A transparent WDM network featuring shared virtual rings

**P.P. Iannone, K.C. Reichmann, A. Smiljanic, N.J. Frigo, A.H. Gnauck,
L.H. Spiekman*, and R.M. Derosier**

AT&T Laboratories – Research, 100 Schulz Drive, Red Bank, NJ 07701
***JDS Uniphase, Prof. Holstlaan 4, Bldg. WAM, NL-5656AA Eindhoven, The Netherlands*
Phone: (732) 345-3216, Fax: (732) 345-3040, Email: ppi@research.att.com

ABSTRACT

We propose and demonstrate a transparent WDM metro ring architecture in which optics enables simultaneous provisioning of dedicated wavelengths for high-end users, while low-end users share wavelengths on "virtual rings". All wavelengths are sourced by the network and remotely modulated at customer "End Stations" by low cost semiconductor optical amplifiers, which also serve as transmission amplifiers. Measured bit-error rate data for three different ring configurations show negligible power penalties associated with upstream transmission around a 120-km 3-node ring through as many as 4 cascaded semiconductor optical amplifier/modulators. Preliminary tests indicate that the SOAs are suitable modulators for packet transmission over the network.

Keywords: Optical fiber communication, Wavelength division multiplexing, Semiconductor optical amplifiers, Optical modulation, Metropolitan area networks

I. INTRODUCTION

Optical networking has shown itself to be valuable in core transport networks, in large part due to the use of wavelength division multiplexing (WDM) [1]. More recently, WDM optical networking has also become important in the metro access arena [2]. The DARPA Next Generation Internet (NGI) initiative, for example, is especially interested in investigating the issues surrounding transport of Internet Protocol (IP) traffic over such networks [3], and recent vendor activity in this space is also evidence of a growing awareness of the need for flexible and high-capacity solutions.

The demands on metro networks are stringent, since the traffic is more diverse than that in the core, yet the system costs must be smaller. Ring architectures are generally preferred since they offer more cost-effective management and protection options, as evidenced by the performance and acceptance of SONET systems that have been used in office rings [2,4]. As in the core networks, WDM is expected to play an important part for several reasons. First, it allows existing fiber to be "mined" for more bandwidth by using additional wavelengths of light. This prevents "fiber exhaust" on existing routes, defers the need to deploy more fibers, and permits more flexible transport solutions. Traditional SONET equipment, for example, could be used on each wavelength, forestalling the need to go to higher rates in the hierarchy. Second, more sophisticated optical networking (i.e. more sophisticated than simply increased transport capacity) can be performed by utilizing the wavelengths as optical channels [5] which can be provisioned, added, dropped, routed, and managed as individual entities, independent of the data format they carry. A third, and corollary reason is that WDM allows service transparency, permitting new services with independent formats to be developed and distributed. The extra dimension in wavelength also permits efficient terminal solutions through using transparency to transport data in native format [6], rather than requiring conversions and multiplexing.

In this paper, we propose and demonstrate an architecture suitable for metro access networks, which exploits the features just mentioned. Specifically, (1) it is a WDM ring, using individual wavelengths to provision services to a geographically diverse set of users. Each wavelength forms a virtual ring and operates independently of the other rings. The architecture is discussed in Section II. (2) It uses optical networking to allow users to participate on different virtual rings. That is, neighboring users could be on the same or different virtual ring, by virtue of the fiber optic connections to the ring nodes. They can be provisioned to share a wavelength with other users, if costs or common channels dictate it, or could have dedicated wavelengths if demanded. Over time, the connections can change or new wavelengths can be added. Each virtual ring forms a network of users connected to a common central hub. (3) The virtual rings are independent, and can support packet-based traffic. Each is amenable to a variety of known [7] or new protocols. We discuss in Section III how two protocols could be used. Optical technology permits the use of a standard optical unit that is not necessarily registered to the wavelength of the user's virtual ring, and could be used to access any fraction of a wavelength's bandwidth, up to the entire channel capacity. We show in Section IV how the bandwidth can be partitioned in a variety of ways. Finally, we note that while we present it as a single architecture, it is possible to consider this as an overlay. That is, all the wavelengths described in the paper can be considered to be some subset of the wavelengths carried on the ring: the other wavelengths might be bearing more conventional circuit-switched traffic, for example.

