost equal to the packet transmission time. For medium throughput values, the mean packet delay increases with increasing burst factor. This is due to the fact that for a fixed throughput value and an increasing burst factor, the number of packets contending for the channel increases. As a result, the number of collisions on the channel

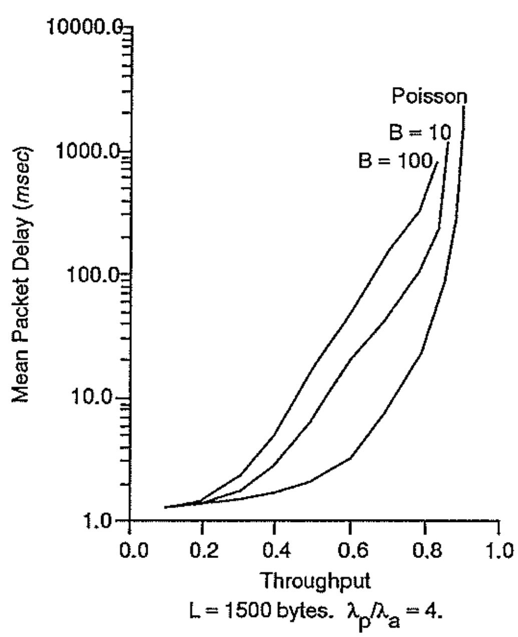

Figure 7: Mean packet delay vs. throughput and burst factor

increases, and consequently the queueing delay experienced by packets increases. As the throughput reaches the channel capacity, the difference in delay performance for different burst factors becomes more significant. Note that all curves converge together when throughput approaches channel capacity.

Effect of Varying Burst Factor on Packet Loss

The effect of burstiness is more pronounced on packet loss, as shown in Figure 8 which depicts packet loss as a function of throughput and burst factor for a peak-to-average ratio of 4. For low throughput values, the contention for the channel is minimal (see Figure 7). As a result, packet loss is zero irrespective of burst factor. However, as both throughput and burst factor increase, packets experience more collisions, and consequently the number of packets dropped by the Ethernet protocol increases. As a result, packet loss becomes more significant for higher burst factors. An interesting result to note here is that for high throughput values, the difference in packet loss for different burst factors is small, whereas the difference in packet delay is significant (see Figure 7). This is due to the exponential backoff mechanism of the Ethernet which in this case acts as a trade-off between delay and packet loss. At high load, burst collisions occur more frequently, which causes bursts involved in the collision to be delayed for a long time.

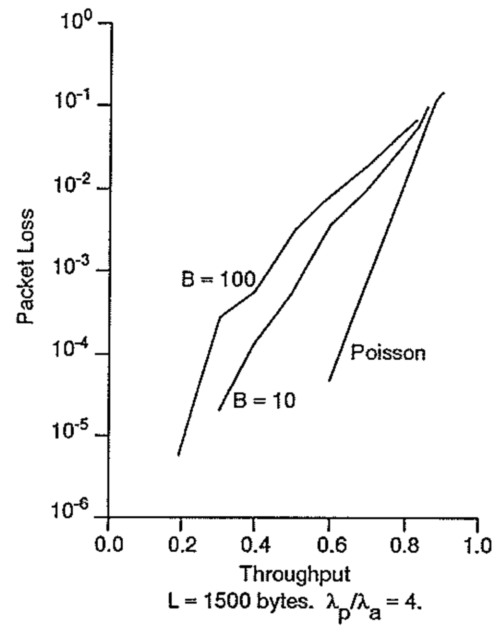

Figure 8: Packet loss vs. throughput and burst factor

Effect of Varying Burst Factor on Buffer Requirements

The mean queue length per station as a function of throughput and burst factor for a peak-to-average ratio of 4 is depicted in Figure 9. The curves in this Figure show that the mean queue length is higher for larger burst factors. However, the queue length remains relatively stable until the throughput gets close to the channel capacity, at which point the queue length increases abruptly to many times the packet length.

