

Reprint

Efficient Bandwidth Allocation in Buffer Insertion Ring Networks

J. Koutsonikos, T. Antonakopoulos and V. Makios

The 5th Canadian Conference on Electrical and Computer
Engineering

TORONTO, CANADA, SEPTEMBER 1992

Copyright Notice: This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted or mass reproduced without the explicit permission of the copyright holder.

"EFFICIENT BANDWIDTH ALLOCATION IN BUFFER INSERTION RING NETWORKS"

J. Koutsonikos, T. Antonakopoulos and V. Makios

Laboratory of Electromagnetics, Department of Electrical Engineering
University of Patras, 26500 Patras, Greece

Abstract

A new access method for the high-speed LAN environment is proposed in this paper. The access method is based on the Buffer Insertion Ring protocol using fixed length packets in unslotted operation. The method also uses a simple scheduling of traffic to adapt the offered load to the network conditions and to distribute the available bandwidth in different classes by monitoring the instantaneous traffic flow. The synchronous traffic has the highest priority and is served with bounded transfer delays by managing the allocated bandwidth as it is allocated to the different traffic classes.

1. Introduction

The support of various traffic classes in the LAN environment requires the use of high transmission speed, small and constant length packets and the adaptability of the service provided by the access method to the characteristics and constraints of the supported traffic class. The networks of today are very specialized and suffer from a large number of disadvantages, like service dependence, inflexibility, inefficiency etc. [1]. In the LAN/MAN environment, new access methods have already been proposed and developed. These methods can support various traffic classes but their disadvantages are their high cost and their complexity arisen when they are connected with a B-ISDN network.

The Buffer Insertion access method was initially considered to support data exchange in the LAN environment [2], [3]. When new 'real-time' services were developed and new performance requirements were to be satisfied, the use of the Buffer Insertion method was not proven to be appropriate without basic modifications. The stations of a Buffer Insertion Ring 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. The packets are removed from the network by the destination station achieving an increased network throughput, which approaches 400% in dual ring structures. Its main disadvantage is the variable transfer delays occurring due to the intermediate ring buffers, resulting to inefficient support of synchronous class of traffic.

The initially proposed Buffer Insertion method did not include any mechanism to allocate the network bandwidth according to specific service requirements, but each station transmitted, if it had a packet to transmit, without taking into

account the behavior of the residual stations.

As John O. Limb has shown in [4], the management of the network load using traffic measurements, like the proposed Load-Controlled Scheduling of Traffic (LOCOST), can be used to improve the network performance and probably to satisfy the requirements of 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 propose a network access method which is based on the Buffer Insertion method and includes traffic monitoring functions to fulfil the Quality of Service (QOS) requirements of various traffic classes. In section 2, we describe the access method in relation to the basic station architecture while in section 3 details of the traffic monitoring function as well as the load control algorithms are presented. Finally, the ability of the new access method to support synchronous traffic is studied and some results are discussed.

2. The Access Method

In the integrated services LAN environment three classes of traffic are usually considered, the isochronous, the synchronous and the asynchronous class [5]. The isochronous traffic, which has the most strict requirements, requires constant delay while the synchronous traffic requires bounded transfer delays. The asynchronous traffic has no delay limit and usually is supported by the network bandwidth not used by the other two classes of traffic. In our discussion, the isochronous traffic will be considered as a special case of the synchronous traffic, where the delay bounds approach the mean delay value by using a type of elastic buffer at the receiver side.

Considering two classes of traffic, synchronous and asynchronous, the access mechanism devotes a Service Access Point (SAP) to each class and a First Come First Serve queue is related with each SAP. 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 can transmit only one packet per access.

The network under consideration uses small and fixed length packets with a destination release method. Although the destination release scheme increases the MAC complexity, it provides improvement in delay and throughput levels, allowing spatial reuse of the available bandwidth. 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 packet transmission, as will be shown later. The network uses a protocol architecture like the one described in [6], where variable length connectionless MAC protocol data units are segmented into multiple fixed length packets. For this reason, the protocol profile contains a segmentation/reassembly sublayer to implement these functions. From now on our discussion will be concentrated on the access method itself, and its services provided to the upper layers, and we will not discuss the functionality of the various submodules of each traffic class at the various layers.