II. ARCHITECTURAL DESCRIPTION

Fig. 1 shows a schematic of the basic layout [8]. It is a hubbed ring with a Network Node (NN) connected via “feeder” fiber to several Access Nodes (AN), which in turn serve subscribers at End Stations (ES) over “distribution” fibers. The distribution scheme is based on an earlier “RITE-Net” WDM star architecture for passive optical networks (PONs) [9]. In that PON, individual users communicate to the hub by using wavelength-independent modulators to impress data on the optical carriers provided to them by the network [9]. In this work, we extend that architecture in several dimensions, as described below.

First, we apply the concept to a ring, rather than a PON star [9]. This maps the PON’s hub or central office to the ring’s Network Node and maps the PON’s remote node to the ring’s Access Node. The NN contains (see inset, Fig. 1) a WDM transmitter (represented as multiple sources and a multiplexer) to create data at the l_i to be sent downstream for each virtual ring, and a corresponding WDM receiver to detect each of the upstream signals. We implement the Access Node by using a waveguide grating router. Light from the feeder ring (left top fiber in AN inset) enters a waveguide grating router (WGR) input port and is demultiplexed according to its wavelength. According to the “routing property” of the WGR, if light of a given wavelength, after exiting a WGR output port, re-enters an adjacent output port (l_1 and l_2 loopbacks in AN inset), it will emerge on an input port adjacent to the feeder input (right top fiber in AN inset), similar to the operation of RITE-Net [9]. In this way, the AN provides (to the distribution loops) access to the virtual rings in the feeder fiber. The light on the distribution loops provides both downstream information and the ability to send upstream information, as shown in Fig. 2. Modulated light enters the End Station for user N, where a portion is split off to a receiver, which recovers the downstream information, and the remainder is passed to a semiconductor amplifier (SOA)/modulator. The light contains both downstream data and long unmodulated portions upon which upstream data can be impressed. A protocol establishes, through control bits, whether or not user N can seize this “chalkboard” for modulating upstream data. The light, modulated or not, re-enters the access node and continues onto the feeder ring towards the next access node or the hub. This technique is a generalization of the PON architectures and, unlike an earlier approach [10] does not require two separate wavelengths per user for reading and for writing data, thereby simplifying the optical infrastructure.

Second, instead of having each user, by virtue of the star topology, consume a unique wavelength, we demonstrate that wavelengths can be shared among either several end stations at an Access Node or among end stations at different Access Nodes, or both. This permits a group of users to share bandwidth by using packets, and requires a medium access control (MAC) protocol to determine which users are permitted to overwrite the chalkboard in which time slots. In Fig. 1, "ES K_i" is meant to signify that the end station is the K^{th} end station on the virtual ring served by wavelength l_i. Note that in the Figure, l_1 serves J end stations associated with Access Node 1 and one end station associated with Access Node 2, so these J+1 end stations are all on the same virtual ring. Also note that some high-end users may consume an entire wavelength themselves. Even in this dedicated wavelength situation, there are advantages to our approach, since, for example, the wavelength control problem does not need to be solved by the End Station.

Third, instead of a simple modulator as in the PON, we demonstrate the use of a semiconductor optical amplifier (SOA) to both optically amplify the signals and to modulate optical chalkboards with upstream data as shown in the End Station inset. Optical amplification is necessary to support the ring lengths (e.g. 200 km [3]) expected in some metro applications and to overcome excess losses incurred by traversing successive WGRs in the Access Nodes. When the light entering the SOA is bearing data for this or another end station, the SOA acts simply as a single channel optical amplifier. However, when the end station must modulate an optical chalkboard, its data is pre-equalized to compensate for the SOA's frequency response, and applied to the SOA. Like a broadband optical modulator, the SOA is essentially indifferent to the wavelength of the light used: the stimulated emission will be at the same wavelength as the input light. This feature makes it possible to have a universal design so that each ES would be identical, regardless of which virtual ring it occupied. As in all cascaded systems, the buildup of noise must be considered. While the limitation imposed by amplified spontaneous emission is dependent on the loss between amplifiers and the amplifier noise figure, it is reasonable to expect that a properly engineered system could support at least 8 cascaded ES's on each virtual ring [11].