In order to determine the buffer requirements per station, we examine the maximum queue length per station, which is shown in Figure 10 as a function of channel throughput. As expected, the maximum queue length per station increases with the traffic burstiness. Figure 10 shows that to keep the buffer size per station less than a few packets (say ≤ 50 packets), the Ethernet should always be under-utilized. For instance, when the burst factor is 100, the Ethernet utilization should be kept below 4 Mbps in order to guarantee a maximum queue length of less than 50 packets.

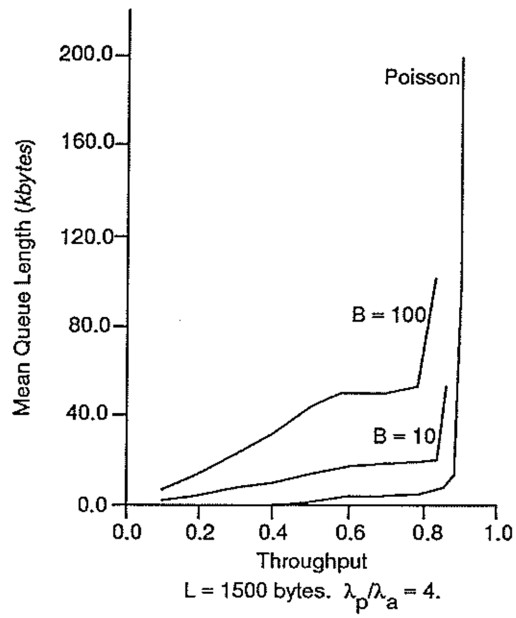

Figure 9: Mean queue length per station vs. throughput and burst factor

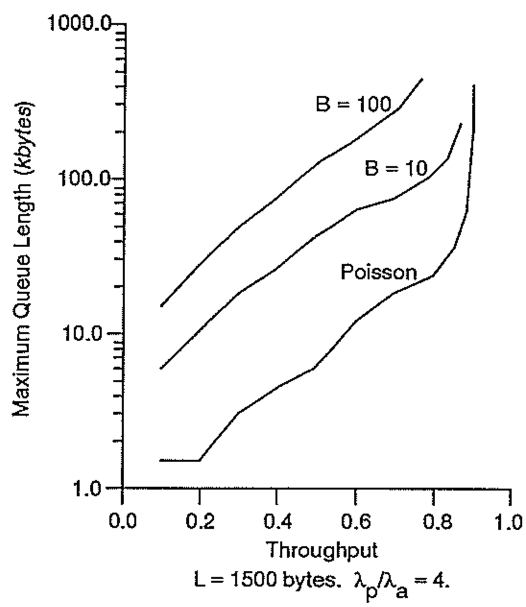

Figure 10: Maximum queue length per station vs. throughput and burst factor

Summary

The results presented in this section clearly show that the Ethernet performance under conditions of bursty traffic is different from the performance obtained when using traditional analyses with Poisson traffic. The simulation results show that the performance of the Ethernet deteriorates as the burst factor increases. For instance, assuming a peak-to-average ratio of 4 and 1500 bytes packets, Figure 7 shows that in order to guarantee less than 10 msec mean packet delay, the Ethernet utilization has to be considerably lower than the channel capacity of 90%. Under conditions of Poisson traffic, the Ethernet utilization should be lower than 72%. However, under conditions of bursty traffic, this value deteriorates quickly to 57% and 43% for burst factors of 10 and 100, respectively. Bursty traffic has a more dramatic effect on packet loss. Figure 8 shows that in order to guarantee less than 10^{-4} packet loss under conditions of Poisson traffic, the Ethernet utilization should always be less than 62%. Under conditions of bursty traffic, this value decreases quickly to 39% and 27% for burst factors of 10 and 100, respectively.

7.2. Effect of Varying Peak-to-Average Ratio on the Ethernet Performance

In the following set of experiments, we examine the effect of the peak-to-average ratio on the Ethernet performance. We do that by keeping the burst factor and the average packet arrival rate fixed, and varying the peak packet arrival rate. The resulting changes in the traffic characteristics are illustrated in Figure 11, which shows that as the peak-to-average ratio increases, the holding time in the active state decreases to keep the burst factor constant, and the holding time in the idle state remains constant to keep the average packet arrival rate constant.

