

LABORATORY FOR
COMPUTER SCIENCE



MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TM-292

PACKET TRAINS: MEASUREMENTS AND A
NEW MODEL FOR COMPUTER NETWORK TRAFFIC

Raj Jain and Shawn Routhier

November 1985

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Packet Trains: Measurements and a New Model For Computer Network Traffic

Raj Jain
Shawn Routhier

November 1985

Massachusetts Institute of Technology 1985

This research was supported in part by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under Contract No. N00014-83-K-0125.

Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, Massachusetts 02139

Packet Trains: Measurements and a New Model For Computer Network Traffic

Raj Jain¹
Digital Equipment Corporation
77 Reed Road (HLO2-3/N03)
Hudson, MA 01749

Shawn Routhier
Massachusetts Institute of Technology
545 Technology Square
Cambridge, MA 02139

ABSTRACT

Traffic measurements on a ring local area computer network at Massachusetts Institute of Technology are presented. The analysis of the arrival pattern shows that the arrival processes are neither Poisson nor Compound Poisson. An alternative model called "packet train" is proposed.

In the train model, the traffic on the network consists of a number of packet streams between various pairs of nodes on the network. Each node-pair stream (or node-pair process, as we call them) consists of a number of trains. Each train consists of a number of packets (or cars) going in either direction (from node A to B or from node B to A). The inter-car gap is large (compared to packet transmission time) and random. The inter-train time is even larger. The Poisson and the Compound Poisson arrivals are shown to be special cases of the train arrival model.

Another important observation is that the packet arrivals exhibit a "source locality". If a packet is seen on the network going from A to B, the probability of the next packet going from A to B or from B to A is very high.

Implications of the train arrivals, and source locality on the design of bridges, gateways and reservation protocols are discussed. A number of open problems requiring development of analysis techniques for systems with train arrival processes are also described.

Key Words: Computer Networks, Local Area Networks, Ring Networks, Computer Communication, Network Monitoring, Traffic Measurements, Network Performance, Workload Characterization, Performance Modeling.

¹ This work was done while Raj Jain was on a sabbatical at Massachusetts Institute of Technology.

INTRODUCTION

Many system design problems are essentially resource management problems which can be done more efficiently if the resource requirements can be accurately predicted. In computer networks, a common assumption is that packet arrivals are independent and unpredictable. However, If we could somehow predict something about future arrivals, we could design better packet handling strategies or at least implement current ones in a more effective manner.

The problem of predicting packet arrivals in a computer communication system is analogous to that of predicting memory page references in a paged memory computer. If the page references are assumed to be independent, analysis would indicate that a random page replacement strategy is as good as any other. On the other hand, actual page references have been observed to be correlated such that the probability of a page being referenced decreases as the time to its previous reference increases. This observation leads to a *least recently used* (LRU) policy, which is at present the most commonly used page replacement policy. Similarly, if we find that the probabilities of packets going to different destinations in a computer network are not the same, it may lead us to use different strategies, than if we assume the probabilities to be the same.

In designing computer networks we have a choice of at least two models of packet arrival patterns: a "car model" which assumes independent single packet arrivals, and a "train model" which assumes that a group of packets travel together. A protocol design based on the assumption of a train arrival would be quite different from one based on independent arrivals. In the car model, each car has to decide at each intersection (or exit) whether to take that exit or not. Even if all packets are going to one destination, they each make an independent decision, which may result in unnecessary overhead. The overhead is apparent on computer networks in which all intermediate nodes (routers, gateways, or bridges [Hawe 1984]) must make this decision for all packets, therefore resulting in long queues at each node. In a train model, on the other hand, the locomotive (the first packet of the train) may make the routing decision and all other packets may follow it.

The size of data objects being transported over computer networks has increased substantially compared to the increase in packet sizes. Packet sizes have generally been limited by the buffer sizes and by the need to be compatible with old networks. Transfer of a graphic screen may involve on the order of two million bits. This

increase in information size means that most communications involve a train of packets, not just one packet. This fact precipitated a closer look into actual traffic on the networks to determine whether there is a "train" effect. This paper is a result of that inspection.

A number of studies of LAN traffic exist in the literature [Shoch 1980, Rake 1984, Herskovitz 1982, Feldmeier 1984, Lloyd 1982, Murray 1984, and Terry 1983]. Also a number of authors have discussed issues related to measurement of LAN traffic [Chlamtac 1980, Amer 1982]. For wide area networks traffic studies we refer readers to an excellent survey by Tobagi et al [1978]. The measurements presented in this paper differ from other measurements in that we are looking for burstiness, predictability, locality and correlations in the traffic.

In the next section, we describe the network on which the measurements were done. We then describe the commonly used arrival patterns such as Poisson and Compound Poisson and introduce the concept of packet train arrivals. In the third section, we present an analysis of actual data that shows the existence of the train phenomenon. Finally, we present many applications of the train concept in the design and implementation of protocols.

ENVIRONMENT

All the measurements presented in this paper were done on a token ring network [Clark 1978] at Massachusetts Institute of Technology (M.I.T.) Laboratory for Computer Science. The ring configuration is shown in Figure 1. The ring operates at 10 megabits per second, connecting 33 Computers and 5 gateways on four floors of a building. It has a star shape with one wire center on each of the four floors [Saltzer & Pogran 1980]. It has two gateways to ARPAnet, one gateway to a 3 megabit per second Ethernet™, another gateway to a 10 megabit per second Ethernet, and a dial up gateway used by personal computers. There are three disk servers, which are used by many time-sharing VAX™ systems that use a *remote virtual disk* (RVD) protocol [Greenwald 1985].