Figs. 3-5 show the flexibility in provisioning that the architecture enables due to the devices. In all cases, the shaded plane represents the physical layout of a simple network: two ANs, each with two users nearby. In the vertical direction, we plot different virtual rings for several wavelengths. Thus, in Fig. 3, all four users are on the same virtual ring. That is, the users served by each AN are connected on a single loop. By virtue of the operation of the ANs, this makes all 4 users connected on the same virtual ring. The upper plane shows the connection at l_1, demonstrating the virtual ring directly. This is a low performance case in which multiple users will all share their single wavelength's bandwidth, according to their protocol. An example would be a network of remote campus locations. This is clearly a case which exploits the fact that fiber can support multiple wavelengths simultaneously, and permits a service which might be difficult to justify economically if conventional transport solutions were required, but would be an additional service opportunity as an overlay. Fig. 4 shows a slightly more aggressive implementation. Now, two networks are formed, one for users 1 and 4 and the other for users 2 and 3. The optical hardware is essentially the same in this case, but now the architecture at the ANs looks as if it is in a star configuration locally, while it is still in a virtual ring configuration globally. The reduced sharing indicates that the traffic loads are higher in this case. An even more aggressive situation is shown in Fig. 5. Here, users 2 and 3 consume so much bandwidth that they are unable to share it with other users. At the same time, user 4, perhaps a computer facility, is both a solitary user on l_4, communicating to remote sites, and also shares a ring with user 1, perhaps an administration center, on l_1.

This approach has some potential pitfalls, following from the fact that the virtual rings are unbroken light paths which can be shared over a wide area with terminal equipment that is not wavelength-specific. First, re-provisioning is more complicated than if users were connected to a physical star architecture.

Second, because the virtual ring is continuous, it must be linked through ANs regardless of the number of End Stations served by that AN. All virtual rings not serving End Stations on a given AN must be completed by looping the demultiplexed wavelengths back into the WGR, as shown in the AN detail in Fig. 1. These looped-back wavelengths must be individually amplified, adding additional cost per subscriber in a sparsely populated scenario. This problem could be solved by moving the transmission amplifiers to the ring fiber and trimming the ES gain so that the ES and distribution fiber have zero net loss (in this way, a simple fiber jumper could be used to loop-back wavelengths that are not serving an ES on a given AN). Third, because of the terminal's simplicity, there is no wavelength reuse. Thus, if ES 2 sends a packet to ES 3, that time slot is unusable for users 4 through N. This entails a tradeoff between system simplicity and throughput. Finally, the failure of a single ES could disrupt service for all users on that virtual ring, thus potentially increasing the complexity and cost of the network protection scheme.

Our architecture has several expected operational advantages. First, the wavelength independence of the end station optical unit should reduce the installed first cost as well as the operational burden. Wavelength control only needs to be performed at the NN, instead of at each of the end stations. Since many multi-wavelength WDM transmitters have the feature that they tune as a comb, it is likely that wavelength can be controlled at the hub with a single degree of freedom. Second, since all wavelengths are sourced and terminated at a common location, management functions which normally require telemetry, such as performance monitoring and fault detection, should be simplified. Third, the nature of the AN makes provisioning this network extremely flexible, as we will show shortly. Fourth, the ring nature of the connection makes it possible to use protocols that have been optimized for ring performance, as discussed in the next section. Newer protocols, based on others that were not specifically designed for rings, are also possible, utilizing the ring, modulation, and packet nature of the network. We give an example of each type of protocol in the next section.

III. MEDIUM ACCESS CONTROL PROTOCOLS

Although many previously reported medium access control (MAC) protocols can be modified to serve as the MAC for our proposed ring network, we focus on two particularly suitable candidates: Fiber distributed data interface (FDDI), which includes a standardized medium access control (MAC) protocol for optical ring networks [12], and ADAPT, a MAC proposed for networks with tree topologies [13] and later incorporated in the IEEE 802.14 standard.

