To be presented at 17th Conference on Local Computer Networks, Minneapolis, MN, Sept. 1992.

April 1992
Revised July 1992

LIDS-P-2103

A Method for Isochronous Traffic in a WDM Star Network

Ghassan Semaan *

Laboratoire de Téléinformatique
Ecole Polytechnique Fédérale de Lausanne
El - Ecublens, 1015 Lausanne - Switzerland
Email : semaan@ltisun.epfl.ch

Pierre Humblet †

Lab. for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge 02139 - USA
Email : humblet@lids.mit.edu

Abstract

The need for high speed networks is continuously increasing, resulting in many proposals for new access protocols. Amongst these, Wavelength Division Multiplexing *(WDM) is one of the most promising. Due to the physical characteristics of WDM, only the star architecture is considered practical. Despite the great number of proposed access protocols based on this architecture, isochronous traffic was not, or badly, managed. In this paper, we present a new access protocol that handles isochronous traffic in a much better way. This protocol allows the stations to reserve part of the bandwidth for their isochronous traffic while the rest is used for asynchronous traffic. In addition, this technique requires only one transmitter and one receiver per station and is free from the constraint of a control channel.*

1 Introduction

Extremely high bandwidth applications such as connecting thousands of users for videoconferences, HDTV, transfer of big files, etc. exceed the capacities of traditional LAN technology. Relatively new access protocols for fiber-optic networks such as FDDI [1], [2], DQDB [3], and ATM [4] have increased the throughput of LANs from around ten Mb/s to over one hundred Mb/s. Recently, networks with as much as 1 Gb/s, (e.g. CRMA [5]) have been proposed. However, these proposals fail to exploit the maximum possible bandwidth of the fiber, which can be as high as 30 THz [6]. Approaching this rate is difficult due to the limitations of the electronic components, which cannot yet be avoided. One solution is to use *Wavelength Division Multiplexing* (WDM), which is sometimes called *Frequency Division Multiplexing* (FDM)

In WDM, the data of the stations are transmitted on different wavelengths of light. Each wavelength is considered as an independent channel and stations usually have to process only the data travelling on one channel. Thus the limitations of the electronic components constrains the bandwidth of only a portion of the whole. Laboratory prototypes using this technique have achieved throughputs as high as a few hundreds of Gb/s, with each station's processing rate around 1 Gb/s [6].

Most proposed WDM protocols use a star architecture, the bus architecture being unsuitable for large networks due to attenuation in the couplers [7]. In its basic form, the stations, wired to a central coupler, transmit their data on a given wavelength. The coupler mixes the transmissions on the different wavelengths and broadcasts them back to the stations. An optical filter is used to select the wavelength to be received.

Since the coupler is a passive optical device, it cannot function as a master device managing the selection of wavelengths for transmission and reception. The stations must therefore decide amongst themselves how to handle this task. In order to do so, almost all proposals dedicate a special wavelength, λ_0, to play the role of a *Control Channel* [8]-[12]. When needed, stations send control slots on this channel informing the destination stations of the transmission and indicating on which wavelength the communication will take place. The choice of the wavelength to be used can be random or subject to special algorithms [11, 12].

There are several disadvantages to this solution. First, each station needs four devices, two transmitters and two receivers, to be able to read and write on the control and one data channel simultaneously. Second,

*Part of this research was conducted when the author was visiting MIT.

†This research was supported by ARO under contract DAAL03-86-K-0171.

the throughput of the network depends heavily on that of the control channel: as more stations try to send data, only those who successfully transmit their control information will be able to transmit on the data channels, even though the bandwidth of the network can handle a much larger number of communications. This problem becomes more dramatic when the control channel is accessed in a random manner (ALOHA or slotted ALOHA protocols). Finally, even though access schemes on the control channel may handle asynchronous[1] data in an adequate manner, they do not offer good support for the isochronous[2] case. Since many future applications such as video conferences, HDTV, and so on, will generate isochronous data, it is important that the network be able to handle both types. One possible solution for this problem is to dedicate a whole channel for an application that needs to transfer isochronous data. This solution will suffer from a great loss of bandwidth as very few applications need 1 Gb/s, which is the usual channel rate in WDM. Therefore the channels should be shared by different applications generating either asynchronous or isochronous data.

