

Architectural Principles of Optical Regional and Metropolitan Access Networks

Adel A. M. Saleh, *Fellow, IEEE, Member, OSA*, and Jane M. Simmons, *Senior Member, IEEE, Member, OSA*

(Invited Paper)

Abstract— High-end access networks that serve large businesses and campuses will greatly benefit from the introduction of WDM technology, in terms of greater bandwidth, increased flexibility, and enhanced services. We refer to such networks as optical regional and metropolitan access networks (ORMA-Nets). Here, we qualitatively and quantitatively investigate many important principles, as well as challenges, in deploying ORMA-Nets. Access networks in general are functionally comprised of a feeder network, which is responsible for traffic aggregation, and a distribution network, which directly interfaces with the customer premises. We present several configurable, scalable designs for the feeder network that are capable of aggregating a range of traffic types and rates. We also present architectures for achieving a high degree of functionality using relatively low-cost, passive optical components in the distribution network. We explore topics such as optimal switch placement and wavelength banding, and emphasize the technologies that are needed to deliver advanced capabilities. Various underlying themes run throughout the paper, such as optionally not always using bandwidth as efficiently as possible in order to simplify the architecture, and the importance of transparency in providing enhanced services and architectural flexibility.

Index Terms— Configurability, metropolitan area networks (MAN's), regional access networks, transparency, wavelength division multiplexing (WDM).

I. INTRODUCTION

WAVELENGTH division multiplexing (WDM) technology will dramatically improve the capabilities of access networks and enhance the range and quality of services that can be delivered to customers. First, the huge bandwidth afforded by WDM is critically needed as access networks have become one of the chief bottlenecks in the delivery of data services. Current access networks, based on conventional SONET technology, for example, do not allow users to take full advantage of the huge transport pipes being deployed in long-haul backbone networks. Second, WDM technology provides a high degree of configurability, thereby enabling efficient sharing of resources and graceful service upgrades. The resulting flexibility and scalability is highly desirable in an access network. Finally, WDM provides a degree of

transparency, where services can be carried by the network independent of data rate and format. This appears to be very important in an access environment, where the network interfaces with individual customers with a wide array of service demands, as will be further explored below.

One of the barriers to introducing new technology in a network is clearly cost. This is especially significant in an access network where the costs are shared by a much smaller customer base than in a backbone network. Thus, while WDM is currently dominating long-haul networks, it is only recently that the technology has sufficiently matured to justify its deployment in high-end access networks that serve large businesses, campuses, service providers, etc.

In this paper, we present an array of architectural principles and designs for WDM high-end access networks. We refer to these networks as optical regional and metropolitan access networks (ORMA-Nets). (We do not address residential access, which requires an even lower price-point and poses different topological constraints.) Although some of the established technology of WDM backbone networks can be applied to ORMA-Nets, there are many special requirements and critical challenges of delivering WDM in an access environment. At the periphery of an access network, the level of granularity is often an individual traffic stream, with its attendant highly variable characteristics. A backbone network carries traffic that, for the most part, has already been multiplexed and groomed, and hence sees less variability. Thus, an access network needs to handle finer granularity traffic, and be more agile; aggregation is a more significant function than transport. An access network must directly interoperate with a range of traffic types (e.g., IP, ATM, Gigabit Ethernet, etc.), whereas a backbone network is typically dedicated to operating on a more limited number of formats. While the importance of format transparency in backbone networks is still a controversial issue, it seems to play a more significant role in access networks. Also, the distances covered in an access network are generally much shorter than in a backbone network, allowing the use of lower-cost optical technologies, for example directly modulated lasers, less precise wavelength control, arbitrary fiber types, reduced need for regeneration, etc.

Several research projects have specifically addressed WDM regional or metropolitan access. A recent relevant architectural study was the optical regional access network (ORAN) Project [1] (an offshoot of the DARPA-sponsored multiwavelength

Manuscript received February 8, 1999; revised July 30, 1999. This work was performed while the authors were at AT&T Laboratories.

The authors are with Corvis Corporation, Columbia, MD 21046 USA (e-mail: asaleh@corvis.com; jsimmons@corvis.com).