In the usual FDDI implementation, the optical signal is regenerated at each end station, requiring a separate laser diode per site. This is especially onerous in a WDM ring, in which each transmitter must have the correct wavelength. However, the basic concept of FDDI can be readily applied to our network in which the Network Node (NN) provides End Stations (ES) with an optical chalkboard. One token is passed from station to station, and the station that possesses the token at a given time transmits data as long as allowed by its counters. It will hold the token at most for the negotiated time duration, but if the token has arrived earlier than scheduled (evidence that the previous stations have not used their negotiated bandwidths) a station can transmit for a longer time. The NN also may transmit data when it holds the token. If it is not transmitting data, it transmits an optical chalkboard (during which time it does not pass any incoming packets). But while it is transmitting an optical chalkboard, some incoming packets may arrive at the NN on their way from source to destination (for instance, when ES $(J+1)_1$ sends a packet to ES 1_1 or ES J_1). For this reason, FDDI should be modified somewhat in our architecture. For example, packets from ES $(J+1)_1$ to ES 1_1 are sent to the NN first, are stored, and then retransmitted later when the token arrives at the NN. This preserves the unidirectional FDDI character. Another possibility is to allow transmissions in both directions using either two different fibers or two different wavelengths on the same fiber. The Network Node feeds counter-propagating rings with two optical chalkboards in this case. Then, for instance, station ES $(J+1)_1$ transmits packets to station ES 1_1 in the clockwise direction, while station ES 1_1 transmits packets to station ES $(J+1)_1$ in the counter-clockwise direction.

We note that the network efficiency in FDDI on a unidirectional ring is lower than that of FDDI on a bi-directional ring. In FDDI on a bi-directional ring, packets traverse the ring only once. In FDDI on a unidirectional ring, however, those packets that must pass through the hub station will traverse the ring twice, otherwise they will traverse the ring only once. We assume that r/2 of the packets pass through the

Network Node, where r is the fraction of traffic remaining in the ring. Then packets (on average) pass (1+ r/2) rings. If ν_b denotes the efficiency of the bi-directional FDDI, and ν_u denotes the efficiency of the unidirectional FDDI, then it holds that

$$\nu_u = \nu_b/(1 + r/2) = 2\nu_b/(2+r), \quad (1)$$

so that the unidirectional case is as much as 33% less efficient than the bi-directional case. Also, for unidirectional FDDI, packets that traverse the NN may experience an additional delay waiting to be retransmitted. On the other hand, a unidirectional FDDI might have a favorable optical implementation. In both proposed FDDIs the NN might continually negotiate the bandwidth that will be used for packets entering and leaving the metro-ring, and therefore it might flexibly follow the change in the local-to-backbone traffic ratio.

Control in ADAPT is centralized, and in our architecture it would be performed by the Network Node. End Stations use upstream bandwidth (e.g. parts of time slots in a slotted ring) to send requests to the NN. The NN schedules transmissions and sends acknowledgements to End Stations by using down-stream bandwidth. ADAPT can also be implemented on unidirectional and bi-directional rings. If applied on a unidirectional ring, the efficiency in ADAPT degrades similarly to the FDDI case above.

In comparing FDDI and ADAPT, there are some trade-offs. The benefits of FDDI are twofold. Because of its simplicity FDDI can be implemented at high bit-rates supported by developing optical technology. At the same time, FDDI guarantees end stations negotiated bandwidth and access delays satisfying the requirements of most multimedia applications. FDDI-II has been developed to support isochronous circuit-switched traffic as well. An advantage of ADAPT is that the Network Node might meet more sophisticated service requirements since it has complete information about the traffic in the network. On the other hand, more complex processing might be a burden at very high bit-rates. Our optical architecture allows the MAC decision to be made based on the application.

IV. EXPERIMENTAL RESULTS

Bit-Error Rate Performance

Three experimental configurations of a 120-km three-node ring demonstrate the flexibility of the architecture and the feasibility of using SOAs as remote modulator/amplifiers [8]: The first, shown schematically in Fig. 6a, is a single-wavelength virtual ring shared by two users, with one user per AN serving area; the second, shown in Fig. 7a, demonstrates an increase in the number of users per virtual ring from two to four; the third configuration (Fig. 8a) shows the same four users sharing two virtual rings (two users per virtual ring), thus permitting higher per-user data rates than in the previous single virtual ring configuration. The relationships between the physical network layouts for the last two configurations, shown in Fig. 7a and Fig. 8a, and their corresponding virtual ring networks are shown in Fig. 3 and 4, respectively.