The purpose of this paper is to present a new access scheme to control data transmission in a WDM star network. In this protocol, asynchronous and isochronous data simultaneously and dynamically share the same channels. Moreover, this access scheme requires only one tunable transmitter and one tunable receiver and does not need a control channel, contrary to the main protocols in [13]. In the following sections, we outline the basic properties of the network, give a detailed description of the access scheme, and discuss some aspects of the data transmission.

2 Network Description

2.1 Physical Architecture

The network has the star architecture proposed in figure 1. Each station is assigned an independent wavelength and is connected to a passive star coupler. The wavelengths are mixed at the coupler and broadcast to all stations. Only one tunable transmitter is needed by each station to transmit its data and one tunable receiver to select the wavelength to be read. Both devices are capable of tuning on all wavelengths. If all the low loss regions of a single-mode fiber are used, up to 1000 stations can use the network simultaneously.

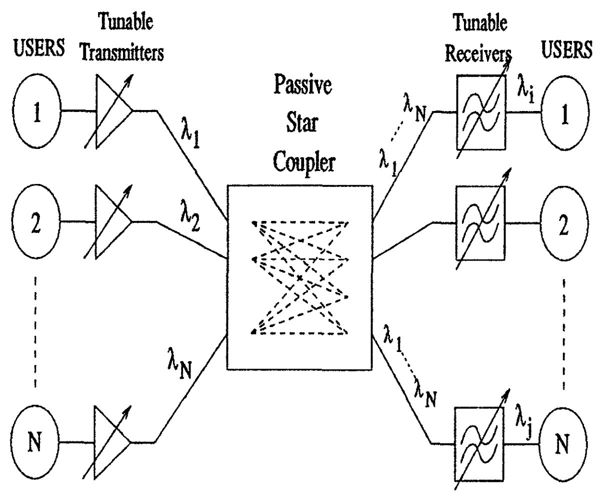

Figure 1: The star architecture

2.2 The Frame

Data are transmitted in slots. Slots are multiplexed into frames as the one presented in figure 2. Frames can have, for example, a duration of 125 μs to facilitate interconnection with public switched networks.

Each frame is made of (N+5) slots. The first five are called *Control Slots* (CS) and serve for the transmission of control information. The remaining N are called *Data Slots* (DS) and serve for the transmission of data.

The DS can be of two types : *Reserved* and *Free*. When Reserved, a DS can be accessed by only one station to transmit isochronous data. When Free, it can be accessed by any station to transmit asynchronous data. For reasons that will be explained in section 4, the first M DS are always Free and their type cannot be modified. The last (N-M) DS can be Reserved or Free and their type is indicated in the control slot *Status*. The values of M and N are implementation dependent.

Each station will deal with two different types of frame. The first, seen by the transmitter, consists of the transmitted slots. Thus, it is called *Transmitter Frame* (TF). The second, seen by the receiver, consists of the received slots and is called *Receiver Frame* (RF). The difference between these two frames is due to the

[1] Asynchronous data : data that is send at irregular period of time. In this case the bandwidth can be shared statistically between different stations. These data are generated typically by applications such as distributed operating systems, Network File Systems (NFS), ...

[2] Isochronous data : data that is send at regular period of time and needs guaranteed bandwidth to meet performance requirements. These data are generated by applications such as video and voice communications