In Figure 1, the architecture of the Access Control Layer is given. The 'Network Buffer' (NB) is of constant length and a pipeline structure. When there is no station transmission, the Network Buffer acts as a constant delay line providing also buffering utilities for various station functions. This operation of NBs which are inserted in the traffic flow, form a type of Buffer Insertion Ring network. The NB length is related with the fixed packet length and is equal to the packet length plus the length of the packet header. The NB stores temporarily the incoming traffic and inserts a constant delay, irrespective if there is a station transmission or not. This delay is added in each passing-through packet, so a packet transmitted from the source station will arrive at the destination station with constant delay, irrespective of the network traffic conditions. The total packet delay is composed of two parts, the 'queuing' delay and the 'passing-through' delay. The 'queuing' delay is determined as the time elapsed from the reception of a packet at the Tx-Buffer up to the transmission of the packet into the network. The 'passing-through' delay is determined as the time a packet needs to pass through all the intermediate stations from the

source to the destination station. The 'passing-through' delay d_{pt} is constant and is given by the following formula:

$$d_{pt} = \frac{(2L_H + L_I)}{M} n + \sum_{i=0}^n l_i / u$$

where L_H is the packet header length, L_I is the length of the packet information field, M is the transmission speed in the medium, n is the number of stations which are between the source and the destination station, u is the signal transmission speed in the physical channel and l_i is the distance between adjacent stations where the source station is marked with 0 while the destination station is marked as $n+1$. In long interstation connections the delay imposed by the 'Network Buffer' is small compared to delay caused by the point-to-point connections. In Figure 2 a network operation example is given. The delay of the packets of a synchronous traffic from station A to station E has a constant factor, affected by the Network Buffers of the three intermediate stations and a variable factor of the 'queuing' delay at the source station. The jitter of the 'queuing' delay will determine the ability of the network to support synchronous traffic.

Each traffic class has one dedicated buffer for reception and one for transmission. These buffers are controlled by the 'Load Controller and Traffic Scheduler (LCTS)' unit. The LCTS receives information about the 'Network Buffer' status from the 'Packet Validator' unit and controls the flow of information inside the station based on some load control algorithms. The 'Packet Validator' is controlled by the Layer Man-

