A New Access Control Protocol For High-Speed ATM LANs

Martin Götzer

University of Stuttgart
Institute of Communications Switching and Data Technics (IND)
Seidenstraße 36
D-70174 Stuttgart
email: goetzer@ind.uni-stuttgart.de

Abstract

Protocols for shared media Local Area Networks (LAN) are optimized for high asynchronous traffic throughput. Multimedia applications, e.g. interactive video, require low traffic delay with low delay jitter. Peak bit rate reservation for this kind of traffic guarantees the required quality of service but usually results in a waste of bandwidth due to the bursty nature of the traffic produced by multimedia applications.

In this paper, a new MAC protocol for ring LANs is presented which integrates isochronous and asynchronous traffic. Fixed-length frames are divided into slots that are designated for carrying ATM-sized cells. By reserving the same set of slots in every frame, a constant bandwidth can be provided for realizing isochronous channels. All slots that are not reserved for isochronous channels are available for asynchronous traffic. In addition to that, asynchronous traffic may use slots that are reserved for but currently not used by isochronous channels. In this way, the protocol combines lossless data transfer and a high degree of network utilization.

1 Introduction

Protocols for shared media Local Area Networks (LANs) are facing three major challenges: The growing demand for high throughput to the desktop, upcoming multimedia applications which require low traffic delay with low delay jitter at high data rates, and integration of ATM-type traffic which will be applied in Wide Area Networks. Traditional LANs, i. e. Ethernet or Token Ring, do not fulfill the requirements to meet these challenges. Even FDDI is not suitable for high-speed ATM-like traffic: its efficiency decreases when the medium bitrate increases since only a single node can send data at a time [1].

High-Speed LAN protocols use slots as "data containers" which are filled with fixed size data [2]. All nodes can send data concurrently by using free slots. If fairness is desired, a medium access control (MAC) protocol must be applied to share the bandwidth between nodes according to a given traffic shape. Several proposals have been made for HSLAN MAC protocols. An overview can be found in [2], while in [3-6] a selection of protocol definitions can be found.

If isochronous channels are needed, slots are grouped into frames. For an isochronous channel, a slot is reserved in each frame which is characterized by its position within the frame. Thus the node accessing the isochronous channel is guaranteed to get an empty slot within a period

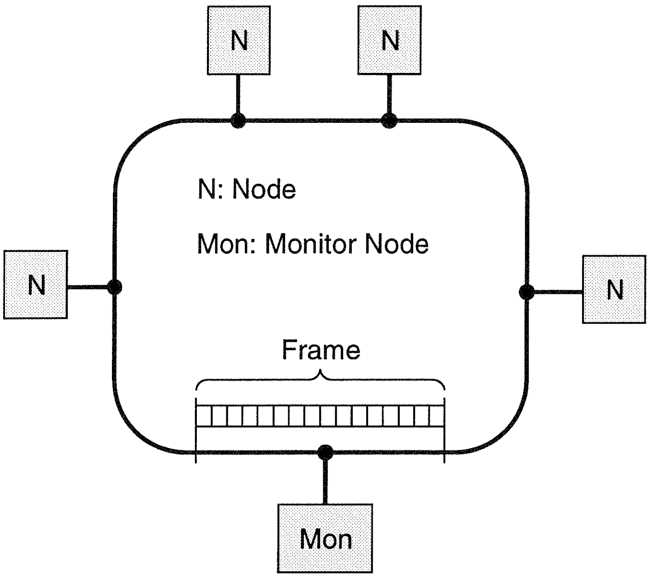

Figure 1: Frame-based Ring Structure

which is equal to the frame duration. If a node dosn't use its reserved slots, the bandwidth will be lost. Compressed video, for example, can be transmitted easily with high quality by using isochronous channels which can hold the peak bit rate, but a high percentage of the overall bandwidth is wasted due to the bursty nature of the traffic.

In this paper a new MAC protocol is presented which allows usage of reserved bandwidth currently not being used by isochronous channels. The protocol combines the benefits of peak bit rate reservation and the efficient usage of the overall available bandwidth.