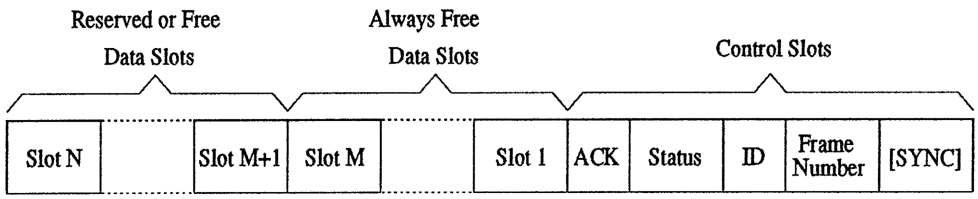

Figure 2: The structure of the frame

fact that the stations do not transmit and receive data on the same wavelength. Only the CS are transmitted on the dedicated wavelength of a station. The DS are transmitted on the wavelength of the destination station (made possible by the use of a tunable transmitter). On the other hand a station tunes its receiver to the different wavelengths to read the CS of any station then tunes back to its dedicated wavelength to read the data slot sent to it. This approach is illustrated in figure 3. In this figure stations 1, 2 and 3 are assigned wavelengths λ_1, λ_2 and λ_3 respectively (the frames are simplified to one CS and four DS). If we consider station 1 for example, this station will send its control slot (slot with a "C" is the middle) on its dedicated wavelength, λ_1, then tune its transmitter to wavelength λ_2 to send slots 1 and 3 to station 2 and to wavelength λ_3 to send slot 4 to station 3. For reception, it will tune its receiver to λ_2 to read the control slot of station 2 then tune back to its dedicated wavelength (λ_1) to read the data slots sent to it by station 2 (slots 1 and 4) and station 3 (slot 3). An equivalent behaviour is executed by stations 2 and 3.

For ease of reference we define we define the *Source* as the transmitting station and the *Destination* as the station to whom the data is addressed.

It should be noted that a station does not necessarily read the CS of only one station each frame duration. For any frame, it can for example read the Status slot on λ_i and the ACK slot on λ_j with $i \neq j$.

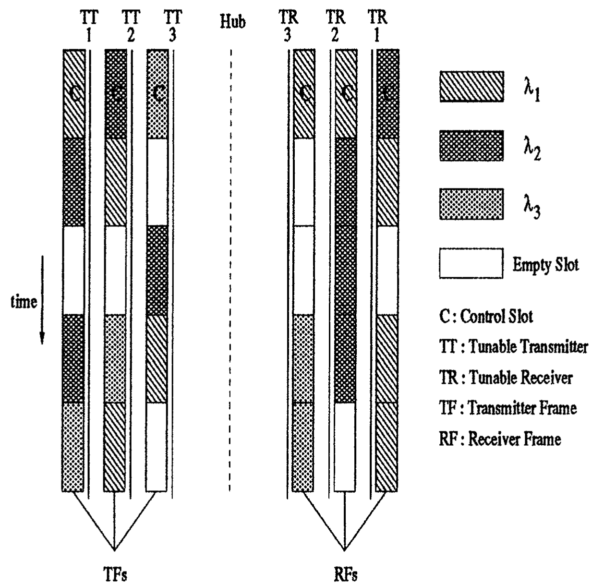

Figure 3: The Transmitter and Receiver Frames

3 The Access Scheme

In this section we will describe the access scheme used to control transmission on the network. To economize space only the fundamental parts are shown.

The principal components of the access scheme are the CSs. CSs are exploited in different ways depending on the type of data to send. Three different types can be handled by the protocol. In addition to the asynchronous and isochronous ones introduced earlier in this paper, we add the Bulky Data. Bulky Data is a special case of the asynchronous type, one where the amount of data to send from one station to a destination is big. Therefore, three types of transmissions are possible, one for each type of data. We will explain first how the CS work, then we will show how they are used with these different types of data.

3.1 The Control Slots

Before any new transmission, the station will read the Status slot of the Destination. The Status slot indicates the state (Reserved or Free) of the DS in the TF and RF. Based on this, the Source checks which slots are *Useful Slots*. A Useful Slot (US) is a slot that is free in the TF of the Source and the RF of the Destination simultaneously. For example, if station A has slots 3, 4 and 8 free for transmission and station B has slot 4, 7 and 8 free for reception, the US for a transmission from A to B are 4 and 8. Once the US have been checked, the station will choose the slots to use for its transmission.