The 120-km ring employs two ANs and a NN, each separated by 40 km of conventional single-mode optical fiber. The average loss per 40-km span is 8.3 dB. In the NN, multiple single-wavelength sources (here we demonstrate two) are multiplexed onto the ring. The launched optical power is 6 dBm per wavelength. At each AN, the light enters a (2 x 16)-port WGR, with 50-GHz channel spacing, is demultiplexed, and distributed to the user End Station(s). Each ES includes a 3-dB splitter, which directs half the light to a polarization-insensitive (< 1dB) 1.5-μm semiconductor optical amplifier (SOA) and half to a PIN-FET receiver. The SOAs are pre-biased and directly modulated to 100% modulation depth with a 622 Mb/s pseudo-random pattern of length $2^{23}–1$. The typical fiber-to-fiber gain is 14 dB. After traversing the ES(s), light re-enters the AN and is multiplexed back onto the fiber ring, via the "routing property" of the WGR as described in section II. At the receiver in the NN, the demultiplexer is simulated by a tunable optical bandpass filter with a 3-dB bandwidth of 1.3 nm. A variable optical attenuator and in-line power meter are inserted before the PIN-FET receiver to measure sensitivities.

Fig. 6b shows the bit-error rate (BER) performance of the associated network configuration (Fig. 6a), as measured at the Network Node. A single wavelength (λ_1=1544.7 nm) serves two End Stations, each connected to a distinct AN. The open circles correspond to modulation of the network-provided optical chalkboard at the first End Station (in effect, a 622-Mb/s link from the first ES to the NN, with the second ES unmodulated and therefore serving as a transmission amplifier). The solid triangles show performance for modulation at the second ES (in this case, a 622-Mb/s link from the second ES to the NN, with the first ES serving as a transmission amplifier for the optical chalkboard as it traverses AN 1). In both cases, the performance is nearly identical, with less than 0.1-dB difference in sensitivity at 10^{-9} BER. The power penalty relative to the SOA/Modulator and PIN-FET "back-to-back" (solid circles) is less than 0.3 dB.

Fig. 7b shows BER performance for the same 120-km ring network, for which one additional ES has been added to each AN (Fig. 7a). Since BER performance does not vary greatly from ES to ES, we plot only the extreme cases of 622-Mb/s modulation at the first ES (triangles) and 622-Mb/s modulation at the last ES (squares). Again, the total power penalty relative to "back-to-back" is less than 0.3 dB.

Fig. 8 represents a scenario in which the demand for bandwidth is increased, resulting in the provisioning of two virtual rings (Fig. 4) on λ_1=1544.7 nm and λ_2=1549.1 nm to serve the same set of ES's shown in Fig. 7. In this case, each AN is connected (in a local sense) to its associated ES's in a two-wavelength distribution star configuration (rather than the distribution loops associated with each AN in Fig. 7a). Note that each virtual ring in Fig. 8a is equivalent to the virtual ring in Fig. 6a. Thus, in the absence of crosstalk between wavelengths, we expect the BER performance plotted in Fig. 8b to be identical to that in Fig. 6b. The squares represent 622-Mb/s data added at the second ES on λ_1 and are also indicative of BER performance with the first ES modulated. In this case, λ_2 was modulated at 622 Mb/s at the NN and did not cause a measurable crosstalk penalty.

The BER data plotted in Figs. 6 – 8 demonstrates upstream transmission for one or two wavelengths. However, based on the performance of our commercially available WGRs, and on the gain bandwidth of the SOAs, we expect that a fully populated system (8 wavelengths spread over 6.4 nm for our 16-port WGRs) would also operate without significant crosstalk penalties. Our data also shows up to four SOAs in cascade without appreciable penalty. Based on recently published results [11] we expect that a properly engineered system could support at least 8 ES's per virtual ring. While we have not as yet demonstrated downstream transmission (from the NN to a user's ES), downstream is not as challenging, since the NN transmitters can employ conventional external modulators, which generally outperform SOA/modulators.

