Reprint

Bounded Transfer Delays in Buffer Insertion Rings using Load-Controlled Scheduling of Traffic

T. Antonakopoulos, J. Koutsonikos and V. Makios

The 5th IEEE Workshop on Metropolitan Area Networks

TAORMINA, ITALY, MAY 1992

"Bounded Transfer Delays in Buffer Insertion Rings using Load-Controlled Scheduling of Traffic"

T. Antonakopoulos, J. Koutsonikos and V. Makios

Laboratory of Electromagnetics, University of Patras, 26500 Patras, Greece

Abstract: The use of a modified Buffer Insertion Access method in the Metropolitan Area Network environment is examined in this paper. This method has been modified using a simple Load-Controlled Scheduling of Traffic protocol, which overrides its intrinsic disadvantages and achieves the adaptation of the offered load to the network conditions. Each station independently adjusts its asynchronous traffic to the available bandwidth by monitoring the instantaneous traffic flow. Although the performance of the asynchronous class of traffic decreases, the synchronous traffic is served with bounded transfer delays and the total network performance makes it appropriate for use in the MAN environment.

I. INTRODUCTION

The development of the Metropolitan Area Networks (MANs) is characterized by high transmission speeds, long inter-station distances, use of small and constant length packets (usually called cells) and the support of various services with a variety of characteristics and constraints. For MAN access methods, various techniques have been proposed and some of them (like FDDI and DQDB) have been already used by standardization bodies. These techniques are extensions of some very popular techniques used in the LAN environment and are based on the new technological achievements and the new service requirements.

One of the most interesting topics in the communications networks development has been the algorithm of the medium allocation among the various users and their services. From the early stage, where the supported services differed slightly in their characteristics, the network utilization has been an important performance measure. This measure is mainly a consideration from the network point of view and expresses the fulfilment of the user requirements only under specific service-related conditions. As far as the supported services diverse in their characteristics and different requirements, usually contradictory, must be satisfied, new performance measures have been considered [1]. These measures are related with the set of Quality of Service (QOS) requirements of a specific class of traffic and express the network consideration from the individual user point of view. These QOS are mainly the maximum delay, the maximum delay jitter, the average throughput, the acceptable bit error rate and the acceptable packet error rate [2].

The Buffer Insertion technique in ring topologies was initially considered for use in the LAN environment [3], [4], but its performance to support 'real-time' services was unacceptable, without basic modifications. During Buffer Insertion Ring operation, the station interfaces are connected by point-to-point links and ring buffers are used to control the flow of information inside the network. The used access technique is very simple: a station transmits its own packets if there are no data in its ring buffer, otherwise the station defers its access until the ring buffer is empty. Such a technique can provide an increased network throughput, which approaches 400 percent in dual ring structures, but its main disadvantage is the variable transfer delays occurring due to the intermediate ring buffers, resulting to inefficiency for use in the MAN environment, especially to support synchronous class of traffic.

The main problem of the Buffer Insertion technique is that it does not include any mechanism to allocate the network bandwidth according to the service requirements of all stations, like the Target Token Rotation Time (TTRT) of FDDI [4], but each station transmits, if it has a packet to transmit, without taking into account the behavior of the residual stations.

For network access methods which are intrinsically unstable and unfair, a type of Load-Controlled Scheduling of Traffic (LOCOST) [5] can be used to arbitrarily manage the traffic inside the network, to improve its performance and probably to satisfy the requirements of the various traffic classes. The LOCOST method measures the traffic flow on the transmission medium in each station independently and adjusts the station transmission rate in order to maintain the total traffic flow at a target value. In this work we adapt the LOCOST concept to a Buffer Insertion Ring type of network, in order to fulfil the requirement of the synchronous class of traffic for bounded transfer delays.

In section II, we describe the basic station architecture and the various types of services supported by the network, while in Section III the developed type of LOCOST is described and three load control algorithms are presented. In Section IV, the used simulation parameters are described and a comparison of the performance of the simple Buffer Insertion Ring to a modified one by LOCOST is given. The influence of the LOCOST parameters to the network performance is also presented. Finally, we evaluate the usage of the proposed scheduling of traffic in Buffer Insertion Rings, in accordance with the requirements of a MAN environment.

II. THE SYSTEM MODEL

We will start our discussion on the proposed station model and network access protocol by determining the classes of traffic serviced by the network and their requirements. In many works [6] it is assumed that messages arriving at a station belong to a specific class of traffic. This consideration is no valid when an integrated services network is analysed. In HSLANs and MANs three classes of traffic are considered, the isonchronous, the synchronous and the asynchronous class [2]. The isochronous traffic has a constant delay and will not be explicitly considered in the following discussion, while the synchronous traffic, which requires a bounded transfer delay, will be the main subject of this work. The asynchronous traffic has no delay limit and usually receives the network bandwidth not used by the other two classes of traffic. A complete table of the various traffic classes and their QOS parameters is given in [2]. (This discussion is related with the MAC sublayer and the isochronous traffic will be considered to be a special case of the synchronous traffic, where the delay bounds approach the mean delay value at the upper layers, probably using a type of elastic buffer).