However, due to the data rate and size of the net-

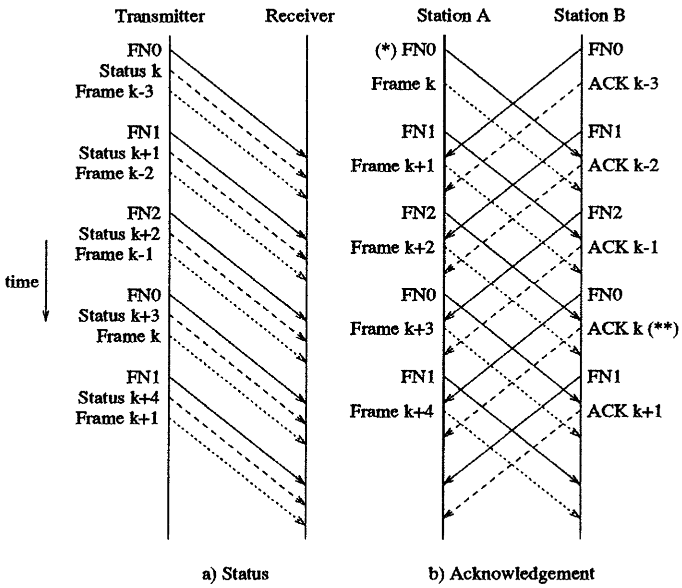

Figure 4: Status and ACK, (R=3);

work, the Status slot cannot be used to represent the state of the DSs of the physical frame to which it belongs. It cannot because the transmitters will complete transmission of the DS of a frame before the receivers are able to read its Status. There must therefore be a delay between the transmission of the Status and its corresponding DSs, large enough to insure that all receivers have received the Status before any of the DS is transmitted. This delay can be insured by introducing the *Cycles*.

Let T be the propagation delay from the transmitter to the receiver of the point that is at the maximum distance allowed from the hub. We define the *Cycle* as consisting of R frames where R is the smallest integer bigger than (1 + T/(frame duration)). In the CS *Frame Number* (FN), each frame is given an incremental number modulo R that indicates the position of the frame in the Cycle. Therefore, to be useful, the Status should indicates the type of the DS of the frame with the same FN but in the next cycle. In this case, all the receivers will receive the Status before any of the DS concerned are transmitted.

In figure 4.a, a timing diagram showing how the Status works is given. In this example R=3. In frame FN0, the transmitters will send the status of the DSs to be sent three frames later, which is the next cycle. When the Status reaches the receivers, none of the DSs of the next FN0 frame has been sent. Therefore, using the Status, the stations can decide in which slot they can send their data.

The CS *ACK* is used to acknowledge the reception of the DS. Its use is almost similar to the Status and is explained in detail in section 3.2.

There remain two CS that have not yet been presented: the *SYNC* and the *Identification*. The *SYNC* slot is used for the synchronization of the network. When using the slot switching technique with WDM, the alignment of the slots becomes a major problem if different stations are sending on the same wavelength. In [14] we solved this problem using the SYNC and gave a structure for the slots. The *Identification* (ID) slot identifies the stations that is using a given wavelength. It should be read periodically to allow the stations to maintain a table that indicates which wavelength is used by which station.

3.2 Asynchronous Data

The access for asynchronous data is based on the slotted ALOHA protocol [15]. After checking which are the US, the station will randomly choose amongst them the number necessary for the transmission, and will then transmit its data. If two stations choose the same slot, a collision will occur and both slots will be lost. As the stations do not read the slots they have sent, they are not able to detect the collisions. Therefore, the Destination should acknowledge the reception of the DS by using the ACK.

For each DS received without an error the Destination will set the corresponding bit in the ACK. Once all the DS of a frame are received, the Destination will send the ACK slot in the next frame with the same FN as the current frame. As with the Status, the ACK cannot be used to acknowledge the DS belonging to the same physical frame. ACK acknowledge the DS received in the frame with the same FN but in the previous cycle. Therefore, the stations, remembering in which frame and which slot they have send their data, read two cycles later the corresponding ACK slot and get the results of their transmissions. In figure 4.b a timing diagram resumes how the ACK works. In frame FN0, station A send the data of frame k (referenced with a "*"). Once all the DS have reached station B, this station will acknowledge the reception of the slots in the ACK of the next FN0 frame (referenced with a "**").