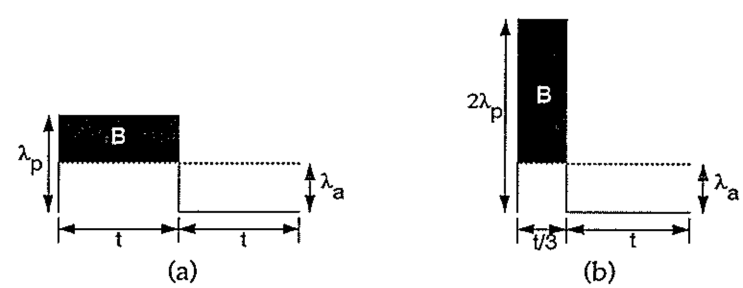

Figure 11: Traffic characteristics when the peak-to-average ratio is doubled

Effect of Varying Peak-to-Average Ratio on Packet Delay

The mean packet delay as a function of throughput and peak-to-average ratio for a burst factor of 100 is shown in Figure 12. A study of the curves in this Figure shows that the mean packet delay increases with increasing peak-to-average ratio, and that the increase in delay is significant in the case of high peak-to-average ratios, even for very low throughput values. When the peak-to-average ratio increases, the number of packets generated per active unit time also increases. As a result, the contention on the channel becomes higher, and equivalently the average delay experienced by packets increases. Since the maximum attainable throughput is constrained by the channel capacity, all delay curves converge together as throughput approaches channel capacity.

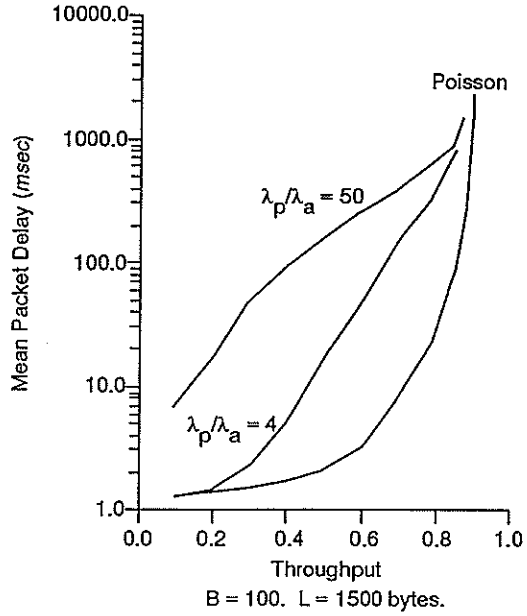

Figure 12: Mean packet delay vs. throughput and peak-to-average ratio

Effect of Varying Peak-to-Average Ratio on Packet Loss

Packet loss as a function of throughput and peak-to-average ratio is shown in Figures 13. The burst factor is equal to 100. A study of the curves of Figure 13 shows that packet loss is higher for large peak-to-average ratio and low throughput (< 0.6). When throughput approaches channel capacity, all curves converge together. The interesting result to note here is the behavior of packet loss curves for high throughput values (> 0.6) for which packet loss is lower for high peak-to-average ratio. In order to justify such unexpected result, the behavior of bursty sources and channel contention is examined as the peak-to-average ratio and throughput change.

Figure 11 shows that as the peak-to-average ratio increases for fixed average bandwidth, both the mean holding time in the active state and the source cycle time decrease. This can also be seen in Figure 14 which shows that the mean holding time in the active state ($1/\beta$) for a peak-to-average ratio of 50 is lower than that for a peak-to-average ratio of 4. What is important to note here is that for high throughput values, bursty sources are characterized by a very short active state and a relatively long idle state. The active state is even shorter when the peak-to-average ratio is high, as shown in Figure 14 for a peak-to-average ratio of 50. These observations imply that when both the peak-to-average ratio and throughput are high, large numbers of packets within a burst contend for the channel during a very short period of time. In the event of a collision, bursts involved in the collision get delayed for some period of time. However, since the idle state in this case is much longer than the active state, delayed bursts have a better chance of capturing the channel. As a result, when throughput is high there will be fewer collisions for a high peak-to-average ratio than for a low peak-to-average ratio. However, when throughput is low and the active state is long enough compared to the idle state, we should obtain more collisions for higher peak-to-average ratio. This phenomenon is depicted in Figure 15 which consists of histograms of the number of collisions per source cycle time for two throughput values and for peak-to-average ratios of 4 and 50. They show that for a throughput of 0.2, which is below the throughput value that corresponds to the crosspoint in Figure 13, packets corresponding to the peak-to-average ratio of 50 experience more collisions. This behavior is reversed for a throughput value of 0.8, which is above the throughput value corresponding to the crosspoints.