In our model, we consider two classes of traffic, synchronous and asynchronous, and a queue devoted to each one. For the synchronous class of traffic, we also consider that each station can handle at most a single synchronous message at any time. Each queue is served by a 'limited' service discipline, where each station transmits up to m packets per access (in our case m equals 1).

The network under consideration uses small and constant length packets with a destination release method. The use of constant length packets allows easy implementation of the access mechanism and, for high speed networks, this can be used to achieve constant delay in the intermediate stations during the transmission of a packet, as will be shown later. Although the destination release scheme increases the MAC complexity, it provides improvement in delay and throughput levels [4], allowing spatial reuse of the available bandwidth.

The MAC sublayer of each station consists of three types of buffers, the ring Buffer, the Asynchronous Traffic Buffer and the Synchronous Traffic Buffer. The 'Ring Buffer' (RB) stores temporarily the incoming traffic and inserts a constant delay, irrespective if there is a station transmission or not. The length of this buffer has been set to 1.5 times the constant packet length and that results to a delay of 1.5 times the packet duration per station to the packets passing-through. This delay is small compared to the delay imposed by other network parts and we will give an example to make it clear. Suppose that the transmission speed is 1 Gbit/sec, the packet length is 1000 bits and the distance between adjacent stations is 5 km, then the delay imposed by each station due to its RB is 1.5 μsecs, while the time needed to traverse the distance between adjacent stations is 25 μsecs.

For each class of traffic a different buffer is used, both at the reception and the transmission path. The Ring Buffer acts as a constant delay line and its pipeline structure can be easily implemented, while it allows the easy implementation of the scheduling of traffic algorithm (e.g. using Content Addressable Memories in conjuction with this buffer). During a packet shift inside the RB, its destination address and type are examined and if the destination address matches with the station address, the packet is removed from the network and, according to its type, is stored into the respective buffer. A diagram of the station architecture is given in Figure 1. When the station has a packet to transmit, it examines the RB and, if it is occupied less than half of a packet, the station starts transmission immediately. If there are packets both in the asynchronous and the synchronous traffic buffers, the synchronous buffer has the priority to transmit. The configuration of the four switches in Figure 1 can easily demonstrate the access method and the way the traffic is handled inside the lower part of a station. Due to the constant length of the Ring Buffer, it is obvious that there is a constant delay to a passing-through packet, even if the station starts transmission when half of the packet is inside RB. This mechanism guarantees that packets of the same message have the same delay during their transmission from the source to the destination station and the only cause of delay variation is the variable queueing time in the source station. Especially for the synchronous traffic, this queueing time is mainly affected by the upstream traffic and the MAC protocol described up to now includes no mechanism to prevent these queueing delay variations. The adaptation of the LOCOST method to the Buffer Insertion Ring is intended to overcome these disadvantages, in order to make the Buffer Insertion Ring capable to support synchronous class of traffic.

In Figure 2, a timing diagram of the way the station handles incoming traffic is shown. A message arrival of asynchronous traffic is presented with a batch of packets, while the synchronous traffic is presented as a constant rate packet stream. When the station gains the right to transmit into the network, a transmission from the asynchronous traffic buffer happens only if there are no packets inside the synchronous traffic buffer. The lengths of these two buffers are given under specific traffic conditions, as well as the inter-arrival packet times of the synchronous service at the receiving station. These inter-arrival times must be bounded in order to reliably support synchronous traffic and fulfil their QOS requirements.

III. THE ALGORITHM OF THE LOAD-CONTROLLED SCHEDULING

The basic concept of LOCOST, presented by John O. Limb [5], states that, in intrinsically unfair access methods, the appropriate scheduling of traffic can improve the network performance and override the disadvantages of the access method. The scheduling is based on traffic observations and a distributed algorithm adapts the offered load to the current traffic conditions. The main advantage of this method is that it is independent of the access method and operates in each station independently from the others. In this work we have used this basic idea to control the load in a Buffer Insertion Ring network, but many modifications have been done to adapt it to the specific application requirements and to make it appropriate for hardware implementation. Although these modifications have changed the nature of the method, we will also use the name of LOCOST for our version of this method, for simplicity purposes.

Before starting the description of the applied LOCOST, we will take a closer look on the two classes of traffic and on the way they generate packets for transmission. An asynchronous traffic source generates a batch of packets at the message arrival time and the load at the respective buffer increases abruptly. The number of generated packets depends on the message length (e.g. file length). The synchronous traffic generates packets with constant rate for the duration of a connection. This packet rate depends on the used packet length and the service bit rate, while the packet generation time is equal to the time needed to formulate a packet from the specific bit stream.

For proper LOCOST operation, during the network initialization each station informs the other stations about its synchronous packet generation time and the minimum declared time is used as the basic LOCOST parameter. The used LOCOST method is applied to the asynchronous traffic offered load by affecting the service policy of the respective buffer, following the model described further.