2 Architecture Concept

Figure 1 shows a frame-based slotted ring as it is used in this protocol. The monitor node plays the role of the ring master, providing functions for reservation, supervision, and error recovery. Frames consist of a frame header and a fixed number of slots (Figure 2). Each slot can hold a cell with a size of 53 octetts.

Two traffic classes are defined: The Isochronous Traffic Class (ITC), and the Asynchronous Traffic Class (ATC). In the ITC, the monitor node grants slots to requesting nodes. These slots are determined by their position in the frame. Peak bitrate reservation is used, but nodes can leave reserved slots empty. Slots that are not currently occupied by the ITC can be used by the ATC.

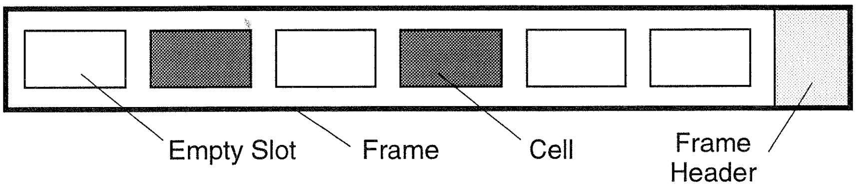

Figure 2: Frame Format

2.1 Slot Format

Figure 3 shows the slot format. The length of the slot is 53 octets. The Access Control Field (ACF) is used by the MAC protocol presented in Section 2.2 and is shown in Figure 4. The Busy/Free, Available, Active Node, and Reset bits are associated with the slots, while Syn/Asyn, Packet, EOP, and Mon bits are cell specific and only valid in a busy slot which is determined by its Busy/Free bit set.

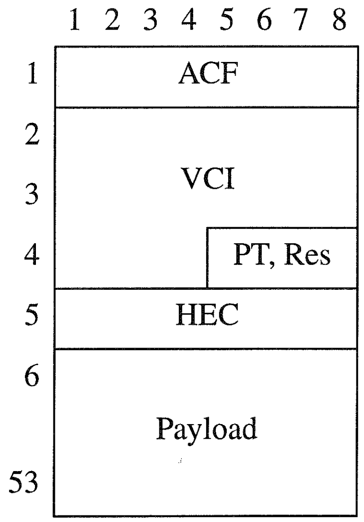

ACF:	Access Control Field
VCI:	Virtual Channel Identifier
PT:	Payload Type
Res.:	Reserved
HEC:	Header Error Check

Figure 3: Slot Format

Busy/ Free	Available	Active Node	Reset	Syn/ Asyn	Packet	EOP	Mon

Busy/Free = 1: Busy Slot
0: Empty Slot
Available = 1: Slot available for ATC
0: Slot unavailable for ATC
Active Node = This bit is set by an active node
and reset by the monitor node
Reset = 1: Reset command issued by the monitor node
0: No Reset
Syn/Asyn = 1: STC cell content
0: ATC cell content
Packet = 1: Cell is part of a packet
0: Single cell
EOP = 1: Last cell of packet
0: Not last cell of packet
Mon = Monitor bit

Figure 4: ACF Format

2.2 Basic Medium Access Protocol

Figure 5 shows the node structure for cell handling. Each node has a list of "usable slots in frame" to determine which slots in each frame are currently reserved for this node's ITC connections. The node owns these slots for the time the ITC connection is established, i. e. it decides on how the slots are to be used. Slot Reservation is handled by the Monitor Node. The node can decide whether the own slots can be used for ATC traffic by any node within the next ring cycle by setting the Available bit in the slot. Incoming cells in the node's slots - which can only be of ATC type - are stored in the node's insertion buffer. Unreserved slots are always available for ATC traffic.

To avoid an insertion buffer overflow in case of heavy ATC traffic, an Available Counter is introduced. The counter value represents the sum of all the node's slots that are marked "available" on the ring, in addition to the number of cells in the insertion buffer. The maximum value of this counter is the number of cells the insertion buffer can hold. If the Available Counter reaches the maximum value, the node's outgoing slots are marked as "unavailable", i. e. the Available bit is not set. Slots with the Available bit cleared must not be used for ATC cells. This mechanism ensures that the insertion buffer cannot overflow.