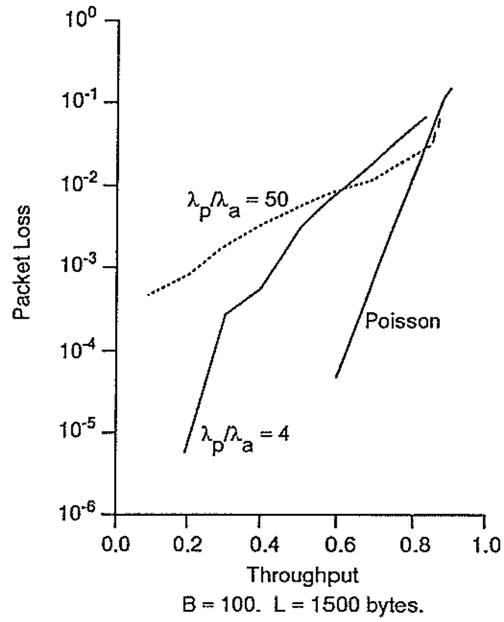

Figure 13: Packet loss vs. throughput and peak-to-average ratio

The buffer requirements in this study have similar trends as in Figures 9 and 10. The mean queue length per station under conditions of bursty traffic is larger than what is predicted under conditions of Poisson traffic. In all cases, however, the buffer requirements are adequate for the current memory technology.

The same set of experiments were run, but using only 5 stations. The simulation results are consistent with the results obtained using 30 stations. That is, as the peak-to-average ratio increases, Ethernet performance deteriorates, but only at low throughput values. When throughput is high, the Ethernet performs better for higher peak-to-average ratio.

7.3. Effect of Varying Load with Fixed Source Cycle on the Ethernet Performance

So far, we have examined the effect of varying the burst factor and the peak-to-average ratio on the Ethernet performance. In both cases, we have noticed that as the load on the channel varies, the source cycle time also varies (see Figures 6 and 11). In this section, we keep the source cycle time fixed, and examine the effect of varying the load on the Ethernet performance. In this case, as the load varies, we are in effect varying all three parameters: peak bandwidth, average bandwidth, and burst factor. All simulation results presented here were obtained using 1500 bytes packets.

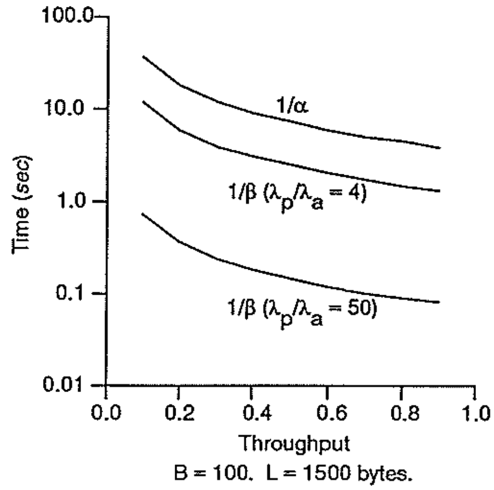

Figure 14: Mean holding times in the active and idle states as a function of throughput and peak-to-average ratio