Each station calculates the number of the connected network stations and the minimum packet generation time. From the constant packet length and the medium transmission speed, the station calculates the length of a 'traffic window'. The 'traffic window' represents the latest part of the traffic and is used to monitor the network conditions and to estimate the scheduling of the asynchronous offered load.

Assuming that the network speed is M bits/sec, the packet length is d bits long, and R_{max} is the fastest bit rate of the synchronous service, then the minimum packet generation time g_{min} is given by:

$$g_{min} = \frac{d}{R_{max}} \tag{1}$$

while the 'traffic window' length (w) is given by:

$$w = \frac{g_{min} M}{d} = \frac{M}{R_{max}} \tag{2}$$

The 'traffic window' is used as an indirect expression of the available bandwidth and for that reason the LOCOST handles it as if it could be divided into two parts: the first part to be devoted to the synchronous traffic only and the second part to be used to support both classes of traffic. The portion of the bandwidth available for each part is determined by a threshold value (t_h), a first approximation of which is given by:

$$t_h = w - \alpha \frac{N}{4} \tag{3}$$

where N is the number of stations connected to the network and α is a dumping factor in order to have a more realistic scenario. The α factor is an expression of the event that normally all stations do not have synchronous traffic to transmit and that synchronous services have not the same bit rate with the highest rate service. The number of stations is divided by 4 because the network uses two counterdirectional rings, so each station transmits its packets on the ring that provides the shortest path to its direction, and on the average a packet goes through the half of the shortest path.

Each station counts the number of asynchronous packets transmitted during the 'traffic window' and, if this number is equal or greater to t_h, it prohibits its asynchronous traffic buffer to transmit a packet, even if it gains the right to transmit into the network. If the number of asynchronous packets is lower than t_h, the asynchronous traffic is served with a variable probability which depends on the used traffic scheduling algorithm.

In Figure 3, three different traffic scheduling algorithms are shown. In the first case, the LOCOST uses a step-like function, where the asynchronous packets are served with probability one if the asynchronous traffic is less than the predetermined threshold. This algorithm gives a portion of the available bandwidth to synchronous traffic, while the rest of the bandwidth is used by both classes of traffic as in the normal access method. In the second case, the asynchronous packet service probability is a linear function of the asynchronous traffic; it becomes one when there is no such traffic in the network and approaches zero when the asynchronous traffic is near the threshold. This algorithm also gives a portion of the available bandwidth to the synchronous traffic but the rest of the network capacity is used by the asynchronous traffic of each station partially, depending on the current load of the total asynchronous traffic. As the total asynchronous traffic increases, each station decreases its transmission probability, while if the total traffic has been decreased each station transmits with higher probability. In the last case, which seems to be more fair and efficient, the asynchronous packet service probability is also a linear function of the asynchronous traffic, but becomes one when the asynchronous traffic is less or equal to the half of the threshold. This case is similar to the previous one, with the following differences; first, when the total traffic load is low, the LOCOST does not affect the access method and second, because the slope of the traffic scheduling is more sharp, the algorithm adapts the asynchronous traffic faster to the total traffic requirements.

It must be mentioned that the value of the 'traffic window' is updated in each packet time interval by using a fixed duration timer, irrespective if there is traffic into the network.

The proposed LOCOST method has the following differences from the originally proposed by Limb:

- The traffic is monitored on a per packet basis instead of a time period, resulting to more efficient adaptation to the network requirements.
- It does not keep each particular class of traffic to no more than a specific percentage of the total capacity of the medium, but distributes the available bandwidth between them, according to the load requested by the high priority traffic.

In the next section some simulation results of the application of the LOCOST method to a Buffer Insertion Ring are given. For these simulations we have considered that the traffic is uniformly distributed between the stations and that each station transmits its packets to every other station with equal probability. This consideration is very important for destination release networks, where each station does not monitor all the traffic but on average only the half.

IV. THE SYSTEM PERFORMANCE ANALYSIS

In this section the simulation parameters are given and the respective results are presented. Initially, we describe the network performance in terms of interarrival times in the receiving station under heavy load conditions and we determine the best LOCOST algorithm for this purpose. We then examine the influence of the dumping factor on the performance of the previously determined algorithm. Finally, we indicate the way the LOCOST affects the asynchronous traffic and the overall network performance.

The Simulation parameters

Our simulation model considers a Metropolitan Area Network, based on optical fiber with transmission speed at 500 Mbits/sec. 50 stations are connected to the network, which covers an area of 50 km, thus the time needed for a packet to traverse two stations is on the average 5 μsecs. The stations transmit constant length packets in the network. The packet length, d, is equal to 1000 bits, which consists of I bits of information ($I = 960$ bits, in our model), 16 bits for the source address, 16 bits for the destination address and 8 bits for the packet control field (the packet type is included in this field). The packet transmission time equals 2 μsecs.

Two traffic scenarios are mainly considered in the simulation model, the first has a 'light' synchronous throughput (20% of the total station throughput) and the second has a 'heavy' synchronous throughput (80% of the total station throughput), while the rest of the station throughput is for the asynchronous data.