An incoming cell is copied into the appropriate receive buffer if the the node is a receiver for the particular cell. Cells are deleted by the destination node, which is the last receiver in case of multicast. By means of these independent mechanisms, broadcasts as well as multicasts can be realizsed in an efficient way.

If the node is allowed to send an ATC cell and if the insertion buffer is not empty, a cell from the insertion buffer is sent. The node can send a cell from its ATC transmit buffer if the insertion buffer is empty. Keeping the Available Counter value as low as possible, this simple insertion buffer priorization scheme guarantees maximum overall throughput. Alternative schemes can be considered where nodes can send cells from their ATC transmit buffer even if the insertion buffer is not empty. This may lead to a shorter waiting time in the transmit buffer, but the overall throughput may decrease in case the Aavailable Counter reaches its maximum value.

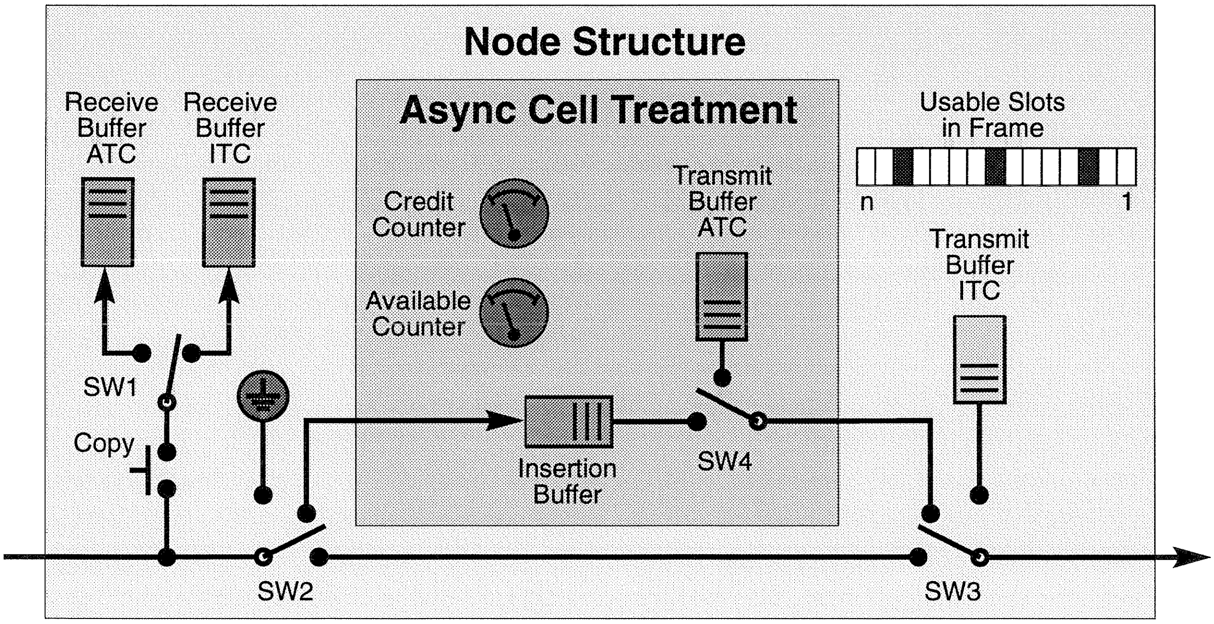

Figure 5: Node Structure

2.3 ATC Access Control

As a fairness protocol for the Asynchronous Traffic Class, the ORWELL protocol [3] is a simple but effective solution. It is a credit-based mechanism where all nodes get a fixed - but not necessarily the same - amount of credits at the beginnig of a cycle. A credit is required to send a cell. If all nodes have used up all their credits or if none of the nodes wants to send a cell, a reset command is issued by the Monitor Node which marks the beginning of the next cycle. The time between two reset commands can be comparably long because all timing critical traffic is sent in the Isochronous Traffic Class, leading to a high overall ATC throughput.