Publisher Item Identifier S 0733-8724(99)09677-2.

optical networking (MONET) Consortium [2]), which qualitatively investigated some of the issues involved with delivering WDM to high-end customer premises. Several metropolitan area testbeds have been built, including: all-optical network (AON) [3], [4], Rainbow [5], [6], optical networks technology consortium (ONTC) [7], lightwave exchangeable add/drop ring network (LEARN) [8], metropolitan optical network (METON) [9], and PROMETEO [10]. AON has a tree-based hierarchical architecture; Rainbow has a broadcast-and-select architecture; the remainder are ring-based and mostly focus on the feeder portion of an access network as opposed to directly interfacing with the customer premises.

The work presented here is architecturally more extensive, extending all the way to the customer premises, and includes quantitative as well as qualitative analysis. We emphasize various architectural options, along with the relative benefits of each, as opposed to presenting a single design. We enumerate technology requirements and functional descriptions rather than detailed component architectures. Various underlying themes run throughout the paper, such as optionally not always using bandwidth as efficiently as possible in order to simplify the architecture, and the importance of transparency in providing enhanced services and architectural flexibility.

Throughout this paper, we refer to the two functional portions of an access network as the *Distribution Network* and the *Feeder Network*. The distribution network directly interfaces with the customer premises and is responsible for delivering and collecting traffic. Some amount of traffic aggregation may occur in the distribution network as well. The feeder portion of the network aggregates traffic, delivers traffic to an appropriate egress point, and transfers traffic from one portion of the distribution network to another. In many instances, the design of the distribution and feeder portions of the network can be treated independently. For example, different levels of redundancy can be present in the feeder and distribution networks, or different wavelength granularities can be deployed in the two portions of the network.

In the next section, we discuss the services required of an ORMA-Net, with an emphasis on addressing the need for transparency. In Section III, we present a high-level architectural description of the network. Section IV explores the details of the feeder network, including topology and protection. Several options for traffic aggregation and switch placement are presented, which demonstrate the architectural flexibility that can be gained from transparency, ultimately leading to more scalable and economic networks. Optical bypass and how it can be efficiently realized through *wavelength banding* technology is explored. We also discuss the flexibility that can be achieved through a hierarchical approach to configurability.

The distribution network is described in detail in Section V. We explain the motivation for wanting it to be totally passive, and describe architectures for achieving a high degree of functionality using low-cost, passive optical components. Some of the features include routed and broadcast wavelength delivery, and survivability. In addition, we enumerate technologies that appear to be crucial for the distribution network, with an emphasis on passive add/drop technology.

Overall, this body of work presents important principles, as well as challenges, in meeting the growing access needs of high-end customers.

II. REQUIRED SERVICES

A high-end access network must support a wide range of data services and rates. First, it must serve the traditional role of an access network in aggregating electronic traffic and delivering the traffic to backbone networks such as an Internet backbone or an ATM network. The rates of such connections are typically T1, T3, or a standard SONET rate, e.g., OC3, OC12, or even OC48. In addition, there are a host of new applications that will stress the capabilities of the network, e.g., video distribution, video conferencing, large image transfers, remote medicine, etc. Huge bandwidth with low latency is a necessity as access networks evolve.

In addition to requesting permanent connections, customers may request that some connections be established on-demand. For example, a customer may request additional bandwidth for a two-week period; or, a customer may need video conferencing capabilities every week for just a short period. The access network must be agile enough to dynamically and efficiently provide the required resources. In general, our philosophy is to deploy infrastructure that will enable rapid provisioning of new connections. While this incurs some additional up-front costs, the incremental cost of provisioning a new connection is relatively low. We expect customers will find fast provisioning to be highly desirable, thus attracting more customers, so that the up-front cost of deploying the network will be more than recovered.

Another formidable challenge is providing virtual LAN connectivity within an access network. Assume that two customer endpoints in a particular ORMA-Net are remote locations belonging to the same company. It is often desirable to create a virtual LAN environment between the two locales so that servers, databases, and peripherals can be easily shared. As the number of data formats and rates deployed in the LAN environment rapidly multiplies, the access network transport capabilities will be required to keep pace. Some examples of current formats and rates are: 100 Base-T (100 Mb/s), High-Speed Token Ring (100 Mb/s), ESCON (Enterprise System Connectivity, 200 Mb/s), and Gigabit Ethernet (1 Gb/s).