The predominantly used higher level protocols on the ring include: DARPA Internet's *transmission control protocol* (TCP) [Postel 1981] used mostly for remote terminal (TELNET) [Davidson 1977] applications, *remote virtual disk* (RVD) protocol used by the disk servers, and *user datagram protocol* (UDP) [Postel

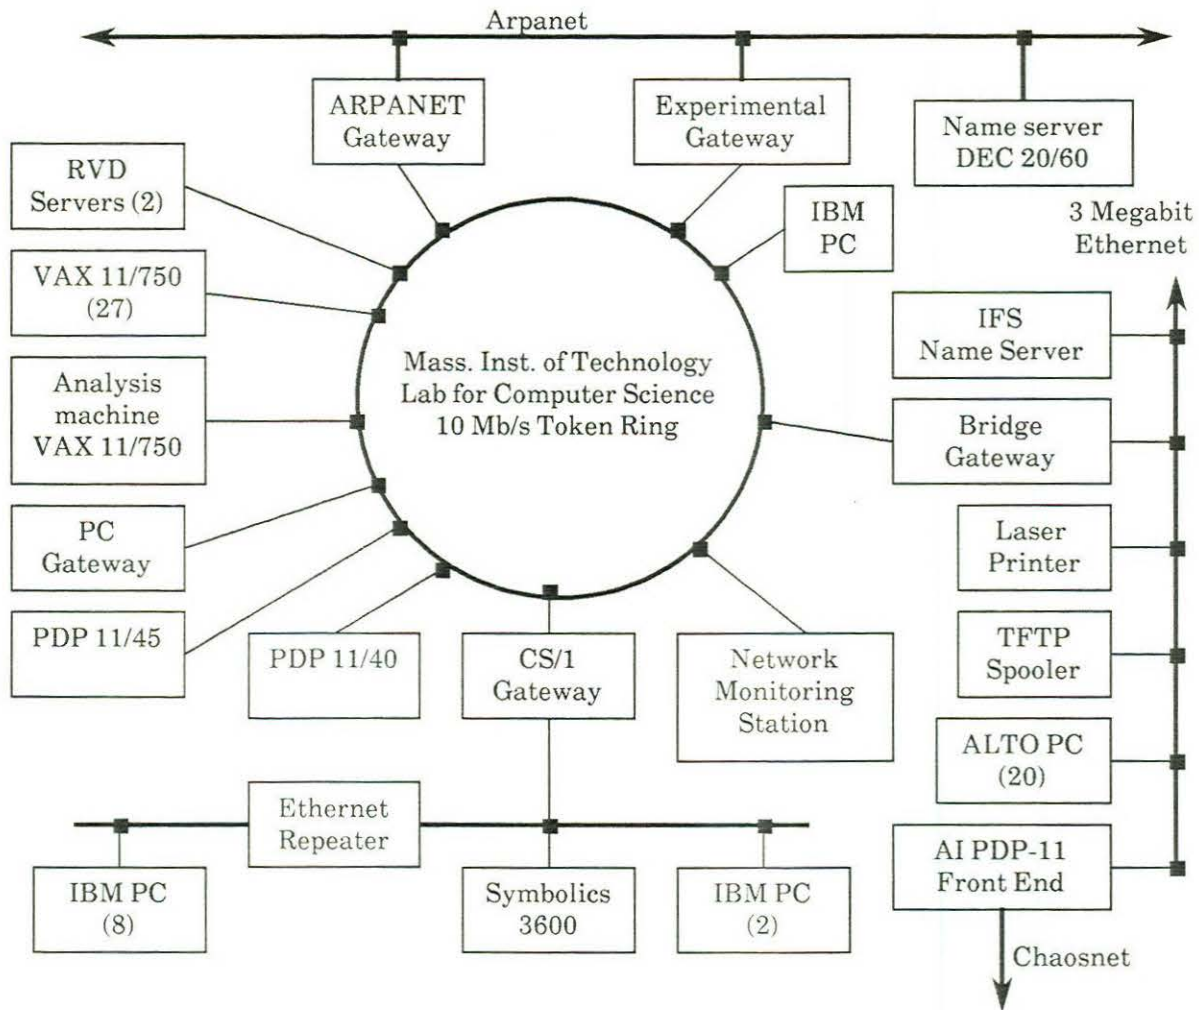


Figure 1: The MIT LCS 10 Megabit Ring Network

1980] used generally in a request-response mode by network inquiries and for a file transfer protocol named *trivial file transfer protocol* (TFTP) [Sollins 1981].

The traffic on the ring is continuously monitored by a station which is described in [Feldmeier 1984]. The monitor extracts the first 16 bytes of each packet's header, which contains source, destination, packet length and protocol type. Since the monitoring station does not have the power necessary for detailed analysis, it combines headers of 67 successive packets and prepares a *monitor packet* which is sent over the ring to a more powerful analysis machine. The monitor packets are sequentially numbered and therefore the analysis machine can recognize any packets that are sent from the monitoring station and lost on the way. We lose about 1 to 9% of the monitor packets. This is because the monitor uses a simple protocol

without retransmissions, it sends most of its packets when the network is busiest, and the receiver is a timesharing Unix™ system that may be busy with other tasks and its buffers may overflow. This introduces some discrepancies in the numbers presented but this should not change any of our conclusions since they do not depend upon exact numeric values.

We analyzed the traffic on the ring at different occasions in many different ways. The analysis presented here is for data collected during the week of December 10-17, 1984. During that week the ring carried a total of 11 million packets.

MODELS OF PACKET ARRIVAL

Model 1: Poisson Arrival Model

The most commonly used model for arrivals in analytical modeling is "Poisson Arrival" [Tobagi 1979, Liu 1982, Marathe 1982, Ramakrishnan 1982]. In this model,




Figure 2: The Poisson Model treats each packet as a black box.

the inter-arrival times t_i (between arrival of packet i and $i + 1$, see Figure 2) have the following characteristics:

1. *They are independent.*
2. *They are exponentially distributed. i.e., probability density function*

$$p(t) = \lambda \exp(-\lambda t)$$

If we plot a histogram of the inter-arrival times, it would be an exponentially decreasing function as shown in Figure 3a. There are many statistical techniques to verify if a particular arrival process is Poisson. One simple way to visually verify whether the inter-arrival times are exponentially distributed (the second condition above) is to plot a log histogram as is shown in Figure 3b. Since the probability is an exponential function, the logarithm of the probabilities would be a linear function:

$$\text{Log}\{p(t)\} = \log\{\lambda\} - \lambda t$$

Another property of the exponential distribution is that its coefficient of variation (the ratio of standard deviation to the mean) is one.