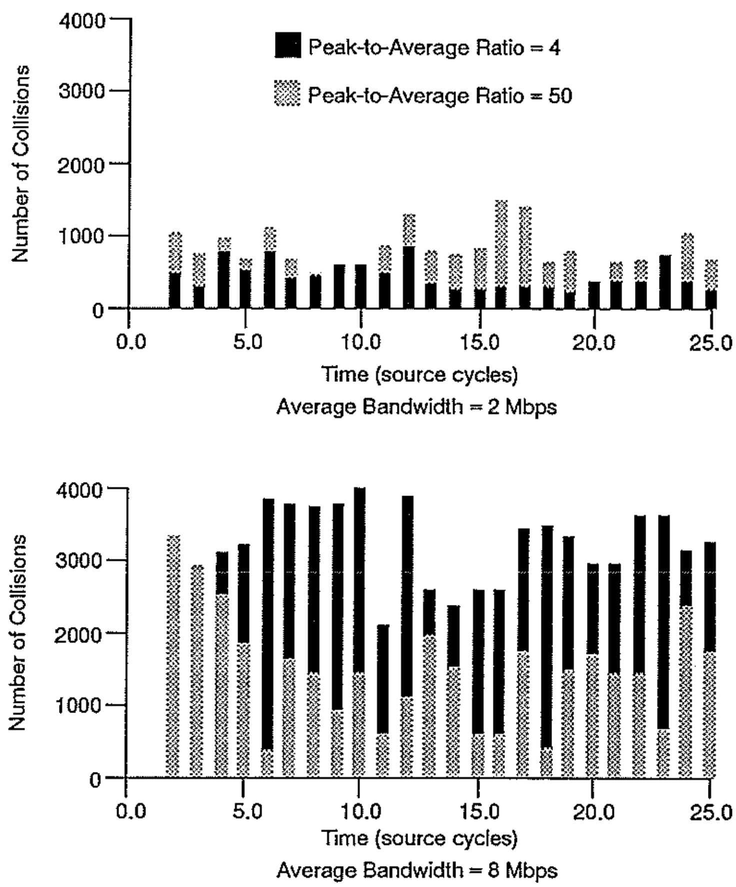

Figure 15: Number of collisions as a function of cycle number, peak-to-average ratio, and average bandwidth

Packet Delay

Mean packet delay as a function of throughput and source cycle time is shown in Figures 16 and 17. Recall that $1/\alpha$ and $1/\beta$ are the mean holding times in the idle and active states, respectively. Note that burstiness is directly related to both $1/\alpha$ and $1/\beta$. As the ratio of the idle state to the active state increases, the traffic burstiness increases. Figure 16 shows the mean packet delay obtained using sources with low burstiness traffic. As expected, the mean packet delay increases with the traffic burstiness. This increase in delay is much more significant under conditions of high burstiness, as shown in Figure 17. In this experiment, the ratio of mean holding time in the idle state to the mean holding time in the active state is equal to 50 and 100.

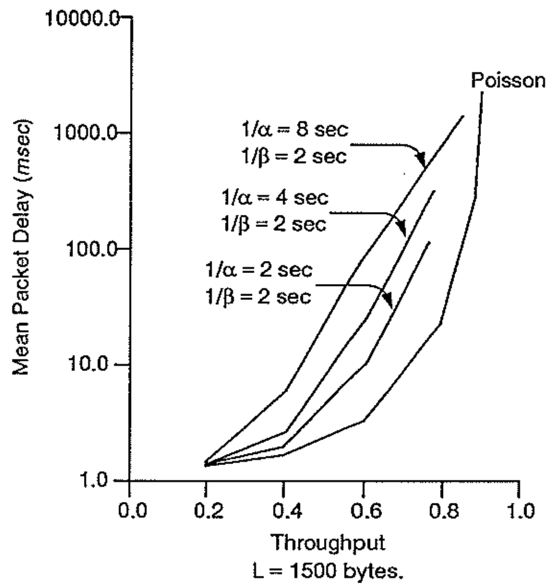

Figure 16: Mean packet delay vs. throughput and source cycle time

Packet Loss

Packet loss as a function of throughput and source cycle time is shown in Figures 18 (for low burstiness traffic) and 19 (for high burstiness traffic). Similar to the results presented in Section 7.2, packet loss deteriorates as the traffic burstiness increases, but only for low throughput values. For high throughput values, stations tend to be backlogged more often, and consequently the number of stations contending for the channel decreases. In this operation region, the Ethernet performance is chiefly dictated by the number of collisions on the channel, which is smaller for high burstiness traffic (see Figure 15). As a result, packet loss for high throughput values improves as the traffic burstiness increases.