In each station the asynchronous traffic consists of variable length file transfers, where the mean length is 10 kbytes and the length distribution is exponential. The whole file is considered to be available for transmission immediately after its generation time in the station. The file generation procedure is a Poisson process, with mean value equal to λ, where λ takes the values 200 files/sec and 50 files/sec for the 80/20 and 20/80 scenarios respectively. The synchronous traffic consists of connections established in time instants according to a Poisson process, while their duration is uniformly distributed between 7 and 9 secs. The synchronous data stream has a constant bit rate (CBR) of R_i bits/sec. Thus, synchronous packets are available for transmission every I/R_i secs.

The Simulation Results

Figure 4 shows the packet interarrival time distribution of a synchronous class of traffic at a receiving station, both for a simple Buffer Insertion Ring network and for the same network with our LOCOST protocol. The offered load is approximately 1 Gbit/sec and the synchronous traffic is 80% of the total. The fastest synchronous bit stream is 17 Mbits/sec and the minimum packet generation time is 56 μsecs. From these curves, it is obvious that the performance of the simple network can not support synchronous traffic properly, because the packet interarrival times are spread over a large area around the mean value. On the other hand, the packet interarrival times in a network with our LOCOST protocol are concentrated around the mean value and with this modification, a Buffer Insertion Ring can be used to support synchronous traffic. (In this Figure, the LOCOST uses the case '2' scheduling of traffic algorithm). Assuming that a packet is rejected when the interarrival time between this packet and the previous one is greater than twice the packet generation time, there is no packet rejection using the LOCOST protocol, while for the simple network operation the packet rejection rate is $0.5 \cdot 10^{-3}$. According to [2], for compressed video transmission the acceptable packet error rate must be less than 10^{-9}. It is obvious that a simple Buffer Insertion Ring cannot support services like compressed video.

In Figure 5, the distribution of the interarrival times for the three scheduling algorithms are shown. In these curves the dumping factor α is equal to 1, which means that the portion of the 'traffic window' devoted solely to the synchronous traffic is equal to a quarter of the number of the connected stations. The best performance is achieved when the third traffic scheduling algorithm is used, because it has the fastest adaptation rate to the total traffic conditions.

In the case where the number of stations is large enough and the minimum packet generation time is small, most of the bandwidth which is devoted to the synchronous traffic is not used effectively. This is the reason we have introduced the dumping factor α. In Figure 6, the interarrival time deviation is given as a function of the dumping factor for the second and the third traffic scheduling algorithms. As it is shown, as the available bandwidth explicitly used by the synchronous traffic decreases (as the dumping factor decreases), the interarrival time deviation increases and the interarrival times are spread at a larger area around the mean value. Although the interarrival times' deviation increase, their values remain in acceptable limits and the synchronous traffic is supported properly.

When the synchronous traffic covers the major part of the offered load, the network utilization remains high and approximates 390% of the network capacity. This value has been measured when the dumping factor was set to 0.8, the total offered rate was 2 Gbits/sec and 65% of the total bandwidth was available to be used only by the synchronous traffic. Under the same conditions, the network utilization drops to 270%, when the asynchronous traffic covers the major part of the offered load and only 35% of the network capacity could be used by the asynchronous traffic. This happens because the threshold value of the 'traffic window' depends only on the number of the connected stations and is irrespective of the traffic conditions. A possible solution for this problem is to determine the threshold value periodically using the synchronous traffic statistics.

From the above results, it can be concluded that the application of the LOCOST protocol to a Buffer Insertion Ring, which uses constant length packets, although decreases slightly its performance, makes it appropriate to support synchronous class of traffic.

V. CONCLUSIONS

In this paper, we have described the application of the Load-Controlled Scheduling of Traffic concept to Buffer Insertion Ring networks, in order to make them appropriate to support integrated services in the Metropolitan Area Environment. Each network station has been considered to support two classes of traffic, asynchronous and synchronous, uses different buffers for each class and the information transmission is carried out using small and constant length packets. The use of constant length packets, in combination with the pipeline structure of the ring buffer, guarantees a constant transmission delay to the packets of the same message and the end-to-end delay is affected mainly by the variable queueing at the source station due to the upstream network traffic.

In order to override these delay variations and to achieve bounded transfer delays, we have applied a modified LOCOST protocol to each network station. The LOCOST measures the percentage of the asynchronous traffic to the total network bandwidth and adjusts the asynchronous offered load accordingly, in order to provide the synchronous traffic with the required bandwidth. Various algorithms have been used in the LOCOST protocol and their influence to the network performance has been analysed.

The performance of a Buffer Insertion Ring has been measured with LOCOST and without it and, as the results show, a significant performance improvement can be achieved using the LOCOST protocol with the appropriate traffic scheduling algorithm. Although the total network throughput decreases as the LOCOST is applied, this performance decline is not significant and is overriden by the network ability to support bounded transfer delays for synchronous classes of traffic.

References

[1] European Computer Manufacturers Association: "Requirements for Access to Integrated Voice and Data Local and Metropolitan Area Networks", *ECMA TR/51*, June 1990.