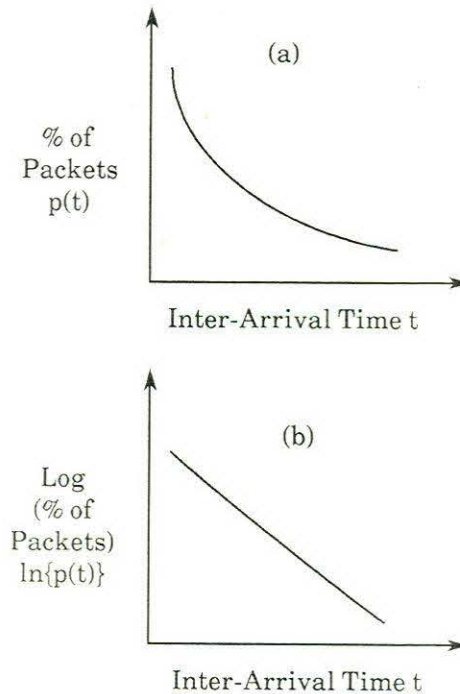


Figure 3: The histogram and log histogram of inter-arrival time of a Poisson process.

In order to verify if the inter-arrival times are exponentially distributed, Feldmeier [1984] plotted a log histogram of an earlier week's activity, which is reproduced in Figure 4. It is obvious from this figure that the log histogram is not linear. Rather it consists of three distinct straight line segments. This deviation from the Poisson is what eventually led us to the research presented in this paper.

Model 2: Compound Poisson Arrivals

An extension of the Poisson arrivals is the compound Poisson arrival process [Mohanty 1978, Meister 1980, Heyman 1982]. As shown in Figure 5, in this model, the arrivals occur in batches. The batch arrival process is Poisson in the sense that the inter-batch times are independent and exponentially distributed. The batch size is random. If the batch size is assumed to be geometric, it is possible to derive simple analytical results for the process.

On a log histogram, compound Poisson arrivals would result in a straight line with a spike near the origin. From Figure 4 we see that this is not the case. Our measurements also confirmed that simultaneous (or back-to-back) arrivals are rare. This is because the time required to prepare packets (in the order of milliseconds) is




Figure 4: The log histogram of inter-arrival times of packets as measured.




Figure 5: A compound Poisson arrival process consists of a sequence of bursts arriving in a Poisson manner. Each burst (batch) consists of several simultaneous arrivals.

generally much longer than the time required to transmit the packet on the network (in the order of 100 microseconds). Furthermore, most network nodes are shared by

non-network activities. This makes the time between successive packets from a single node large as well as random.

Model 3: Train Arrival Process

Our measurements (presented later in this paper) led us to a new model of arrival which we named the *train model* shown in Figure 6. Imagine that every node on the network is connected to every other node via a railroad track (sometimes called logical link). Consider the track between two nodes A and B. All packets on the track are flowing either from A to B or from B to A. A train consists of packets flowing on this track with inter-car time between them being smaller than a specified maximum, referred to as *maximum allowed inter-car gap* (MAIG). If no packets are seen on the track for MAIG time units, the previous train is declared to have ended and the next packet is declared to be the locomotive (first car) of the next train. The inter-train time is defined as the time between the last packet of a train and the locomotive of the next train.

Notice that the train packets flow in both directions, and that there may be several different trains traveling simultaneously on the network. For example, in between packets of a train traveling between nodes A and B, there may be seen packets of another train traveling between nodes C and D.

Before coming to the above definition of trains, we experimented with other possible models, such as: *source trains* - the train of packets starting from a given source, *destination trains* - the train of packets destined to a given node, etc. However, the analysis showed that these alternatives do not characterize the traffic well. This is because the sequence of packets going in one direction on a track is closely related to the sequence going in the reverse direction on the same track. In fact, in many protocols (e.g., in request-response protocols), given the sequence of packets going in one direction, it is possible to predict the sequence of packets going in the reverse direction.

The Poisson as well as compound Poisson models treat packets as black boxes. They do not distinguish between packets coming from different sources or those going to different destinations. They therefore lose some information which is easily available at the network layer. By dividing the packets into different tracks we are trying to use this information. An analogous example is the problem of predicting employee arrival times. If we stand at the gate and measure inter-arrival times of




Figure 6: The packet train model consists of a sequence of packets traveling between a given pair of nodes. The inter-car interval is much smaller compared to the inter-train interval.

employees, we may conclude that the successive inter-arrival times are independent and exponentially distributed, we therefore can not predict arrivals. On the other hand, if we note the badge numbers and their arrival times we can accurately predict arrival times for the next day, as people generally arrive around the same time each day. Ignoring the source and destination of packets on the networks is like ignoring the *badge numbers*. The packets on different tracks are independent, yet packets on the same track may be correlated.

ANALYSES OF MEASURED TRAFFIC

Analysis 1: Packets as Black Boxes

The first analysis that one can perform on a stream of packets is to treat them as black boxes. We do not look into the packet header fields or distinguish packets based on their source or destination. The time intervals between successive packets form a *time series* whose mean, standard deviation, and coefficient of variation, can be calculated. Another important quantity for a time series is its *auto-correlation function* (ACF). This function shows the relationship of an element of time series, say t_i with a previous element, say t_{i-k} . The covariance between t_i and t_{i-k} normalized by the variance of t gives the ACF at lag k [Box & Jenkins 1970]. The ACF always lies between -1 and +1. A negative ACF implies an *inverse relationship*, i.e. when t_{i-k} goes up, t_i is expected to go down and when t_{i-k} goes down t_i is expected to go up. A positive ACF implies a *direct relationship*, i.e. if t_{i-k} is high, t_i is also expected to be high. A zero ACF indicates no relationship or statistical independence.

If the packet arrivals on the network were a Poisson process, the inter-arrival times would have a coefficient of variation of one, and ACF would be zero at all lags.

Table 1 shows the results of such an analysis for our data. From the table we see that the ACF is small, which indicates that successive time-intervals are independent. However, the coefficient of variation is very high compared with unity, leading us to conclude that inter-arrival times are not exponentially distributed. Hence, *packet arrivals are not a Poisson process.*

Analysis 2: Node-pair Processes

Low auto-correlation as well as high variability are both bad news to a network designer as they both reduce predictability. Designers prefer high predictability because it helps improve the efficiency of resource management. If one could exactly determine the future resource requirements, the resources could be assigned optimally. Therefore, in analyzing the packet stream we started looking for ways to increase the *predictability.*

One alternative that comes to mind is to divide the stream into several node-pair processes as is shown in Figure 7. For each pair of nodes, say (A,B) on the network, all packets traveling between A and B form a time series (process) which can be analyzed separately. For example, in a network with four nodes A, B, C, and D, there would be a maximum of 6 node-pair processes, i.e., those belonging to AB, AC, AD, BC, BD, and CD. Given n nodes on the network, the packet stream can be divided into $n!(n-1)!/2$ node-pair processes. However, not every pair of nodes communicate and therefore the actually observed number of node-pair processes is rather low. We divided the packet stream into individual node-pair processes and we analyzed each node-pair process in exactly the same manner as in the previous section. We computed mean, standard deviation, coefficient of variation, and auto-correlation function at lag 1, 2, and 3. In addition, 90-percentiles of the time series were calculated.