7.4. Effect of Varying Number of Stations on the Ethernet Performance

In this section we present a comparison of simulation results obtained using 5 and 30 stations. As stated before, the number of stations used in all experiments is 30. When a large number of stations transmit bursty traffic on the channel, the resulting traffic on the channel is actually less bursty. In order to maintain the traffic burstiness on the channel, we ran also experiments with smaller number (5 in this case) of stations, and examined the changes in delay, packet loss, and queue length.

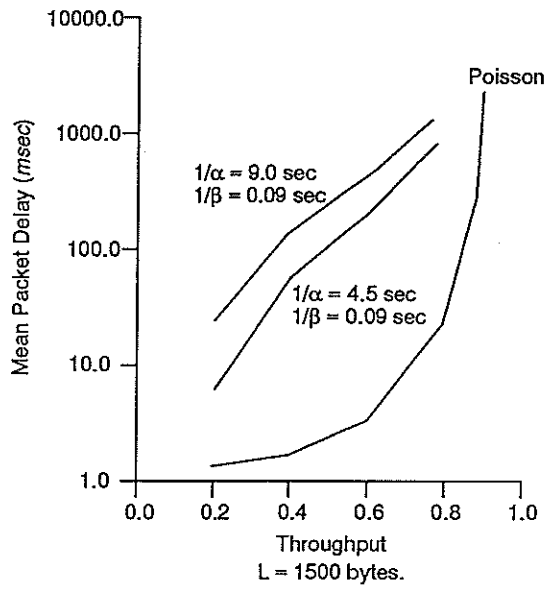

Figure 17: Mean packet delay vs. throughput and source cycle time

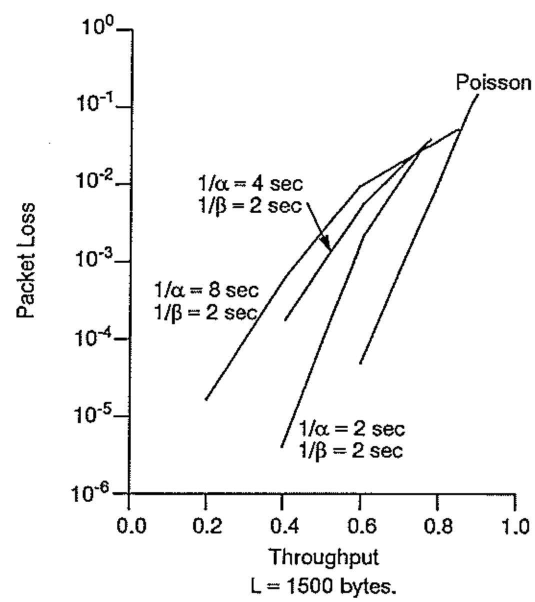

Figure 18: Packet loss vs. throughput and source cycle time

The experiments assume a burst factor of 100, a peak-to-average ratio of 50, and 1500 byte packets. Figure 20 shows mean packet delay, packet loss, and mean queue length as a function of throughput and number of stations. When the number of stations is 5, packets have fewer access points through which they can access the channel. Therefore, a newly generated packet has a very high chance of finding one or more old packets waiting in the queue. As a result, both the mean and the maximum queue lengths obtained with 5 stations are larger than those obtained with 30 stations (see Figures 20(a) and 20(b)). Furthermore, channel contention with 5 stations is lower than that with 30 stations. As a result, the number of collisions experienced by packets in the case of 5 stations is lower, which translates into lower packet loss, as shown in Figure 20(c). Figure 20(d) shows that the mean packet delay is higher for 5 stations because packets get backlogged more often due to fewer number of access points to the channel. The difference in packet delay for both experiments is larger for low throughput values. In this operating region, packet delay is mostly determined by the number of access points to the channel. However, for high throughput values, packet delay due to collisions increases very quickly. And since the number of collisions in both experiments is very significant in this operating region, packet delay for 5 stations becomes comparable to packet delay for 30 stations.

7.5. Performance Comparison with Deterministic and Exponential Arrivals

So far, we have assumed that the interarrival times of packets in the active state are exponentially distributed. There are models of bursty traffic in the literature [10] that use deterministic packet arrivals in the active state. In order to examine the impact of deterministic arrival on the results presented in this chapter, we ran an experiment which uses deterministic packet arrivals, 30 stations, 1500 byte packets, a peak-to-average ratio of 50, and a burst factor of 100. The results show that the difference in delay and packet loss for both packet arrival processes is negligible [9].