In order to be "future proof" in spite of the constant evolution of new data formats, an ORMA-Net ideally should provide connectivity regardless of the protocol and the data rate, i.e., transparency will likely be very important in an access network. Providing format-transparent services is feasible in an access network because many of the impairments that hinder the deployment of a format-transparent long-distance backbone are not as severe in an access network, due to the short distances involved.

Note that even if transparent services are not required by customers, the ability to transparently route wavelengths is still an important feature because it can provide architectural simplicity or flexibility, as will be illustrated by the analysis of optimal switch placement in Section IV-D. We refer to this type of transparency as "internal," because it provides

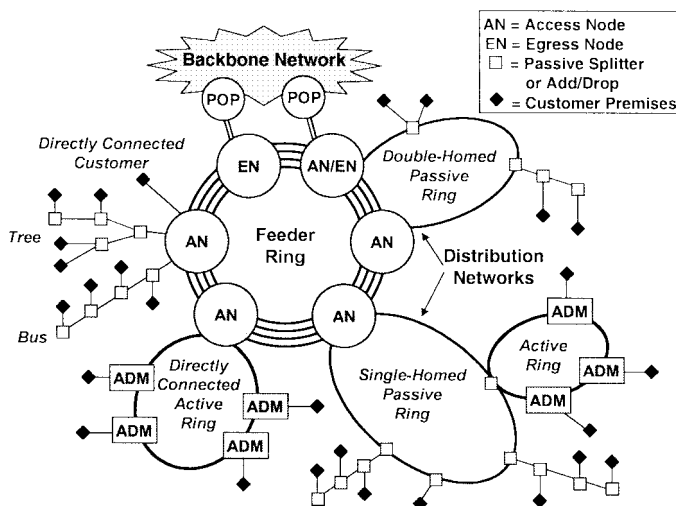


Fig. 1. One example of a regional/metropolitan access network architecture. In this particular example, the feeder network has a ring topology, and is highly configurable; the distribution network can be a tree, bus, or ring, and is totally passive. Legacy active rings, e.g., SONET rings, can also be accommodated. The customers are high-end customers such as businesses, service providers, and campuses.

an economic or operational benefit to the network provider as opposed to enabling new customer services.

III. ARCHITECTURE OVERVIEW

In this section, we present a high-level architectural view of an ORMA-Net, using the architecture shown in Fig. 1 as a representative example. This architecture is a modified version of the architecture developed in the ORAN project [1]. For the remainder of the paper, we use the topology of Fig. 1 for illustrative purposes, however, the principles that we present are fundamental to a range of possible topologies.

In the architecture of Fig. 1, the feeder network has a ring topology, on which are located a set of Access Nodes and Egress Nodes. Access nodes serve as the intermediary point between the feeder and distribution portions of the network. Most of the aggregation and switching occurs in the access node, which is designed to be highly configurable. Egress nodes serve as the interface between the access network and a backbone network. A node can serve as both an access and an egress node. The feeder is multifiber with dense WDM, with several fibers and up to 100 wavelengths per fiber. Optical bypass of a node is supported at a range of granularities, i.e., wavelength level, wavelength-band level, and fiber level (to be discussed further in Section IV-F). The feeder network is shared by all customers, making reliability crucial, while allowing more expensive components to be used (e.g., configurable add/drop multiplexers, dense WDM technology, etc.).

Extending from the access nodes are numerous distribution networks, which can be any of several topologies depending on the required redundancy and the geographical distribution of the customers; e.g., tree, bus, single-homed ring, and double-homed ring. Customers can also directly connect to an access node and thus bypass the distribution network.

WDM is present in the distribution network, although the architecture is flexible with respect to the density of wavelengths used (in the extreme case, only one wavelength is used, i.e., non-WDM). Wavelength granularity will be discussed in Section V-B.

An important feature of the distribution network is that it would be totally *passive*, i.e., contain no amplifiers or active switches. The distribution portion of the network is geographically diverse, making it highly desirable that required maintenance be minimized. Since a significant percentage of network failures are due to power problems, deploying only passive components should greatly improve the reliability of the distribution network. Also, a distribution network, and hence its cost, is shared among relatively few users, which is another motivation for using low-cost, low-maintenance, passive components.