The results for the 10 most active node-pair processes are shown in Table 2. The first two columns of the table list node-identifiers of the two nodes of the pair. The third column is the number of inter-packet intervals for packets that were sent between the nodes. The remaining columns give statistics of the inter-arrival times (measured in milliseconds). From the table, we see that ACF is still small indicating negligible correlation. The coefficient of variation for some processes is more than

Table 1: Characteristics of Network traffic

| Number of Intervals | Mean (ms) | ACF | | | Stand. Dev. | Coeff. Var. |
|---------------------|-----------|-------|-------|-------|-------------|-------------|
| | | 1 | 2 | 3 | | |
| 11,022,088 | 65.8 | 0.015 | 0.046 | 0.043 | 2835.3 | 43.1 |




Figure 7: A node-pair process consists of all packets traveling between a pair of nodes.

that in Table 1 and for others it is less than that in Table 1. Overall, coefficient of variation is now smaller than that in Table 1. However, it is still high as compared to that for a Poisson process. The node-pair processes are therefore neither Poisson nor any more predictable than the packet stream as a whole.

The 90-percentile column in Table 2 provides some new information. Notice that for most node-pair processes, the 90 percentile is lower than the mean. This implies that the distribution of inter-arrival times is highly *skewed to the right* (i.e. with a long tail on the right). Most packets arrive within a short interval of the previous arrival. However, in a few cases, there is a considerable delay leading to a *tail* in the distribution. This raises the mean above the 90-percentile value. This observation leads us to the train model discussed next.

Analysis 3: Train Model

Each node-pair process can be divided into a number of trains by specifying a *maximum allowed inter-car gap* (MAIG). We experimented with a few different MAIG values. Table 3 shows the analysis with a MAIG of 500 milliseconds. Although the choice of MAIG does impact numerical results, the final conclusions about traffic characteristics remain unchanged. We prefer the chosen value primarily because ninety-percentiles for most node-pair processes in our initial measurements were well below this value.

There are a number of observations that one can make from Table 3. First, the coefficient of variation is very near one. Ninety-percentiles are two to three times

Table 2

| Nodes | | Number of Intervals | Mean (ms) | 90- Perc* | Coeff. Var. | ACF | | |
|---------|----|---------------------|-----------|-----------|-------------|-----|-----|-----|
| 1 | 2 | | | | | 1 | 2 | 3 |
| 68 | 86 | 1,320,555 | 123.9 | 70 | 60.0 | 0.1 | 0.1 | 0.0 |
| 4 | 9 | 1,275,500 | 381.3 | 435 | 22.3 | 0.0 | 0.1 | 0.1 |
| 9 | 86 | 1,258,595 | 28.1 | 25 | 85.1 | 0.0 | 0.2 | 0.0 |
| 4 | 75 | 981,888 | 187.4 | 80 | 63.7 | 0.0 | 0.0 | 0.0 |
| 6 | 65 | 427,892 | 447.2 | 500+ | 25.5 | 0.1 | 0.1 | 0.0 |
| 67 | 68 | 412,316 | 317.5 | 230 | 47.7 | 0.0 | 0.0 | 0.0 |
| 65 | 87 | 397,095 | 733.2 | 85 | 32.9 | 0.0 | 0.1 | 0.0 |
| 75 | 86 | 395,635 | 25.9 | 45 | 27.9 | 0.0 | 0.0 | 0.0 |
| 68 | 87 | 349,953 | 988.3 | 275 | 27.2 | 0.0 | 0.1 | 0.0 |
| 6 | 66 | 294,332 | 459.5 | 500+ | 22.1 | 0.0 | 0.0 | 0.0 |
| Overall | | 10,850,688 | 1411.5 | 245 | 20.6 | | | |

*If a 90-percentile value is greater than 500ms, it is shown as 500+.

Table 3
(Divided Into Trains with MAIG = 500 ms)

| Nodes | | Number of Intervals | Mean (ms) | 90- Perc. | Coeff. Var. | ACF | | |
|-------------|----|---------------------|-----------|-----------|-------------|------|------|------|
| 1 | 2 | | | | | 1 | 2 | 3 |
| 68 | 86 | 1,315,298 | 31.0 | 70 | 1.6 | 0.0 | -0.1 | 0.0 |
| 9 | 86 | 1,257,178 | 16.2 | 25 | 1.9 | 0.2 | 0.0 | 0.0 |
| 4 | 9 | 1,177,654 | 85.3 | 255 | 1.3 | 0.1 | 0.3 | 0.2 |
| 4 | 75 | 956,195 | 30.1 | 65 | 2.0 | 0.0 | 0.1 | 0.1 |
| 67 | 68 | 401,008 | 59.1 | 195 | 1.6 | 0.1 | 0.0 | 0.0 |
| 75 | 86 | 395,174 | 22.1 | 40 | 1.5 | 0.1 | 0.0 | 0.0 |
| 65 | 87 | 388,007 | 35.0 | 75 | 1.4 | 0.1 | 0.1 | 0.0 |
| 6 | 65 | 382,232 | 84.5 | 200 | 1.1 | -0.1 | 0.1 | 0.0 |
| 68 | 87 | 332,011 | 55.4 | 155 | 1.6 | 0.0 | 0.0 | -0.1 |
| 6 | 66 | 255,689 | 87.5 | 205 | 1.1 | -0.2 | 0.2 | -0.1 |
| Inter-Car | | 10,228,405 | 51.1 | | 1.6 | | | |
| Inter-Train | | 622,283 | 23,773.0 | | 5.0 | | | |

the mean value. Both these observations lead us to believe that the inter-packet time in a train is exponentially distributed. However, the ACF is now generally non-zero. Non-zero correlations indicate that the inter-car periods are dependent. One explanation for this is that the network nodes have other tasks going on in parallel with networking activities and the nodes have their busy periods and idle periods. During a busy period, it takes long to send/route a packet and all inter-car intervals are long. During idle periods, all successive inter-car intervals are short. Some of

the correlations are negative indicating that sometimes short intervals are followed by large intervals and vice versa. This happens particularly in request-response type of protocols in which there is an alternating sequence of data (which takes a long time to generate) and requests or acknowledgements (which are generated quickly). The bottom two lines of the table indicate that the mean inter-car interval is about 51 ms which is small compared to the mean inter-train interval which is about 23.7 seconds. Although the variance in Table 3 is considerably smaller than that in Tables 1 and 2, it is still far from zero. At this point we wonder if we can further reduce the variance and increase the correlation. Doing so will help increase the predictability of resource demands and lead to design of more efficient packet handling strategies. To see if this can be done, we need to look further into the trains. This leads us to our next model called *tandem-trailer model*.