[2] D.B. Hehmann, M. G. Salmony and H.J. Stüttgen: "High-Speed Transport Systems for Multimedia Applications", *IFIP Workshop on Protocols for High-Speed Networks*, North Holland, May 1989, pp. 303-321.

[3] W. Bux and M. Schlatter: "An Approximate Method for the Performance Analysis of Buffer Insertion Rings", *IEEE Trans. on Communications*, Vol. COM-31, No 1, January 1983, pp. 50-55.

[4] I. Rubin and J. E. Baker: "Media Access Control for High-Speed Local Area and Metropolitan Area Communications Networks", *Proceedings of the IEEE*, Vol. 78, No 1, January 1990, pp. 168-203.

[5] John O. Limb: "Load-Controlled Scheduling of Traffic on High-Speed Metropolitan Area Networks", *IEEE Trans. on Communications*, Vol. 37, No 11, November 1989, pp. 1144-1150.

[6] Z. Tsai and I. Rubin: "Performance of Token Schemes Supporting Delay-Constrained Priority Traffic Streams", *IEEE Trans. on Connunications*, Vol. 38, No 11, November 1990, pp. 1994-2003.

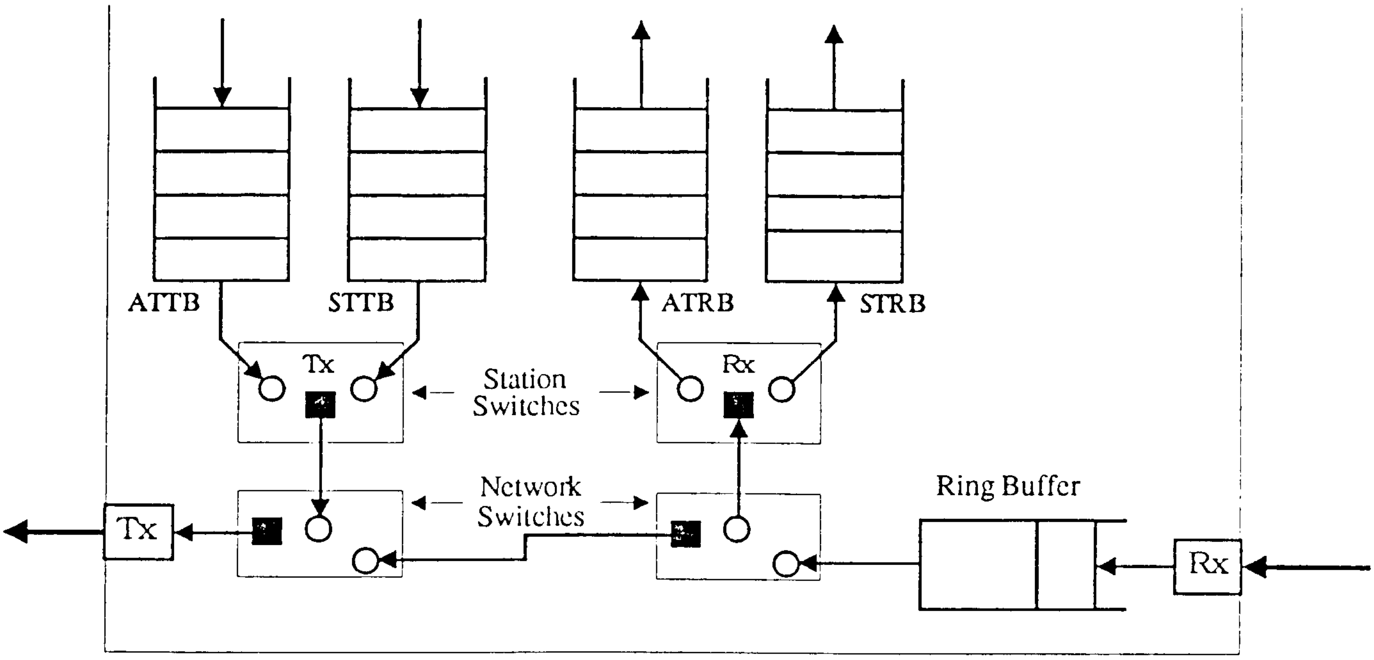

Fig. 1. The model of the ring access method
ATTB: Asynchronous Traffic Transmission Buffer
STTB: Synchronous Traffic Transmission Buffer
ATRB: Asynchronous Traffic Reception Buffer
STRB: Synchronous Traffic Reception Buffer