In some metropolitan areas, SONET access rings have already been deployed, where customers gain access via SONET add/drop multiplexers (ADM's). There are several migration strategies, compatible with the ORMA-Net architectures discussed here, that can accommodate the traffic from these legacy rings. Small legacy active rings can either be directly connected to an access node or can extend from a passive distribution network; these two options are shown in Fig. 1. It may be desirable to maintain such rings for the purposes of aggregating traffic from customers with relatively low data rates.

A legacy SONET feeder ring can be treated as an overlay network on the WDM ORMA-Net feeder ring, where a separate band, e.g., 1300 nm, is used to carry the legacy traffic. Ultimately, however, this traffic will represent a small percentage of the feeder traffic, and thus it will be desirable to integrate this traffic with the ORMA-Net traffic in order to take full advantage of WDM technology and to simplify network operations.

Whether an ORMA-Net is deployed in a greenfield scenario or in an area with legacy rings, there is no escape from installing new fiber. It is envisioned that the feeder fiber will be deployed first, with many strands of fiber in the sheath. The distribution fiber will be installed on an as-needed basis, and will be deployed for the largest customers first. Once distribution fiber is in place, smaller customers near the distribution route can be added with relatively little additional fiber needed.

An ORMA-Net is envisioned to be roughly 100 to 1000 square miles in size, serving about 500 to 2000 high-end customers. The feeder ring circumference is expected to be in the 25–100 mile range, with customers typically located no more than a few miles from an access node. In densely populated areas, it may be necessary to deploy a number of smaller rings. The optimal size of the ORMA-Net region and the optimal size and placement of the feeder ring within that region will depend on factors such as the location of existing fiber, geographic constraints, right-of-way, and customer location. The target customers are high-end users such as businesses, government facilities, hospitals, service providers, and campuses. Note that a head-end that serves thousands of residential customers could be a “customer” of an ORMA-Net.

IV. FEEDER NETWORK ARCHITECTURE

This section presents several qualitative and quantitative analyses pertaining to the feeder network. A ring topology is a good choice for the feeder network due to its inherent reliability. In Sections IV-A and -B, the topology of the ring itself is analyzed, in terms of unidirectional versus bidirectional routing and restoration properties. In Section IV-C, we discuss various options for aggregating traffic, where the suitability of a particular option largely depends on user demand and on the state of switching technology. An analysis of optimal switch placement is presented in Section IV-D, with the inherent tradeoff being fewer electronic switches versus less optics. In Section IV-E, a quantitative analysis of optical bypass through *wavelength banding* is presented. Configurability options at various granularity levels are discussed in Section IV-F.

A. Unidirectional Versus Bidirectional Feeder Rings

There are three standard options for architecting a ring: four-fiber ring with bidirectional traffic; two-fiber ring with bidirectional traffic; and two-fiber ring with unidirectional traffic [11]. In all of these options, we assume half of the capacity is dedicated to service traffic, and half is dedicated to protection. A four-fiber ring is generally “overkill” for an access network, where fiber is often scarce; i.e., having to install fiber in units of four may be too inefficient. This option will not be considered further. In a two-fiber bidirectional ring, the wavelengths are partitioned into two sets, for example odd and even wavelengths. In the clockwise direction, one set of wavelengths is used for service, and the other for protection; in the counterclockwise direction, the roles of the two sets of wavelengths are reversed. The third ring option is a two-fiber unidirectional ring where all the service wavelengths are transmitted on one fiber in one direction and the protection wavelengths on the other fiber in the other direction. A high-level comparison of the two-fiber unidirectional and bidirectional options in terms of capacity, cost, and complexity is presented below.

1) *Capacity*: In a typical bidirectional ring, traffic is routed over the shortest path, where a connection uses a clockwise or counterclockwise wavelength as appropriate. Shortest path routing generally leads to more efficient use of fiber. In a unidirectional ring, service traffic is permitted to travel in only one direction, thus, the shortest path may not always be followed.