Packet Transmission Performance

The BER data reported in the previous section was measured in conventional continuous mode, i.e. using a repeating pseudo-random bit stream continuously clocked at the data rate. Although this is sufficient to demonstrate many key aspects of the ring architecture (such as the performance of the SOAs under high-speed modulation, SOA cascadability without significant power penalty due to ASE accumulation, system power margins around the ring, and tolerance to crosstalk), continuous BER measurements are not a valid test of burst-mode performance. While the entire system was not tested under true burst-mode conditions, due to a lack of both 622-Mb/s burst-mode receivers and a burst-mode bit-error rate test set capable of operating beyond 200 Mb/s, we did modulate SOAs with packet data to test their suitability as burst-mode transmitters. Fig. 9 shows a portion of the pre-equalized electrical drive to the SOA. The constant high level during the initial part of the trace is a segment of the optical chalkboard reserved for overmodulation by one of the downstream end stations. Overshoots during transitions are indicative of the increased amplitude of high-frequency components required to compensate for the high-frequency roll-off of our SOAs. This equalization will not be necessary when the SOA is mounted in a conventional high-speed package, such as those typically used for commercial electro-absorption modulators. Figs. 10 – 11 are oscilloscope traces showing the packet-modulated signal after detection on a broadband dc-coupled optical-to-electrical converter. The trace in Fig. 10 shows the entire repeating pattern consisting of eight 1024-bit slots (this length was limited by the 8192-bit maximum programmable buffer size of our bit-error rate test set). Four consecutive optical chalkboards in slots 5 – 8, comprising 4096 consecutive digital ones, indicate reasonable low-frequency performance of the SOA/modulator. The filtered burst-mode eye diagram (Fig. 11) is wide open, but shows evidence of a slight splitting of the upper rail which should result in less than 1 dB of power penalty. The dark solid portion of the upper rail is due to the long strings of ones from the optical chalkboards. A more thorough investigation of the network's burst-mode performance, which requires burst-mode BER testing, is currently in the preliminary stages and will be reported at a later date.

V. SUMMARY

We have proposed a metro ring architecture capable of supporting multiple virtual transparent rings, each potentially shared among multiple users with arbitrary geographical distribution, and have demonstrated it at a peak rate of 622 Mb/s. Users modulate network-provided and network-controlled wavelengths with inexpensive, polarization-insensitive and wavelength-insensitive SOA/Modulators, which also serve as in-line transmission amplifiers. BER measurements were performed in continuous mode to test the SOAs' response to high-speed modulation and confirmed cascaded operation in these systems. We observed negligible upstream power penalties for three configurations of a 120-km ring: a single wavelength ring serving 2 and 4 users, respectively, and a two-wavelength, two-user-per-wavelength configuration. Time-domain measurements have shown that these SOAs should perform adequately when modulated in burst-mode with packet data. True packet transmission over the network, which is necessary to realize the sharing of virtual wavelength rings among multiple users, will require the implementation of a MAC protocol, such as modified versions of the existing FDDI or ADAPT protocols. A simple analysis shows that the network efficiency could be improved, for both FDDI and ADAPT, by implementing a bi-directional version of the proposed network.

References

1. R. Alferness, P. A. Bonenfant, C. Newton, K. Sparks, and E. Varma, "A practical vision for optical transport networking," *Bell Labs Tech. J*., vol.4 (1), pp. 3-17, Jan-Mar 1999.
2. Y.Chen, M.T. Fatehi, H.J. La Roche, J. Z. Larsen, and B. L. Nelson, "Metro Optical Networking," ," *Bell Labs Tech. J.*, vol.4 (1), pp. 163-186, Jan-Mar 1999.
3. M. Kuznetsov, N.M. Froberg, S.R. Henion, H.G. Rao, J. Korn, K.A. Rauschenbach, E.H. Modiano, and V. S. Chan, "A next-generation optical regional access network," *IEEE Commun. Mag.*, vol. 38(1), pp. 66-72, Jan 2000.
4. J.E. Berthold, "SONET and ATM," Chapter 2 of *Optical Fiber Communications IIIA*, I.P. Kaminow and T.L. Koch, eds., Academic Press, San Diego, 1997.
5. P. Bonenfant and A. Rodriguez-Moral, "Optical Data Networking," *IEEE Commun. Mag.*, vol. 38(3), pp. 63-70, Mar. 2000.
6. J.M. Simmons and A.A.M. Saleh, "Optical Regional Access Network (ORAN) Project," OFC 1999, Feb 22-25, 1999, San Diego, CA, paper WK1, pp 178-180, 1999.
7. A. Tannenbaum, *Computer Networks*, 2nd ed., Prentice Hall, Englewood Cliffs, NJ, 1988.
8. P.P. Iannone, K.C. Reichmann, N.J. Frigo, A.H. Gnauck, and L.H. Spiekman, "A flexible metro WDM ring using wavelength-independent subscriber equipment to share bandwidth," OFC 2000, Mar 7-10, Baltimore, MD., Postdeadline paper, pp. PD38.1-3, 2000.
9. N.J. Frigo, P.P. Iannone, P.D. Magill, T.E. Darcie, M.M. Downs, B.N. Desai, U. Koren, T.L. Koch, C. Dragone, H.M. Presby, and G.E. Bodeep, "A wavelength-division multiplexed passive optical network with cost-shared components," *IEEE Phot. Technol. Lett*., vol. 6, pp. 1365-1367, 1994.