The duration of the Cycle was chosen to be bigger than the maximum propagation delay so that it is independent of the distances of the stations from the hub. The (+ 1) is to allow the stations to receive all the data slots before sending the ACK. This because if we only consider the propagation delay, a station may not have finished the reception of the DS of the frame prior to transmission of a new frame with the same FN (cf. figure 4.b). The slots that are not acknowledged are retransmitted.

Due to the fact that only one Status and one ACK slot can be read each frame, only one asynchronous transfer can be initiated in each frame duration. However, with a frame duration of 125 μs, this would allow eight thousands asynchronous transmissions per second for each station.

3.3 Isochronous Data

Isochronous transfer takes place in reserved slots. If the available bandwidth obtained by the US that can be reserved is sufficient, the station will initiate an isochronous communication by sending an *Iso_Conn_Request* (Isochronous Connection Request) in an asynchronous manner. With the Iso_Conn_Request, the station will indicate the numbers of the slots that it wishes to reserve for its communication. Moreover the Source will reserve a US in its Receiver Frame, the *Answering Slot*, to be used by the Destination to answer the request. Following this, the Source will modify its Status accordingly. It should be noted however, that the slots that are requested by the Source for isochronous transfer but not yet confirmed, can still be used by the Source for asynchronous transfer until the confirmation of the connection.

After the transmission of the Iso_Conn_Request three events may happen :

1. Collision : The request will collide with data transmitted in the same slot by another station. The Source will detect the collision by reading the ACK of the corresponding frame two cycles later. In this case, the Source restarts the process.

2. Change in Status : The Status of the Destination may change during the propagation of the Iso_Conn_Request (cf. figure 5.b). In this case the Destination will send an Iso_Conn_Reject in the Answering Slot indicating the motives of the rejection. After the reception of the rejection, the Source may or may not restart the process.

3. Successful Transmission : The request reaches the Destination without a collision and the slots requested by the Source can be reserved. In this case the Destination will send an Iso_Conn_Confirm in the Answering Slot confirming the reservation. Following this the Destination will modify its Status accordingly (cf. figure 5.a).

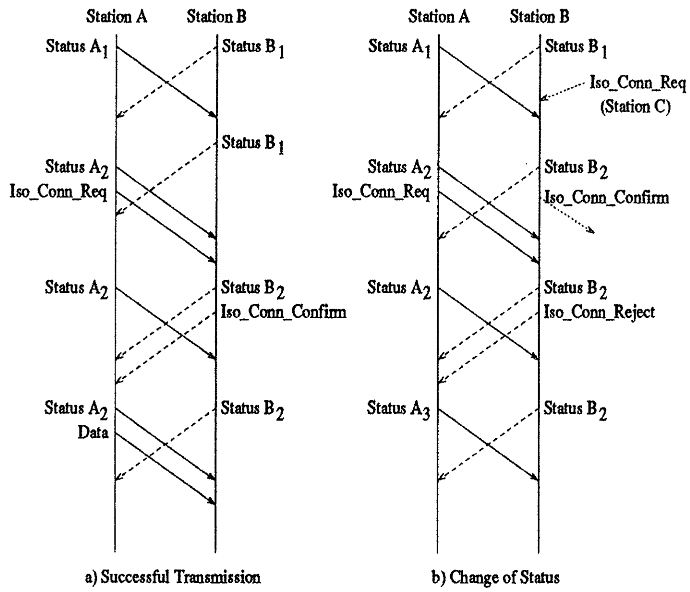

Figure 5: Establishing a isochronous communication

At the reception of the Confirm, the Source will start transmitting its data only in the reserved slots. Once the communication is finished, the Source will send an Iso_Disconnect.

Contrary to the asynchronous data, the slots that are transmitted in an isochronous manner are not acknowledged. This is due to two facts: the first is that in most isochronous applications, acknowledgement is useless, and the second is that isochronous traffic does not get involved in collisions. This allows the stations to have multiple isochronous communications simultaneously.