For simplicity, assume that a feeder ring, comprised of N nodes, supports two types of traffic patterns: hubbed traffic and all-to-all traffic. In the hubbed traffic pattern, we assume there is a single hub on the ring (e.g., an egress node), and that each of the $N - 1$ nonhub access nodes sends H wavelengths worth of traffic to the hub and receives an equal amount of traffic in return. The number of total service wavelengths required to carry this hubbed traffic, in either a unidirectional or a bidirectional ring, is

$$H(N - 1). \quad (1)$$

Thus, for a single-hubbed ring, the number of required wavelengths for hubbed traffic is the same for both unidirectional

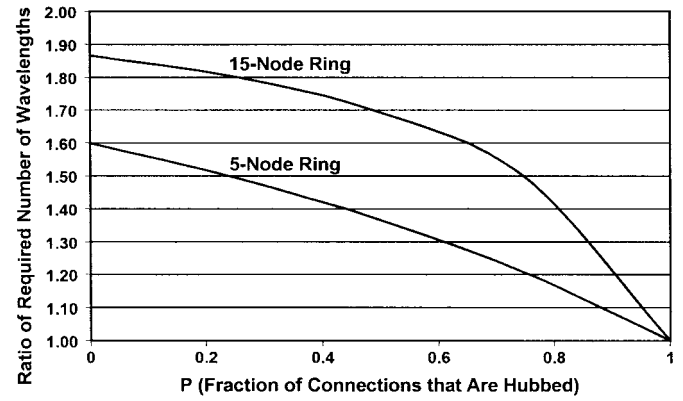


Fig. 2. Ratio of number of required wavelengths in a unidirectional ring versus a bidirectional ring. P represents the fraction of the traffic that is hubbed; the remainder of the traffic is all-to-all.

and bidirectional rings. (If there are multiple hubs, the bidirectional ring may require fewer wavelengths, depending on how traffic is apportioned among the hubs.)

In the all-to-all traffic pattern, we assume each of the N nodes exchange A wavelengths worth of traffic. In bidirectional rings, the total number of clockwise and counterclockwise service wavelengths required to support this traffic is, for N odd, [12]

$$A \frac{N^2 - 1}{4}. \quad (2)$$

In unidirectional rings with all-to-all traffic, a separate wavelength is required for each node pair connection, such that the total number of required service wavelengths is

$$A \frac{N(N - 1)}{2}. \quad (3)$$

Thus, to support all-to-all traffic, the number of required wavelengths for unidirectional rings is approximately twice that of bidirectional rings.

If we assume a fraction P of the total traffic is hubbed, where P equals $H/(H + A)$, then the overall ratio of required number of wavelengths for a unidirectional ring compared to a bidirectional ring is

$$\frac{P + (1 - P)N/2}{P + (1 - P)(N + 1)/4}. \quad (4)$$

This ratio is shown in Fig. 2, for a five-node and a 15-node ring. The greater the fraction of hubbed traffic, the smaller the difference in required number of wavelengths. In an access network, much of the traffic may be sent to/from backbone networks through an egress node, so that a substantial portion of the traffic may be hubbed.

2) *Cost and Complexity*: While two-fiber bidirectional rings use fiber capacity more efficiently than unidirectional rings, they are somewhat more complex and more costly to implement. In the bidirectional case, each fiber carries both service and protection wavelengths, which in general adds to the complexity of demultiplexing and requires more optical components (e.g., filters to split the two sets of wavelengths). In a unidirectional ring, a given fiber carries either service or protection wavelengths, but not both. Thus, it is possible to

architect a unidirectional ring such that the protection fiber does not have to be demultiplexed at all, i.e., all add/drops are from the service fiber. Removing muxes/demuxes from the protection fibers also removes a major source of power loss (on the order of 10 dB of loss per mux/demux pair), thus, potentially fewer optical amplifiers are needed on the protection fiber, resulting in a significant cost savings.

The choice of bidirectional or unidirectional rings also may affect the customer premises. Consider the scenario where a customer's traffic is routed onto the feeder ring without first terminating in an electronic switch in its associated access node, i.e., the customer's traffic is routed transparently from customer premises onto the feeder ring. In order to take advantage of shortest path routing in a bidirectional ring, connections from a given customer may need to be routed using two different wavelengths (one for clockwise-routed connections, one for counterclockwise). Thus, assuming *optical translators* are not used in the access nodes, the customer needs to have two lasers or a tunable laser. In a unidirectional ring, a customer only needs to transmit in one direction and hence at one wavelength. Of course, for flexibility, a high-end customer may need a tunable laser in either case.