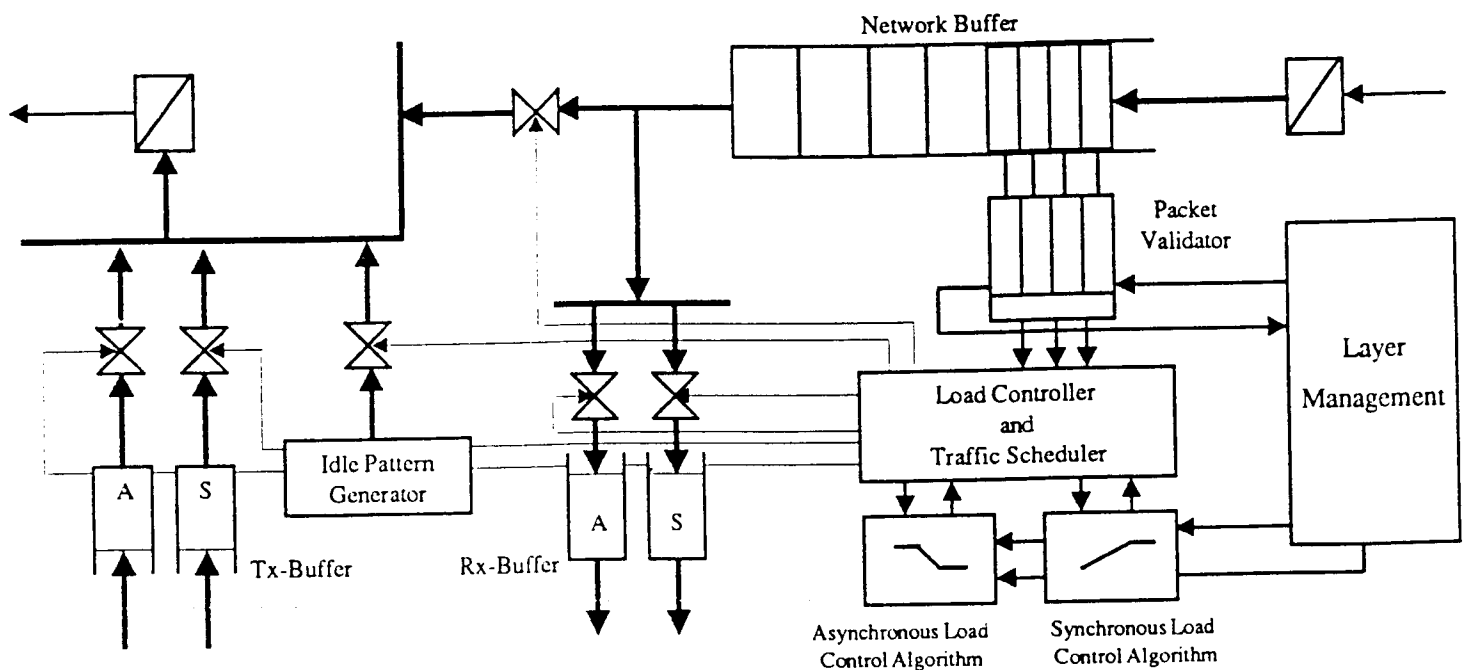


Fig. 1. The Architecture of the Access Control Unit.




Fig. 2. The Buffer Insertion Ring Architecture.

agement unit and is reconfigurable with Content Addressable Memory (CAM) modules. The Packet Validator communicates with the Layer Management using a synchronous serial interface for parameters update. The 'Packet Validator' has also a pipeline structure and its length is equal to the length of the packet header. When a new packet is inserted in the Network Buffer, the 'Packet Validator' specifies its characteristics and informs accordingly the LCTS unit. The specified characteristics of a packet are its type (synchronous or asynchronous), its destination (if this is the destination station), the existence of errors in the header, etc. If the packet was intended to this station the packet is removed from the 'Network Buffer' and is stored in a receive buffer (the asynchronous or the synchronous) according to its type.

When the station has a packet to transmit, examines the Network Buffer and, if it is occupied for less than the packet header length, the station starts transmission immediately. If there are packets both in the asynchronous and the synchronous traffic buffers, the synchronous buffer has the highest priority. The mechanism used in the LCTS to decide if the station will transmit or not and from which output queue, is described in details in the next section. The LCTS also uses two modules to estimate the bandwidth allocation in the two different traffic classes and to provide its estimation to the decision mechanism. In order to keep the receiver of each station always locked on the transmission frequency of the previous station and to avoid the use of preamble in front of each packet, an 'Idle Pattern Generator' is used in the output section of each station. When there are no data to be transmitted either from the internal buffers or from the 'Network Buffer', the 'Idle Pattern Generator' transmits a predefined idle pattern continuously.

The problems which arise in order to support synchronous traffic are due to the delay variation caused by the upstream traffic inside the network. In order to minimize the effect of the network traffic to the delay variations of the same synchronous stream, a traffic monitoring method is used to control the type and the amount of the offered load. For these re-

asons, we have adapted the LOCOST method in order to create an access method which is capable to support synchronous traffic class and to be easily implementable in hardware.

3. The Traffic Scheduler and the Load Control Algorithms

The offered-load control algorithm used in this access method is based on traffic monitoring and is fully distributed because each station decides on its offered load independently from the others and by using only its traffic monitoring statistics [7]. This control algorithm is based on the basic concept of LOCOST which states that the appropriate scheduling of traffic can improve the network performance and if it is properly managed, it can change the characteristics of the access method.

Before starting the description of the load control algorithm, we will examine the way the synchronous packets are generated. 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. During the network initialization, each station informs the others about its synchronous packet generation time and the minimum declared time is used as the basic algorithm parameter. The load control method is based on the collected traffic statistics during a 'traffic window' period, measuring the bandwidth allocation for both synchronous and asynchronous traffic and by affecting the buffer service policies following the model described in the next paragraphs. The control method gives the requested bandwidth to the synchronous traffic by varying the asynchronous offered load. When more synchronous packets are transmitted, more bandwidth is allocated for use by the synchronous traffic only. When the portion of the transmitted synchronous packets to the total traffic decreases, the algorithm assumes that more bandwidth is available for asynchronous traffic, and the asynchronous offered load increases.

After the initialization phase each station knows the number of the connected network stations as well as 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 portion of the traffic under consideration. Assuming that R_{max} is the fastest bit rate of the synchronous services, the 'traffic window' length W_1 , is given by:

$$W_1 = M / R_{max}$$

The 'traffic window' which is an indirect expression of the available bandwidth, is divided in two parts: the first part is devoted to the synchronous traffic only and the second part is used to support both classes of traffic. The portion of the bandwidth available for each part is determined by a threshold value T_3 , which is given by:

$$T_a = W_1 - \alpha N/2$$

where N is the number of stations connected to the network and α is a dumping factor. The α factor is determined by the synchronous load control algorithm and is used to adapt the bandwidth allocation to the current traffic conditions. The number of stations is divided by 2 because the network uses the destination release method to remove the packets from the network so, on the average, a packet goes through the half of the connected stations.