Analysis 4: Tandem-Trailer Model

A packet train consists of packets going in both directions. A train between A and B, as shown in Figure 8, for example, consists of one or more AB packets (packets going from A to B) followed by one or more BA packets, followed again by more AB packets, and so on. A sequence of successive packets going from one source to the same destination is called a tandem-trailer. The train consists of several tandem-trailer trucks going in alternate directions. A tandem-trailer may consist of *segments* of the same user message. It takes some time for a node to generate the data initially; it may have to be read from a disk, for example. However, once it has read a *chunk* of say 8 blocks it can send successive packets (of one block each, say) rather quickly. The leading-packet (first packet) of the truck, therefore, may take longer than the trailing-packets in the truck. Actually, the leading-packets are of two types: simple responses such as acks which required neither I/O nor any significant computation, and complex responses (such as long user messages) which may require I/O, computation and possibly process schedulings. The inter-arrival time before a simple response, complex response, and trailing packets would be expected to be small, large, and medium respectively.

It is obvious that the validity of the tandem-trailer model depends heavily on the networking protocol. Some protocols are purely request-response type with packets going in one direction that are data packets (responses) and the packets going in the other direction are either acknowledgements or requests for more data. To identify

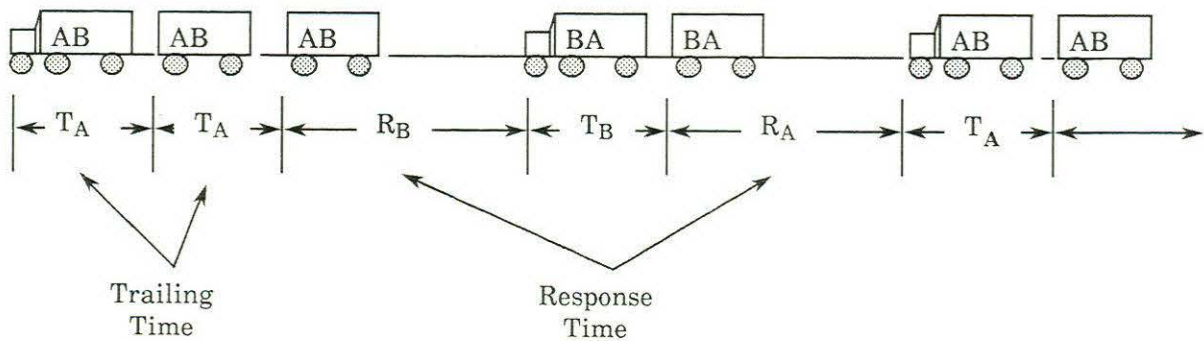


Figure 8: A train can be subdivided into several tandem-trailers. Each tandem-trailer consists of several packets going in the same direction.

tandem-trailers in the network traffic, we have to look in the protocol field as well, and analyze separate packets of different protocols.

Table 4 shows the results for four different protocols. The four most active node-pairs are shown for each of the protocol. The protocols are *remote virtual disk* (RVD), *user datagram protocol* (UDP), *internet control message protocol* (ICMP), and *transmission control protocol* (TCP). For each node-pair AB, the average time between adjacent A-to-B packets is shown under the column marked "Trailing Time A". The average time between adjacent B-to-A packets is shown in the next column. The "Response Time A" column shows average time between a packet going from B-to-A and the adjacent A-to-B response. The truck size is the number of packets going in one direction before a packet is seen in the opposite direction.

Table 4 shows that the average truck size in RVD is 2.4, i.e., every response is followed by one or two trailers. The trailing time is much shorter than the response time in each case.

UDP is the DARPA Inter-net's transport layer protocol used generally for request-response type of applications. In these applications, every packet traveling in one direction is followed by a packet traveling in the other direction. The average truck size is close to one, that is, there are no trailers.

ICMP is another protocol from the DARPA Inter-net suite. Periodically, each node sends a message to a control node. All packets flow in one direction only. There are no responses or leaders. Every packet is a trailer. The truck size is large.

TCP is also from the DARPA Inter-net suite. It is another transport layer. This protocol is used by many different applications. The average truck size is 1.5; that is

Table 4

| Prot. | Node # | | Trailing Time | | Response Time | | Truck Size | | Total Intervals |
|--------------|--------|-----|---------------|------|---------------|-------|------------|-----|-----------------|
| | A | B | A | B | A | B | A | B | |
| RVD | 68 | 86 | 11.3 | 16.0 | 88.0 | 53.1 | 3.8 | 2.6 | 1,314,272 |
| RVD | 65 | 87 | 10.1 | 13.9 | 58.0 | 57.4 | 1.8 | 2.2 | 387,747 |
| RVD | 67 | 68 | 14.8 | 47.5 | 91.6 | 189.9 | 3.8 | 3.1 | 380,766 |
| RVD | 68 | 87 | 12.6 | 10.0 | 169.4 | 66.3 | 1.3 | 2.9 | 331,624 |
| RVD Overall | | | 14.0 | | 74.4 | | 2.4 | | 3,999,668 |
| UDP | 4 | 5 | | | 220.3 | 26.4 | 1.0 | 1.0 | 31,069 |
| UDP | 5 | 75 | | | 51.0 | 27.0 | 1.0 | 1.0 | 9,500 |
| UDP | 5 | 74 | | | 51.0 | 23.0 | 1.0 | 1.0 | 8,392 |
| UDP | 6 | 65 | | | 129.1 | 31.1 | 1.0 | 1.0 | 7,096 |
| UDP Overall | | | | | 51.3 | | 1.1 | | 101,559 |
| ICMP | 195 | 195 | 92.1 | | | | 94759.0 | | 94,759 |
| ICMP | 4 | 67 | 141.1 | | | | 20308.0 | 7.0 | 20,315 |
| ICMP | 4 | 86 | 44.6 | | | | 6883.0 | | 6,883 |
| ICMP | 6 | 65 | 140.0 | | | | 2818.0 | 1.0 | 2,819 |
| ICMP Overall | | | 99.6 | | | | 2679.5 | | 128,615 |
| TCP | 9 | 86 | 7.1 | | 22.9 | 11.5 | 2.7 | 1.2 | 1,256,940 |
| TCP | 4 | 9 | 87.0 | 55.2 | 148.9 | 53.0 | 1.5 | 2.0 | 1,176,827 |
| TCP | 4 | 75 | | 31.2 | 51.0 | 7.6 | 1.0 | 1.3 | 955,656 |
| TCP | 75 | 86 | 14.9 | 83.4 | 23.7 | 12.2 | 1.2 | 1.2 | 395,122 |
| TCP Overall | | | 53.0 | | 61.6 | | 1.5 | | 5,965,835 |
| All Overall | | | 34.2 | | 65.2 | | 1.8 | | 10,202,774 |