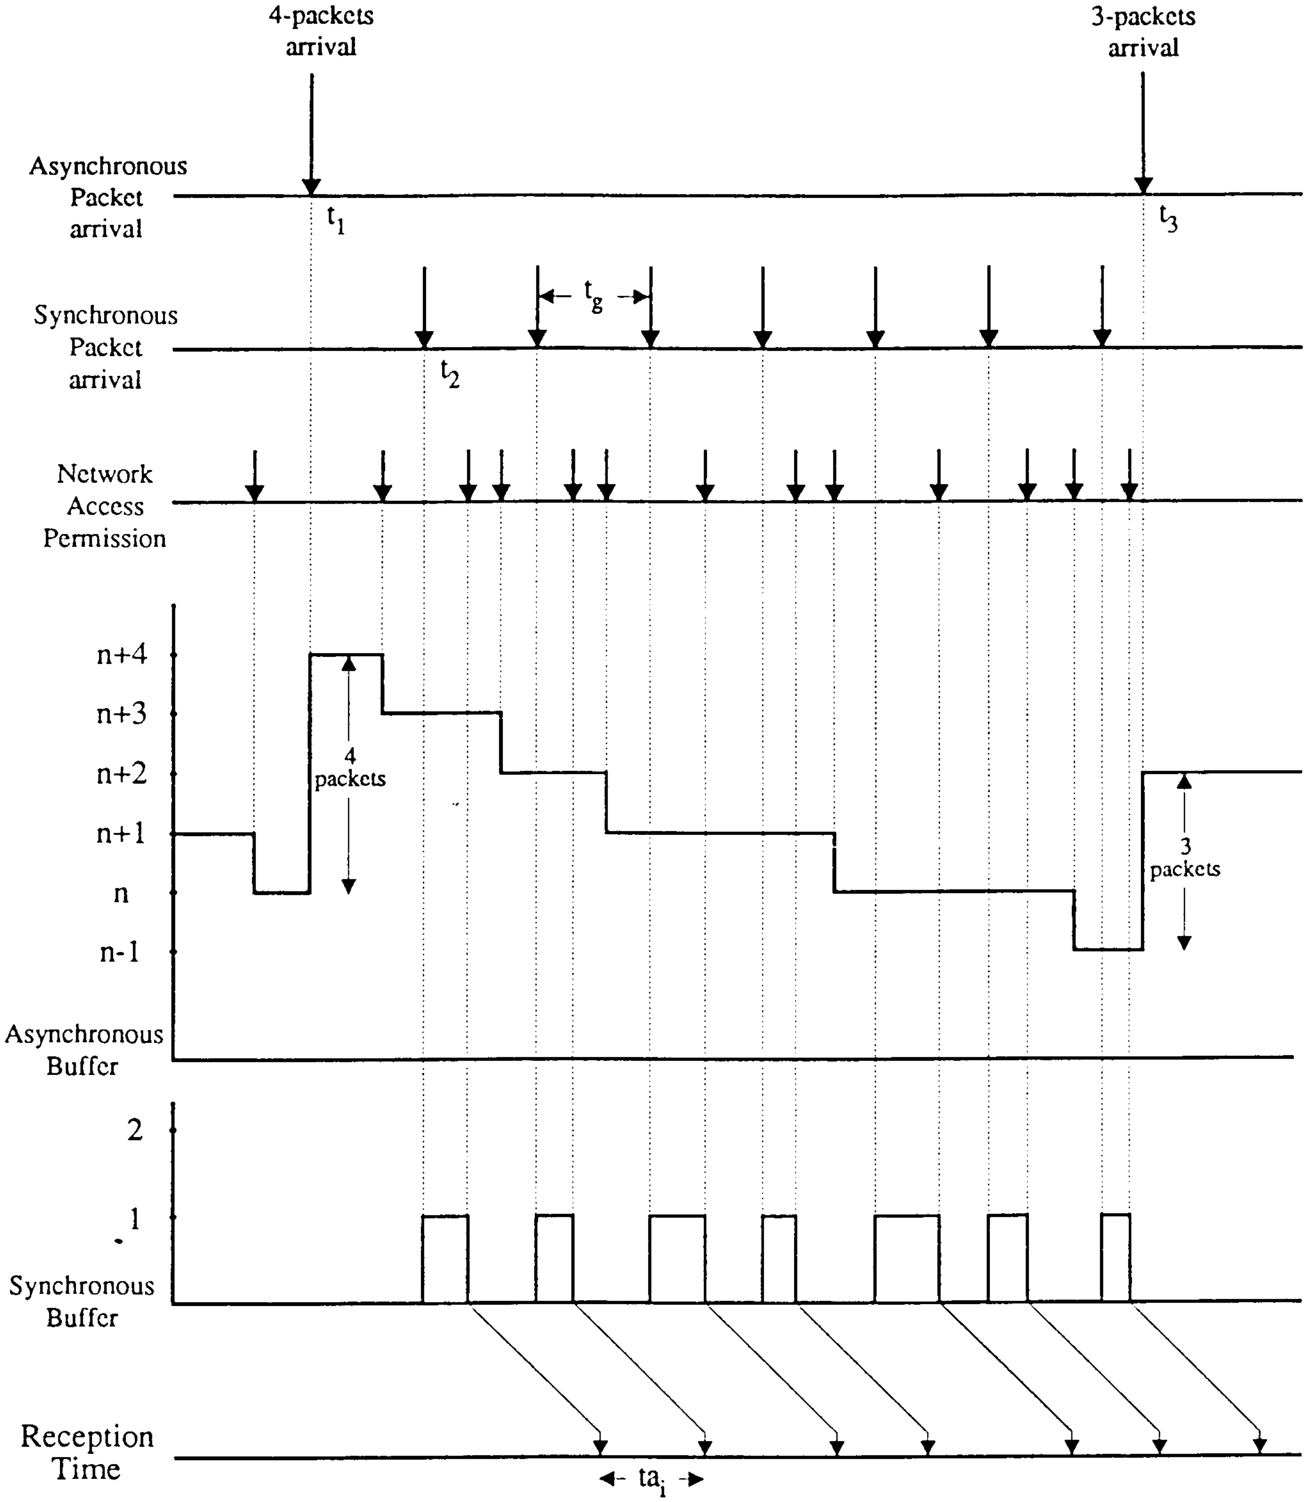

Fig. 2. Typical timing diagram for station operation.