3.4 Bulky Data

Bulky data, as previously stated, is a special case of asynchronous data, where the amount of data to send from one station to a destination is big. Unless the load of the network is low, transmitting this type of data in a slotted ALOHA protocol will decrease the performances of the network due to the rise in the number of collisions. Therefore we transmit Bulky Data in an isochronous manner.

Once a Bulky Data transfer is requested, the station will randomly choose the necessary slots for its transmission amongst the US that can be reserved. Then it will send a *Bul_Conn_Request* (Bulky Connection Request) in an asynchronous manner. The establishment of the communication is identical to the isochronous one with one small difference.

If, during the propagation of the request, the Destination modifies its Status, then it will confirm the reservation of the slots that are still free and does not reject the connection as its does with the isochronous communications. The Source can afterwards decide if the number of slots confirmed are sufficient for its transmission. If not, the Source may try to reserve the remaining necessary slots or may send a disconnect primitive. Once the connection is established, the Source will send its data only in the reserved slots.

As the isochronous data, Bulky Data are not acknowledged in this protocol. This may take place in a higher layer. In [16] different studies of acknowledgement techniques for high speed networks are presented. This option was chosen to allow the stations to have multiple simultaneous bulky communications in addition to the isochronous and asynchronous transmissions. Moreover, with Bulky communications the slots are free from collisions. Therefore, few slots need retransmission and high level acknowledgement can have a very good performance.

4 Discussion

In this section we discuss the options chosen for the data slots. Those chosen for the control slots have been discussed as they have been presented throughout the paper.

As we would like the three different types of data we defined (i.e. asynchronous, isochronous and bulky) to share the same channels simultaneously, a slotted protocol was chosen. A slotted protocol gives us the opportunity to define different types of slots, mainly Reserved and Free. The reserved slots will be used to guarantee a minimum bandwidth for the stations with isochronous data. This guarantee is a mandatory requirement for isochronous transfer. The remaining slots, those that are not reserved, are used for the asynchronous transfer.

In order to always be able to reach any station, we have to guarantee a minimum number of slots that will always be Free. This is necessary, amongst other things, for the network management. This bandwidth is obtained with the first M slots in each frame.

Due to the physical characteristics of the passive optical star, managing the bandwidth is a major problem. For asynchronous transfer we decided to choose a random access which, in this case, is the simplest access. As our network is slotted, slotted ALOHA protocol was chosen. The performances of the slotted ALOHA protocol has been heavily studied in the literature [17, 18], therefore we will not discuss them again here. What should be noticed however, is that this type of protocol offers very good performances (real throughput, access delay, ...) when the load of the network is low. Therefore it will be used only when the stations have few slots to send. For the transmission of a bigger amount of data, we introduced the bulky type.

An acknowledgement for the asynchronous data is necessary because the Sources are not able to detect the collisions by themselves. In the example we show, we presented a very simple type of acknowledgement : Collision or No Collision. However, the ACK can be used to inform the Source of other possible problems. The Destination may loose a slot for reasons such as loss of synchronization, a buffer that is full, etc. By using two bits instead of one for acknowledging each slot, the Destination can indicate for the Source the different reasons for the loss. Depending on the reason, the Source may or may not delay the retransmission of the slot.

The bulky data are transmitted in reserved slots to reduce the amount of data transmitted in a slotted ALOHA manner. This is due to the fact that under heavy load the performances of a slotted ALOHA protocol degrades. As with asynchronous data, the example given is very simple. A lot of improvement can be made for this type of data. For example, with the connection request, the Source can indicate the minimum and maximum number of slots it needs for its transmission. The Destination will reserve continuously the minimum number requested and, depending on the bandwidth available and on the requests from other stations, will reserve an additional variable number of slots. The *Adaptive Bandwidth Allocation* technique is actually under development and will handle this type of problems. A notion of priority can be also introduced.