about one-half of the packets are followed by a second packet which is going in the same direction. Both simple and complex responses can be seen in the table.

Overall, the truck size is 1.8. The trailing times are about half of the response time. The tandem-trailer model seems to be valid, but, the variance is still not zero. The predictability is basically the same as with the train model. Although the truck model gives a little more understanding of the underlying phenomenon, it still does not give understanding sufficient enough for use in protocol design or implementation.

Analysis 5: Source Locality

One additional phenomenon that we observed in our traffic was that of source locality. We have borrowed the term *locality* from the field of memory reference modeling. It is well known that successive references to memory have a tendency to cluster at the *least recently referenced page*. We observed a similar phenomenon in the network traffic. We found that successive packets have a tendency to belong to the same train. We termed this phenomenon source locality.

Most analyses on network modeling assume a uniform probability of a packet coming from all sources on the network. Under this assumption, given a network with m nodes,

Probability{ the next packet will come from a given source i } = $1/m$.

Probability{ the next two packets will come from a given source i } = $1/m^2$,

and,

Probability{ the next two packets will come from the same source } = $1/m$.

In the packets that we monitored, we found that 21 nodes either sent or received more than 100,000 packets. Even if we ignore other nodes, which were not as active, the probability that two packets will come from the same source would equal $1/21$ or 5%. However, we found that the probability of a packet going from A to B being followed by another packet going from A to B is 29% and that of an A to B packet being followed by a B-to-A packet is 31%. Approximately one-third of the packets followed a packet from the same source and another third followed packets from the destination of the previous source. This happened when we did not break the traffic into different node-pairs. This shows that the assumption of uniform distribution does not represent the real world traffic.

High source locality shows that the trains from different source-destination pairs do not overlap much. Of course, the amount of train overlap depends on the total load on the network. During periods of high utilization, one would expect a higher overlap, and less source locality. During periods of low utilizations, there is lower overlap and a high source locality. The networks are designed for peak loads which occur rarely and therefore the network utilization is generally low. The measurements were done several times. Each time we noticed the same

phenomenon including at peak periods, such as just before the final exams when all term papers and theses were due.

APPLICATIONS OF THE TRAIN MODEL

In this section we describe a few potential applications of our findings to protocol modeling/analysis, protocol implementation, and protocol design. We present only a brief discussion of these applications and show that the results based on packet train model would be quite different from those based on traditional assumptions. Solving systems models with train arrivals is left as an open problem.

Application 1: Protocol Modeling

The analysis has shown that an appropriate model of packets on a network is a train model. The traditional Poisson model has only one parameter - *the mean inter-arrival time* which is a cumulative result of many underlying phenomena. The train model on the other hand has many parameters, each of which explains a different phenomenon in the network. The inter-train time depends upon how often users transfer data objects. The parameter depends upon the user behavior. The inter-car time depends not on user behavior, but it depends solely on the system (hardware/firmware/software) and the protocols. The train size is related to data object sizes. Given a distribution of object sizes, one can come up with a distribution of train sizes. The tandem-trailer model provides a yet more detailed insight. The average size of the truck is related to the flow control window sizes used in the protocols. The response time may or may not include I/O time depending upon whether or not it is a simple or complex response. It may therefore depend upon the characteristics of the I/O device. The trailer time, on the other hand, does not depend upon the I/O device but is rather a function of the protocol and system characteristics.

An analog of this problem happens in the modeling of time-sharing systems. The average service time at the system (in the machine-repairman model) was the only parameter that the initial models of time-sharing systems had. This was later replaced by a more detailed (central server) model with explicit modeling of time spent at disks, CPU, paging device, etc. The train model provides a more detailed understanding of user and system effects in a networked environment. The measured train parameters are summarized in Table 5.

Table 5: Measured Train Characteristics

| | |
|----------------------|------------|
| Inter-Car Time: | |
| Mean | = 51.1 ms |
| C.O.V. | = 1.6 |
| ACF(1) | = 0.2 |
| ACF(2) | = 0.3 |
| ACF(3) | = 0.2 |
| Inter-Train Time: | |
| Mean | = 23.8 sec |
| C.O.V. | = 5.0 |
| Number of Cars/Train | = 17.4 |
| Number of Cars/Truck | = 1.8 |

Application 2: Path Caching in Protocol Implementation

The present analysis has shown that successive packets have a source locality. Given a packet, we can predict with high probability that the next packet will be destined either to the destination or the source of the previous packet. Normally, a network has several thousand nodes. Finding the link on which to forward a packet requires sophisticated table search and hashing procedures at each node. The existence of source locality indicates that considerable saving in table lookup can be obtained by simply saving (caching) the table entry for the last packet. The overhead can be further reduced by pre-fetching, the table entry for the source and the destination of a packet. Thus, there is a high probability of having the table entry in our cache even before the next packet arrives. Even a two entry cache would give approximately 60% hit ratio.

Application 3: Number of Buffers in Routers/Gateways/Bridges

A key parameter in network design is the number of buffers required at intermediate nodes. Any node connecting a high-speed network with a low-speed network tends to have a queue of packets to be forwarded to the low speed network. Queuing theorists often use an M/M/1 model for the node to determine the number of buffers that are required to make the probability of buffer overflow (probability of losing a packet) sufficiently small. The result is based on the average inter-arrival time and average service time only. The train model indicates that the buffers would

also depend on the train size. Also, the trains may create more congestion than predicted by an M/M/1 model.