t_1, t_3 : Asynchronous transfer requests (4 packets and 3 packets, respectively)

t_2 : Synchronous transfer request with constant bit rate

t_g : Inter-arrival packet time at the transmitting station (synchronous traffic)

ta_i : Inter-arrival packet time at the receiving station

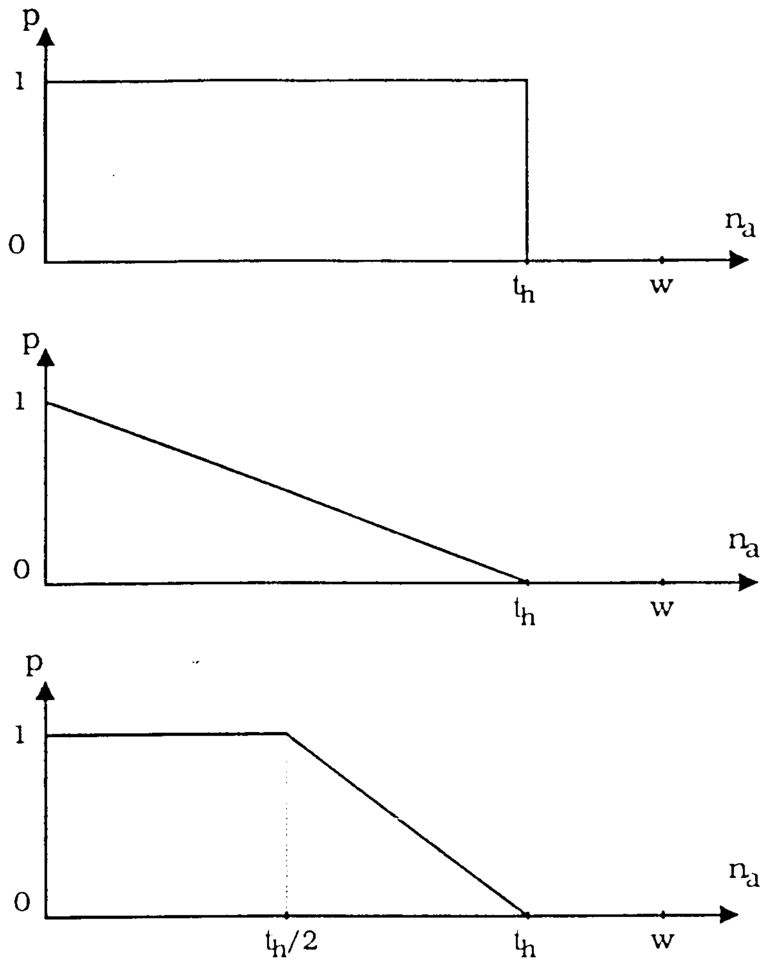

Fig. 3. The LOCOST traffic scheduling algorithms.