Each station counts the number of asynchronous packets transmitted during the 'traffic window' and, if this number is equal or greater to T_a 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_a the asynchronous traffic is served with the algorithm shown in Figure 3. The asynchronous packet service probability P_a is a linear function of the asynchronous traffic, but becomes one when the asynchronous traffic is less than one half of the threshold. The algorithm gives a portion of the available bandwidth to the synchronous traffic exclusively, 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.

The synchronous load control algorithm, which is shown in Figure 4, is used to adapt the bandwidth allocation to the current traffic conditions. As long as the participation of the synchronous traffic to the total traffic decreases, the dumping factor decreases and the available bandwidth for the asynchronous traffic increases. When there are more than $N/2$ synchronous packets in the 'traffic window', the portion of the bandwidth allocated to the synchronous traffic takes its maximum value. As it was explained previously, each station monitors half of the network traffic so when it sees that the synchronous traffic is equal or greater than $N/2$, it means that on the average all the connected stations transmit synchronous data.

It must also be mentioned that the value of the 'traffic window' is updated in each packet time interval, either by the passing-through packet or by using a fixed duration timer.

In the next paragraphs some simulation results concerning the ability of the proposed access method to support synchronous traffic are presented. In our simulation model 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. In table I are given the used simulation parameters [7]. Figure 5 shows the packet interarrival time distribution of a synchronous class of traffic at a receiving station. 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 offered-load control are concentrated around the mean value and it can be used to support synchronous traffic. It becomes appar-




Fig. 3. The Asynchronous Load Control Algorithm.

- P_a : Packet Transmission Probability
- N_p : Number of Asynchronous Packets in the traffic window
- W_1 : Window Length
- T_a : Asynchronous Threshold




Fig. 4. The Synchronous Load Control Algorithm.

- α : Dumping Factor for the asynchronous load control algorithm.
- N_s : Number of Synchronous Packets in the traffic window
- N : Number of Stations

TABLE I: Simulation Parameters

Number of stations: 50
Transmission speed: 500 Mbit/s
Interstation distance: 1 Km
Packet length: 125 bytes
Packet header: 5 bytes
Traffic scenario: Heavy loaded with synchronous data (80% of the total traffic).
Synchronous data bit rate: Constant (CBR) - 17Mbit/s




Fig. 5. The distribution of the interarrival times of synchronous traffic at the receiving station for simple and load controlled network operation.

ent that the application of load control methods to the MAC layer results to a performance like a "resonant filter" around the value '1'. 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 control algorithm, while for the simple network operation the packet rejection rate is $0.5 \cdot 10^{-3}$. According to [5], 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.

4. Conclusion

In this paper, we have described the application of the Load-Controlled Scheduling of Traffic concept to Buffer Insertion Ring networks to make them appropriate to support integrated services. 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 queuing at the source station due to the upstream network traffic.

These delay variations can be smoothened in order to achieve bounded transfer delays by applying offered-load control using traffic statistics. The asynchronous offered load is adjusted in order to provide the required bandwidth to the synchronous traffic. The network performance was measured and significant performance improvement was achieved when the load control algorithm was applied.

References

- [1] Martin de Pryckes: "Asynchronous Transfer Mode: Solution for broadband ISDN", Ellis Horwood Ltd., 1991.
- [2] 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.
- [3] P.A. Pietrzyk and A. Shaout: "New Message Based Priority Buffer Insertion Ring Protocol", *Electronics Letters*, 7th November 1991, Vol. 27, No 23, pp. 2106-2108.
- [4] 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.
- [5] 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.
- [6] H. Ohnishi, N. Morita, S. Suzuki: "ATM Ring Protocol and Performance", *IEEE International Conference on Communications*, Boston, MA, June 1981, pp. 394-398.
- [7] T. Antonakopoulos, J. Koutsonikos and V. Makios: "Bounded Transfer Delays in Buffer Insertion Rings using Load-Controlled Scheduling of Traffic", *Fifth IEEE Workshop on MANs*, Taormina, Italy, May 1992.