For the fairness protocol a Credit Counter is introduced in each node. The Credit Counter initial value is the amount of credits the node gets in each cycle. As long as a node's credit counter value is not zero and the ATC transmit buffer is not empty, the node is called "active node". If its insertion buffer is not empty a node is called "active node" as well, not depending on the credit counter value. An "active node" sets the Active Node bit of any outgoing slot. It can put an ATC cell in any outbound slot that is marked "available" and "free" as described in Section 2.2. Each time a cell is sent from the ATC transmit buffer, the Credit Counter is decreased by one. A node whose credit counter is zero is not allowed to send data from its ATC transmit buffer, enabling downstream nodes to use the empty slots.

The Monitor Node resets the Active Node bit in every slot. Since all "active nodes" set the Active Node bit, a slot with the Active Node bit unset only reaches the Monitor Node if there are no more active nodes on the ring. In this case, the monitor node issues a reset command by setting the Reset bit of an outgoing slot. All nodes set their Credit Counter to the initial value when a slot with its Reset bit set is observed.

2.4 Packet Transmission

With the "Packet" and "EOP" bits, optional packet transmission in the ATC is enabled where nodes can use consecutive ATC slots. A packet will not be interrupted by any other ATC cell, thus enabling simple packet reassembly without sophisticated buffering mechanisms in the receiver.

All cells of a packet are marked by setting the Packet bit. In the last cell of each packet, the End Of Packet (EOP) bit is additionally set. If packet transmission is enabled and the node has taken a cell from the insertion buffer with the Packet bit set, it must send a cell from the insertion buffer whenever the outbound slot is marked "available" and "free" until the EOP bit is detected which marks the last cell of a packet. For the transmit buffer the same scheme is to be applied if a packet is sent by the node.

When transmitting a packet, the insertion buffer must be able to hold an additional number of cells. Logically, the insertion buffer can be seen as two insertion buffers in series. One of them is used for the protocol described in the previous sections, the other is used to store the incoming ATC cells during packet transmission. This insertion buffer portion must be able to store a complete packet before the next packet transmission can take place.

3 Conclusion and Outlook

In this paper, a new Medium Access Control protocol for multimedia HSLANs was presented. Compared to other HSLAN proposals with integrated isochronous and asynchronous channels, the protocol makes use of the bandwidth that is not currently used in the isochronous channels.

Currently a complete protocol specification and evaluation is under development. In addition the protocol is implemented at the IND using standard programmable logic at a link bit rate of 960 Mbps.

Optimization of the ATC access control protocol is subject for further study. Other protocols will be investigated for their ability to replace the simple ORWELL protocol on top of the basic medium access protocol.

4 References

[1] *Tangemann, M.:* "MODELLIERUNG UND ANALYSE VON LOKALEN HOCHGESCHWINDIGKEITSNETZEN MIT ZEITGESTEUERTEN TOKEN PASSING-MEDIENZUGRIFFSPROTOKOLLEN", University of Stuttgart, IND, 1994.

[2] *Lemppenau, W. W.; Götzer, M.; Seibold, W.:* "ACCESS PROTOCOLS FOR HIGH-SPEED LANS AND MANS", Proceedings of the 11th Conference on European Fibre Optic Communicatons and Networks, The Hague, Netherlands, June 1993, pp 253-259.

[3] *Falconer, R. M.; Adams, J. L.:* "ORWELL: A PROTOCOL FOR AN INTEGRATED SERVICES LOCAL NETWORK", British Telecom Technology Journal, Vol. 3, No. 4, October 1985, pp. 27-35.

[4] JTC1/SC6 WG6 N71, "SPECIFICATION OF THE ASYNCHRONOUS TRANSFER MODE RING (ATMR) PROTOCOL", ISO/IEC JTC1 Information Technology, Secretariat SNV, Japanese National Body, Japan, Version A1.0, September 1990.

[5] *Watson, G.; Tohmé, S.:* "S++: A MAC PROTOCOL FOR GBIT/S NETWORKS", Proceedings of the 2nd Télécom Paris /Universität Stuttgart Joint Workshop on High Speed Networks, Stuttgart, Germany, September 1991, pp. 1-4.

[6] *Van As, H. R.; Lemppenau, W. W.:* "CRMA-II: DYNAMIC EMBEDDING OF ISOCHRONOUS CHANNELS", Proceedings of the 12th Conference on European Fibre Optic Communicatons and Networks, Heidelberg, Germany, June 1994.