w = window length

$t_h = w - \alpha * N / 4$

p = asynchronous transmission probability

n_a = asynchronous packets in the window

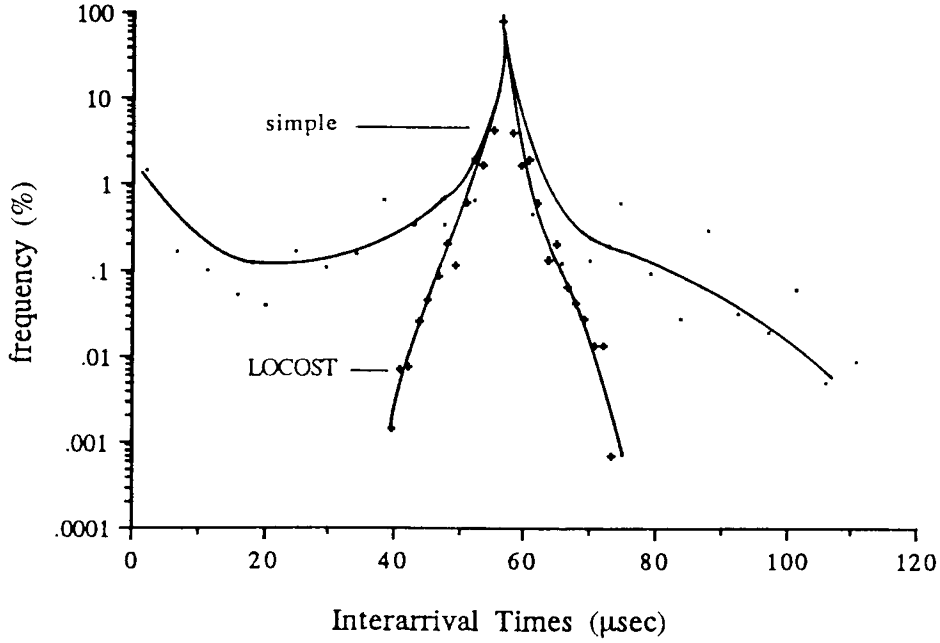

Fig. 4 The distribution of the interarrival times of synchronous traffic at a receiving station for (a) the simple Buffer Insertion Ring and (b) the modified one using the LOCOST protocol.

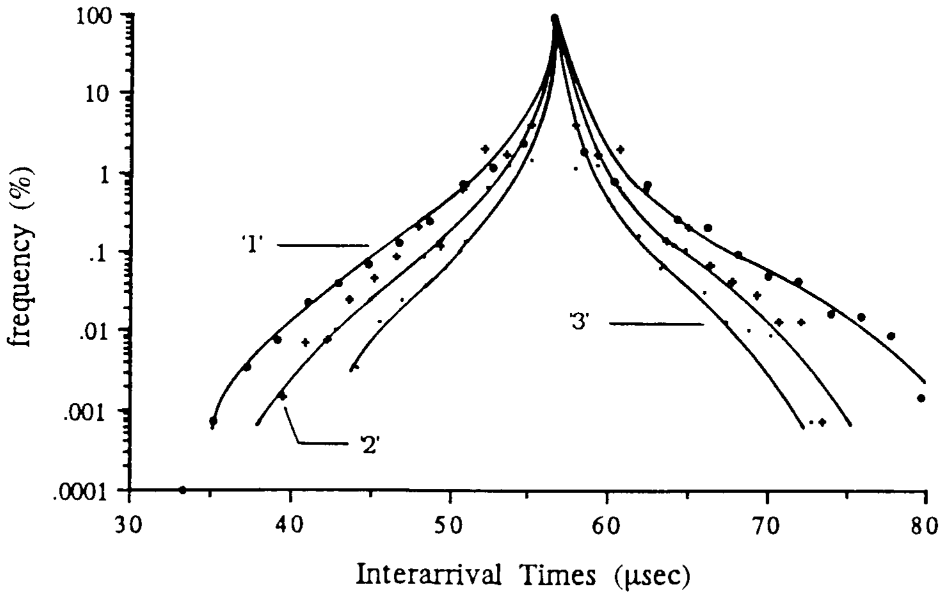

Fig. 5 The three traffic scheduling algorithms of the LOCOST protocol.

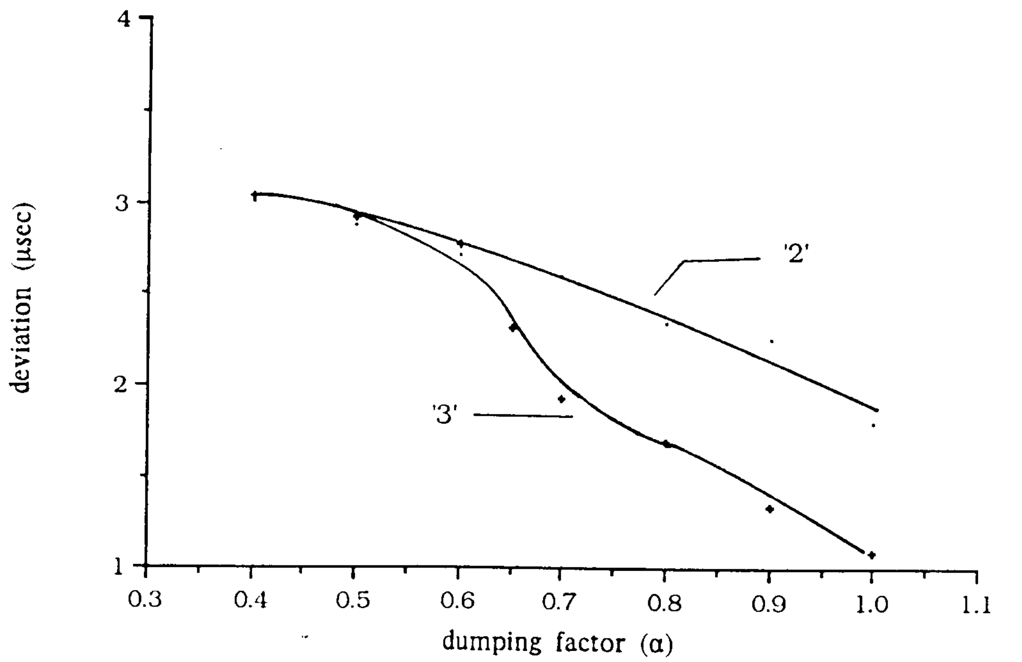

Fig. 6 The interarrival times deviation versus the scheduling algorithm dumping factor.