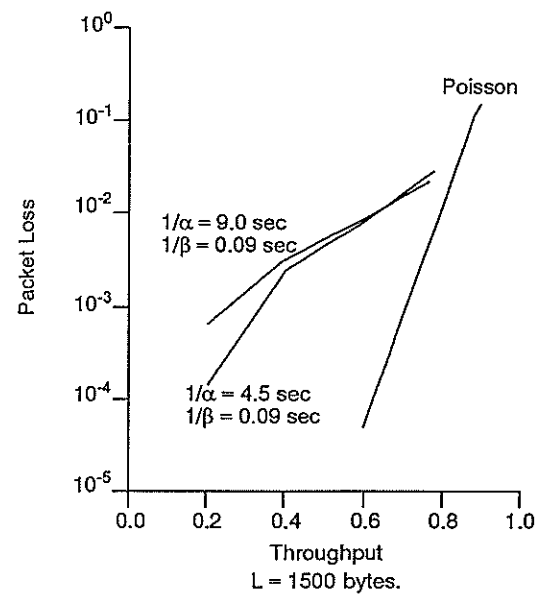

Figure 19: Packet loss vs. throughput and source cycle

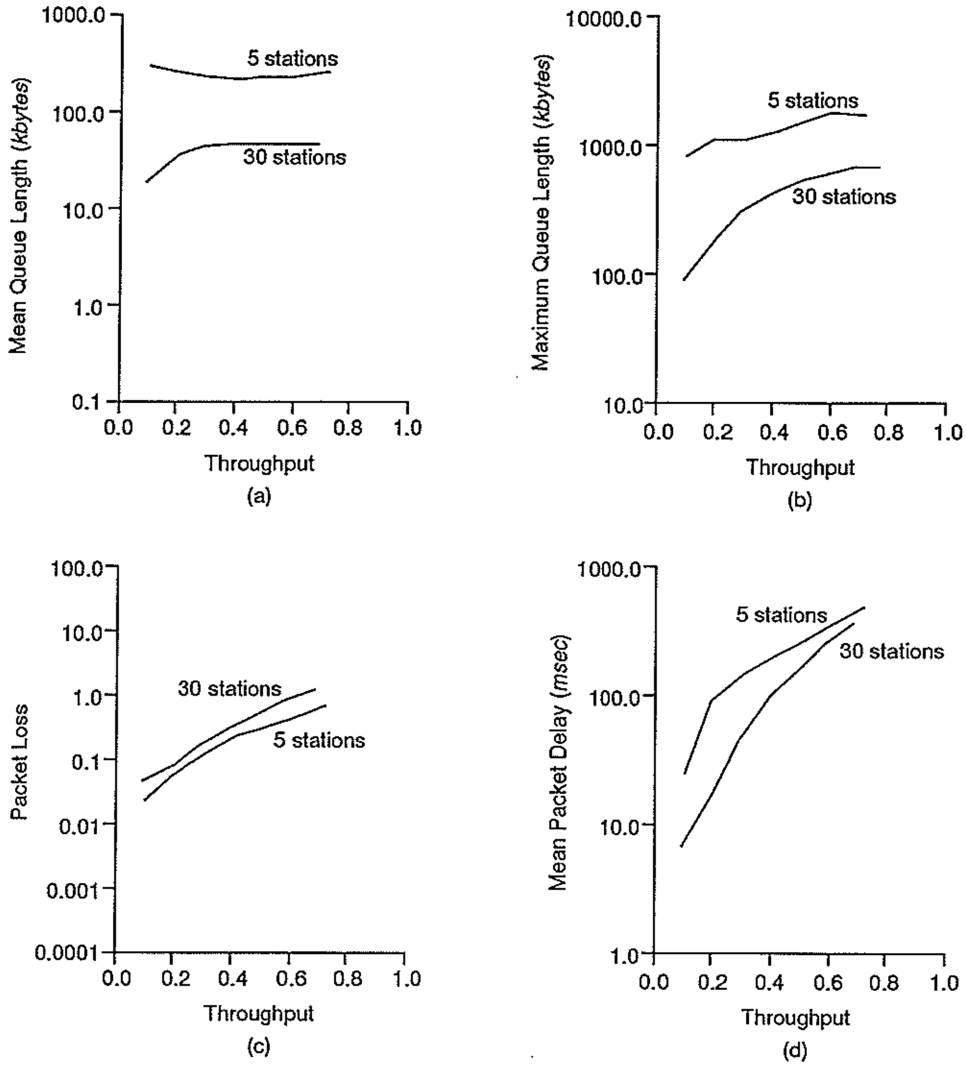

B = 100. L = 1500 bytes. $\lambda_p/\lambda_a = 50$.

Figure 20: Ethernet performance vs. throughput and number of stations

8. CONCLUSIONS

One of the important functions of the Ethernet and other local area networks include providing access to national and international (inter)networks. As providers of the public data networking plan the next phase of (inter)networks based on frame-relay, SMDS, and ATM technologies, they presume that the Ethernet will continue to serve as the access network. One major difference will however be the class of applications supported and their traffic characteristics. We claim that traffic output of most applications can best be modeled as a bursty source described in this paper, and it is essential to characterize the Ethernet performance under such traffic loads. The simulation study reported in this paper achieves this goal.

The simulation study shows that the inherent behavior of the Ethernet does not change with bursty traffic, which is an expected result. That is, as long as the utilization is less than a threshold value, packet delay is almost the same as the transmission time, queue lengths are minimal, and packet loss due to collisions is zero. However, as the utilization increases beyond the threshold, packet delay, queue lengths, and packet loss rate increase very quickly.

However, the important conclusion of this simulation study is that the Ethernet performance deteriorates faster with the *burstiness* of traffic sources, and that bursty sources are more demanding than Poisson sources. For example, packet loss due to collisions, packet delay, and buffer sizes increase with the *burstiness* of traffic sources. The peak to average bandwidth ratio of sources has an unexpected effect on the packet loss rate and queue lengths. At high utilization, packet loss and queue lengths are less for higher peak to average ratio of bursty sources.