10. M.Sharma, H. Ibe, and T. Ozeki, “WDM ring network using a centralized multiwavelength light source and add-drop multiplexing filters,” *IEEE J. Lightwave Tech*., vol. 15, pp 917-929, 1997.

11. L.H. Spiekman, J.M. Wiesenfeld, A.H. Gnauck, L.D. Garrett, G.N. van den Hoven, T. van Dongen, M.J.H. Sander-Jochem, and J.J.M. Binsma, “160 Gb/s DWDM transmission over 160 km of standard fiber using a cascade of semiconductor optical amplifiers,” ECOC ’99, paper PD2-7, Nice, France, 1999.

12. R. Jain, FDDI Handbook: *High-Speed Networking Using Fiber and Other Media*, Addison-Wesley, 1994.

13. J.E. Dail, M.A. Dajer, C.C. Li, P.D. Magill, C.A. Siller, K. Sriram, and N.A. Whitaker, “Adaptive Digital Access Protocol: A MAC Protocol for Multiservice Broadband Access Networks,” *IEEE Communication Magazine*, pp. 104-112, March 1996

FIGURE CAPTIONS

Figure 1: Architecture. The Network Node, comprising WDM sources and receivers, sends WDM signals (l_1 - l_N) along the feeder fiber to Access Nodes (AN), implemented with Waveguide Grating Routers (WGR). Pairs of WGR output ports define distribution loops in which a single wavelength, forming a virtual ring, can be accessed by one or more End Stations (ES). End Stations consist of receivers (Rx) for downstream information and semiconductor optical amplifiers (SOA) which amplify and modulate light to create upstream signals.

Figure 2: Detail of End Station. Incoming light consists of data packets and unmodulated "optical chalkboards" upon which upstream data signals can be impressed. The passive splitter taps a portion of the light for the receiver, to decode downstream information, and passes the remainder to the SOA. If correct permission is granted in the control bits preceding a chalkboard, the ES may write an upstream message by modulating the chalkboard.

Figure 3: Shared Virtual Ring. The shaded plane represents the physical layout, in which two ANs, each with two ESs, share a single virtual ring. The equivalent optical connections are shown in the virtual ring above the plane.

Figure 4: More aggressive configuration. Same users as in Fig. 3, but now only two ESs share the bandwidth of each of two virtual rings. Since each user on each AN is unique, the physical layout at each AN looks more like a tree.

Figure 5: High capacity network. Applications require almost all of the line rate, so users generally do not share bandwidth with other users. Note that it is possible that a single user can use more than one wavelength, and thus be on more than one virtual ring.

Figure 6: (a.) Two ANs, each with a single ES. Single wavelength virtual ring. (b) BER results. Bit error rates for back-to-back, first ES modulated alone, and second ES modulated alone. Inset shows eye diagram after transmission through the system.

Figure 7: (a) Two ANs, each with two ESs. Single wavelength virtual ring, with ESs both on the same AN and on remote AN. (b.) BER results. Essentially no penalty for modulation at any location.

Figure 8: (a) Two virtual rings on two wavelengths. One ES per AN per virtual ring. (b) BER results. Essentially identical results to Fig. 6 and 7.

Figure 9. Pre-equalized SOA drive. The high frequencies are emphasized to equalize the rolloff of the SOA.

Figure 10: Packet format in the ring. Alternating data and optical chalkboard.

Figure 11: The eye diagram for packet experiment shows slight distortion due to the optical chalkboard.