Application 4: Dynamic Circuit Management

A network path using a telephone connection is a dynamic circuit, such that the circuit is established only for the duration of the transfer. To minimize costs, it ought to be closed as soon as we are sure that no immediate traffic is expected to arrive. In network architectures with connectionless orientation (e.g., DNA, and DARPA Inter-net), the end nodes originate the traffic but do not know the path a packet takes. The intermediate node originating the telephone call has to *dynamically* open and close the circuit based on the traffic. With the Poisson model, the packet arrivals are independent and the fact that we have not seen any packet for the last one-hour has no effect on the probability of arrival in the next one-second. With the train model, one could set the cutoff point based on the distribution of inter-car arrivals, and close the connection when the idle time exceeds this cutoff value. The optimal cutoff point will depend on the train parameters as well as on the tariff structure, e.g., cost of opening and closing a telephone call.

Application 5: Reservation Switching

Reservation switching is a form of switching commonly used in satellite links [Kleinrock 1977]. As shown in Figure 9, the time in this switching method is divided




Figure 9: The reservation switching as used on satellites.

into equal size frames and each frame consists of several slots. An empty slot can be obtained by contention. If a node succeeds in obtaining a slot without collision, the slot that is in the same position in the next frame is reserved for the node. If the node does not use the slot in the next frame, the reservation is cancelled and the slot is

again made available to other nodes via contention. Reservation switching is based on the belief that a node sending one packet is very likely to send more packets. This is especially true for voice traffic. Active connections send packets every 20 milliseconds or so.

An example of reservation switching in LAN environments is seen in the Cambridge Fast Ring [Temple 1983]. The ring has two types of slots: normal slots, and channel slots. Once a normal slot has returned to its source after transmission it must be passed on empty. Channel slots remain reserved for a source until the source passes it on empty.

The efficiency of reservation switching depends upon the correct selection of the frame size. For example, if the frame size for voice traffic were chosen to be 10 milliseconds, the nodes will not have a packet to send in successive 10 millisecond intervals, and all reserved slots will go empty. For data traffic, the determination of optimal frame size is not straightforward because the data packets do not arrive in a perfectly regular pattern. Had the train measurements showed a zero variance for inter-car intervals, the ideal frame size would be the fixed inter-car interval. The network measurements have shown that the inter-car intervals are not fixed, nonetheless the ideal frame size is a function of the inter-car interval distribution. Further, the observations show that a packet going from A to B should result in a slot being reserved for packets coming from A *as well as for those coming from B*. The measurements (Table 4) show that the direction changes on the average after 1.8 packets and therefore, for data traffic, we need *bidirectional reservation* protocols.

HIERARCHY OF LAN WORKLOAD MODELS

We have presented several models of LAN traffic. The appropriate model depends upon the level of detail desired. The Poisson model is the least detailed model of the traffic. It can represent the traffic with the single parameter of mean inter-arrival interval. This model treats packets as black boxes in that we do not look into the packets. The next level is that of node-pair processes which requires separating the traffic into several streams based on the source and destination. The train model is the next level down from node-pair processes. At this point, we further subdivide a node-pair sequence into several trains. A new train starts if a car is not seen for MAIG interval. Each train consists of several tandem trailers which is the next level model. Each trailer truck consists of leading-packets (responses) and trailing packets. The responses can be simple or complex. The

complete hierarchy tree is shown in Figure 10. As we go down the tree, the variance decreases, skewness decreases, and correlation increases leading to higher predictability.

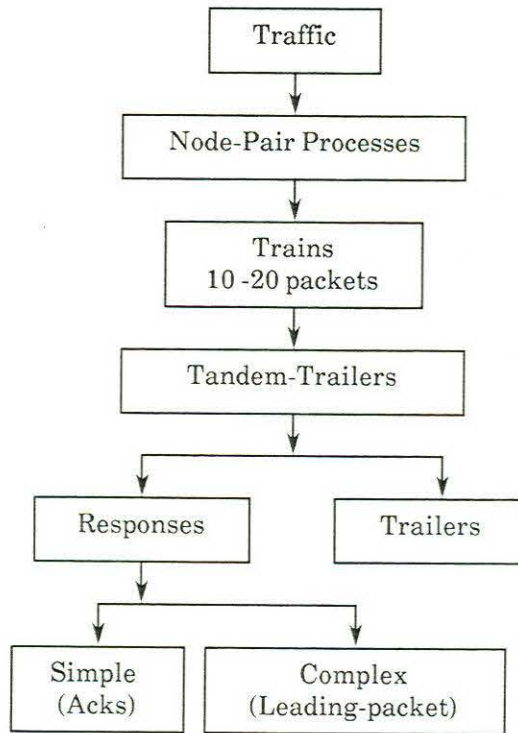


Figure 10: Hierarchy of LAN workload Models.

SUMMARY

The packet train research has shown that the packet arrival process is neither a Poisson process, nor a Compound Poisson process. The packet arrivals follow a train model. A train consists of packets traveling in both directions between a given node-pair. Although, the packets of a train are close to each other, they are too far apart to be considered as simultaneous arrivals. The inter-car interval is much smaller than inter-train interval.

The inter-train time is a user parameter and it depends upon the frequency with which applications use the network. The inter-car interval is a system parameter and depends primarily on the network hardware and software. In Poisson arrival models, these parameters are merged to give a single parameter: *mean inter-arrival time*.

The train model is a generalization of which other models are special cases. If we find ways to analyze network protocols using a train model, the results can be used for other models by simply setting the inter-car and inter-train time distributions to values such as are shown in Table 6. Also shown in the table is a simplification of

Table 6: Special cases of the packet train model

| Inter-Car Interval Distribution | Inter-Car Interval Auto-Correlation | Inter-Train Interval Distribution | Network Traffic Model |
|---------------------------------|-------------------------------------|-----------------------------------|-----------------------|
| Exponential(λ_1) | Zero | Exponential(λ_1)* | Poisson |
| Zero | Zero | Exponential | Compound Poisson |
| Constant | Zero | Exponential | Regular Train |
| General | Non-Zero | General | Train |

* The two exponential distributions have the same mean.

the train model called *regular train*. In this the trains arrive in a poisson process and consist of a random number of cars with constant inter-car intervals. This type of train, which represents voice traffic, is simple to analyze as well as helpful to network designers since the car arrivals can be easily predicted.