These results are very useful in estimating traffic output of the access networks and performance *guarantees* to be expected. Furthermore these results are helpful in engineering the Ethernet gateways in terms of buffer sizing given expected latency and packet loss requirements.

References

[1] Digital Equipment Corp., Intel, and Xerox, "The Ethernet: A Local Area Network: Data Link Layer and Physical Layer Specifications," *Version 1.0*, September 1980.

[2] K. Kummerle, J.O. Limb, and F.A. Tobagi, *Advances in Local Area Networks*, IEEE Press, 1987.

[3] T.A. Gonsalves, "Comparative Performance of Broadcast Bus Local Area Networks with Voice and Data Traffic," *Ph.D. Dissertation, Stanford University*, March 1987.

[4] D.R. Boggs, J.C. Mogul, and C.A. Kent, "Measured Capacity of an Ethernet: Myths and Reality," *Proceedings of the ACM SIGCOMM Computer Communications Review*, 18(4): 222–234, August 1988.

[5] P.J.P. O'Reilly, "A New Technique for Performance Studies of CSMA/CD Local Networks," *Ph.D. Dissertation, Georgia Institute of Technology*, May 1983.

[6] H.D. Hughes and L. Li, "Simulation Model of an Ethernet," *Computer Simulation*, 3(4), December 1982.

[7] M. Decina and T. Toniatti, "On Bandwidth Allocation to Bursty Virtual Connections in ATM Networks," *Proceedings of the IEEE Globecom*, 1990.

[8] G. Gallassi, G. Rigolio, and L. Fratta, "ATM: Bandwidth Assignment and Bandwidth Enforcement Policies," *Proceedings of the IEEE Globecom*, 1989.

[9] T.Y. Mazraani, "High Speed Internet Protocols and Resource Management in the Ethernet," *M.Sc. Thesis, Washington University, St. Louis, Missouri*, August 1990.

[10] R. Jain and S.A. Routhier, "Packet Trains—Measurements and a New Model for Computer Network Traffic," *IEEE Journal on Selected Areas in Communications*, SAC-4(6), September 1986.