As for the isochronous transfer, a *Rearrangement* technique is under development. This technique will handle the cases where the Source and the Destination have sufficient bandwidth for an isochronous communication, but the free slots in each station do not coincide. Therefore a modification of the reserved slots is necessary in order to increase the number of US.

5 Conclusion

We presented in this paper the control access scheme of a new protocol for transmitting data in a passive optical star network. This protocol supports three different types of data, asynchronous, bulky and isochronous sharing simultaneously the same channels. The method presented does not require a Control Channel and uses only one tunable transmitter and one tunable receiver. We introduced the *Cycles* with the Status and the ACK and showed how these notions allows the stations to have a view of the state of the network despite the passive nature of the hub. Then we gave an example of the transmission for each of the three types defined. Finally we discussed the different options chosen. We are in the process of evaluating the performance of this protocol through simulation. Further work remains to be done to provides additional functionalities.

References

[1] ANSI. *FDDI - token ring media access control (MAC).* Technical Report X3.139-1987, ANSI, November 5, 1986.

[2] ANSI. *FDDI - hybrid ring control.* Draft Proposed X3T9.5/89-43 REV5.0, ANSI, February 24, 1989.

[3] IEEE. *Distributed Queue Dual Bus - metropolitan area network.* Unapproved Draft P802.6-/D14, IEEE, July 13, 1990.

[4] CCITT. *Integrated Services Digital Network (ISDN).* Recommendations I.121, Melbourne 1988.

[5] M. Nassehi. CRMA : an access scheme for high speed LANs and MANs. *IEEE Int. Conf. Communications*, 1697–1702, April 1990.

[6] C.A. Brackett. Dense Wavelength Division Multiplexing Networks : Principles and Applications. *IEEE Journal on Selected Areas in Communications*, 6(8):948–964, August 1990.

[7] P.S. Henry. High Capacity Lightwave Local Area Network. *IEEE Comm. Mag.*, 27(10):20–26, October 1989.

[8] I.M. Habbab, M. Kavehrad, C.E. Sundberg. Protocols for Very High-Speed Optical Fiber Local Area Networks Using a Passive Star Topology. *Journal of Lightwave Technology*, LT-5(12):1782–1793, December 1987.

[9] N. Mehravari. Performance and Protocol Improvements for Very High Speed Optical Fiber Local Area Networks Using Passive Star Topology. *Journal of Lightwave Technology*, 8(4):520–530, April 1990.

[10] G.N.M.Sudhakar, N.D. Georganas, M. Kavehrad. Slotted Aloha and Reservation Aloha Protocols for Very High-Speed Optical Passive Star Topology. *Journal of Lightwave Technology*, 9(10):1411–1422, October 1991.

[11] M-S. Chen, N.R. Dono, R. Ramaswami. A Media-Access Protocol for Packet-Switched Wavelength Division Multiaccess Metropolitan Area Network. *IEEE Journal on Selected Areas in Communications*, 6(8):1048–1056, August 1990.

[12] Y. Ofek, M. Sidi. Design and Analysis of a Hybrid Access Control to an Optical Star using WDM. *pp 2A.3.1 - 2A.3.12 Infocom'91*, 1991.

[13] P.A. Humblet, R. Ramaswami, K.N. Sivarajan. An Efficient Communication Protocol for High-Speed Packet-Switched Multichannel Networks. *ACM Sigcomm'92*, August 1992.

[14] G. Semaan. P. Humblet. Issues for Slot Transmission in a Passive Optical Star Network based on WDM. *submitted to Globecom'92 Orlando,* Florida - December 6-9, 1992.

[15] L. Roberts. ALOHA Packet system with and without slots and capture. *Computer Comm. Rev.*, 5(4):28-42, April 1975.

[16] W.A. Doeringer *et al.* A Survey of Light Weight Transport Protocols for High Speed Networks. *IEEE, Transactions on communications*, 38(11):2025–2039, November 1990.

[17] D. Bertsekas,R. Gallager. *Data Networks.* Prentice Hall, 1987.

[18] A. Tanenbaum. *Computer Networks.* Prentice Hall, second edition 1989.