The packets have a high source locality such that given a packet going from A to B, there is a high probability that the next packet will be going either from A to B or from B to A. The probability of packets from other sources is small.

The lessons learned from the train model can be used to improve protocol analysis, design, and implementation. Cars are good for transporting a small number of people to a large number of independent destinations whereas trains are good for bulk transfers.

To verify the existence of trains, we need to repeat the analyses in other environments. That will help determine typical values of train parameters for today's traffic. These parameter values are obviously of interest to network analysts. Even without that verification network designers argue that the amount of information being transported across computer networks is increasing and that they ought to look at ways of making bulk transfers more efficient. We therefore hope that in the near future we will see more *railroad tracks* along with the *highways*.

ACKNOWLEDGEMENTS

Prof. Jerry Saltzer provided ideas and encouragement during the project. Dave Feldmeier designed the monitor and helped it maintain. Dr. Dave Clark helped in getting the resources together.

The following are trademarks of Digital Equipment Corporation: DEC, VAX, VAX 11/750, PDP, PDP 11/45, PDP 11/40, and PDP 11. Ethernet is a trademark of Xerox Corp. PRONET is a trademark of Proteon Associates.

REFERENCES

1. Amer, P. D., "A Measurement Center for the NBS Local Area Computer Network," *IEEE Trans. on Computers*, Vol. C-31, No. 8, August 1982, pp. 723-729.
2. Box, G.E.P., and G.M. Jenkins, *Time Series Analysis: Forecasting & Control*, Holden-Day, 1970.
3. Chlamtac, I., "Issues in Design and Measurement of Local Area Networks," *Proc. Computer Measurement Group Conf. CMG XI*, December 1980, pp. 32-34.
4. Clark, D. D., K. T. Pogran, and D. P. Reed, "An Introduction to Local Area Networks," *Proc. of the IEEE*, Vol. 66, No. 11, November 1978, pp. 1497-1517.
5. Davidson, J., et al, "The ARPANET Telnet Protocol: Its Purpose, Principles, Implementation, and Impact on Host Operating System Design," *Proc. 5th Data Commun. Symp.*, Snowbird, UT, September 1977.
6. Feldmeier, D. C., "Empirical Analysis of a Token Ring Network," *Massachusetts Inst. Technol. Lab. Comput. Sci.*, Cambridge, MA, Tech. Memo., MIT/LCS/TM-254, January 1984.
7. Greenwald, M., "Remote Virtual Disk Protocol Specifications," *Masachusetts Inst. Technol. Lab. Comput. Sci.*, Cambridge, MA, Tech. Memo., in preparation.
8. Hawe, W. R., A. Kirby and B. Stewart, "Transparent Interconnection of Local Networks with Bridges," *Journal of Telecommun. Networks*, Vol. 3, No. 2, September 1984, pp.116-130.
9. Herskovitz, J., "Evaluating Local Network Performance," *Computer Science and Technology*, October 1982, pp. 389-396.
10. Heyman, D. P., "An Analysis of the Carrier-Sence Multiple-Access Protocol," *Bell System Technical Journal*, Vol. 61, No. 8, October 1982, pp. 2023-2051.
11. Kleinrock, L., "Performance of Distributed Multi-Access Computer-Communication Systems," *Proc. IFIP Congress*, 1977, pp. 547-552.
12. Liu, M. T., W. Hilal, and B. H. Groomes, "Performance Evaluation of Channel Access Protocols for Local Computer Networks," *Proc. Compcon 1982*, pp. 417-426.

13. Lloyd, P. J. and R. H. Cole, "Transport Protocol Performance over Concatenated Local Area and Satellite Networks," Proc. Conf. Data Networks with Satellites, September 1982.
14. Marathe, M. and W. Hawe, "Predicted Capacity of Ethernet in a University Environment," Proc. Southcon'82, March 1982, pp. 1-10.
15. Meister, B., "Waiting Time in a Preemptive Resume System with Compound-Poisson Input," Computing, Vol. 25, No. 1, 1980, pp. 17-28.
16. Mohanty, N. C., "Buffer Behavior for Batch Arrivals and Single Output in Satellite Channel," Proc. IEEE 1978 Natl. Telecommun. Conf. NTC'78, Birmingham, AL, Part III, December 1978, pp. 40.4/1-3.
17. Murray, D.N., and P.H. Enslow, Jr., "An Experimental Study of the Performance of a Local Area Network," IEEE Commun. Mag., Vol. 22, No. 11, November 1984, pp.48-53.
18. Postel, J., "User Datagram Protocol," ARPA RFC-768, 1980.
19. Postel, J., ed., "Transmission Control Protocol - DARPA Internet Program Protocol Specification," ARPA RFC-793, September 1981.
20. Postel, J., "Internet Control Message Protocol," ARPA RFC-821, 1982.
21. Ramakrishnan, K. K. and S. K. Tripathi, "A Common Framework for Studying the Performance of Channel Access Protocols," Proc. of the Comp. Performance Evaluation Users Group (CPEUG'82), October 1982, pp. 365-373.
22. Saltzer, J.H. and K.T. Pogran, "A Star-shaped Ring Network with High Maintainability," Computer Networks 5(4):239-244, Oct./Nov., 1980.
23. Shoch, J.F. and J.A. Hupp, "Performance of the Ethernet Local Network," Communications of the ACM 23(12):711-721, December 1980.
24. Sollins, K., "The TFTP Protocol (Revision 2)," ARPA RFC-783, June 1981.
25. Sventek, J., et al, "Token Ring Local Area Networks: A comparison of Experimental and Theoretical Performance," Proc. IEEE Computer Networking Symp., Silver Springs, MD, December 1983, pp. 51-56.
26. Temple, S., "The Cambridge Fast Ring," Proc. IFIP WG6.4 International Workshop on Ring Technology, University of Kent, England, September 1983.
27. Terry, D. and S. Andler, "Experience with Measuring Performance of Local Network Communications," Proc. IEEE Computer Conf., Comcon'83, April 1983, pp. 203-205.
28. Tobagi, F. A. and V. B. Hunt, "Performance Analysis of Carrier Sense Multiple Access with Collision Detection," LACN Symposium, May 1979, pp. 217-244.
29. Tobagi, F. A., et al, "Modeling and Measurement Techniques in Packet Communication Networks," Proc. of the IEEE, Vol. 66, No. 11, November 1978, pp. 1423-1447.