Failure recovery is also impacted by the routing architecture. In a bidirectional ring, a node failure may lead to traffic that was destined for that node being misrouted to another node [13]. This problem arises due to the wavelength reuse that occurs in bidirectional rings; i.e., a single wavelength may support numerous connections. In order to avoid misdelivery, the ring network management must include the ability to "squench" certain connections under node failure conditions [14]. This problem is potentially easier to handle with unidirectional rings carrying symmetric traffic. Under these conditions, only one pair of nodes communicates via a particular wavelength. Thus, while a node failure may lead to traffic destined for that node being routed back to the source of the traffic, no other misrouting occurs. This may obviate the need for squenching, or allow the squenching mechanism to be simpler.

Overall, we see that while bidirectional rings require up to 50% less capacity, they likely cost more and are somewhat more complex to implement. The appropriate architectural choice depends on customer traffic patterns and the scarcity of fiber.

B. Feeder Network Reliability

If the feeder network fails, all customers are potentially left without service, thus survivability is crucial. Restoration in the feeder should ideally be implemented in both the optical and electronic layers with coordination between the two domains. For example, to protect against fiber cuts and complete node failures, 2×2 optical switches can be deployed per fiber such that service traffic can be looped back onto protection capacity in order to avoid the point of failure. This allows an entire fiber's worth of traffic to be restored at once, as opposed to restoring each individual wavelength on the fiber. There are a variety of optical ring protection schemes that can be utilized, as described in [15]–[17].

In addition to optical-layer protection, redundancy can be provided against failures in the electronic domain, e.g., a laser failure. In one protection scheme, each node is equipped with a small number of spare lasers and receivers that can be switched to in the event of a failure. If the laser for wavelength A fails in a node, then one of the spare lasers is used, which we assume to be wavelength B . The node that formerly received wavelength A now receives wavelength B . Another option is to have the spare laser be tunable so that it can tune to whichever wavelength has failed. The tunable option is preferable because other nodes are unaffected. In spite of the fact that tunable lasers are more expensive, this may ultimately be cheaper than sparing each wavelength individually.

It is also important to provide redundant paths into and out of the feeder. For example, there should be at least two points of egress to each backbone network. Also, some customers may desire access through two different access nodes, e.g., through dual-homed distribution rings.

C. Traffic Aggregation in the Feeder Network

An access network must be able to aggregate a wide range of service rates and service types. In general there are two strategies to satisfy this requirement. In one approach, resources are matched exactly to demands. For example, an OC3 connection is allocated 1/16 of an OC48-capable wavelength. While this is efficient, it may be more difficult to upgrade (e.g., increase from OC3 to OC12). The second strategy is to allocate resources in 'simple' units, such as full wavelengths. For example, an OC3 ATM connection and an OC3 IP connection may be each allocated a full wavelength, rather than being more efficiently multiplexed. The cost of the simplicity is that wavelength resources are not used as efficiently. These two strategies are discussed further in the following sections. In general, the huge bandwidth afforded by WDM technology favors "squandering" bandwidth in order to achieve simplicity and flexibility.

1) *Assigning Connection Rates*: High-end customers are expected to request electronic connections at rates from OC3 to OC48, or higher. Such connections must eventually enter an electronic switch in the feeder, where it is necessary that the line card rate match the transmission rate at the customer premises; e.g., an OC3 connection needs to terminate in an OC3 line card. (In principle, line cards may be able to adapt to the connection rate, however, this option may be too costly).

The necessity to match connection rate to line card rate poses several architectural tradeoffs. One solution is to match the desired customer connection speed to an appropriate line card; however, this solution is limiting if the customer wants to change its rate. If rate-change requests are infrequent, they can be accommodated by manual intervention at the access node. If rate changes are more frequent, a more flexible solution is to use a crossconnect in the access node to match connection rate to line card rate. Thus, the data switch in the access node would contain some number of OC3, OC12, and OC48 (or higher) line cards, and the crossconnect would match the customer input to an appropriate line card. One drawback to this approach is the cost of the crossconnect.




Fig. 3. In this solution, a SONET Mux is used to multiplex the various electronic service types at the customer premises. In the access node, a SONET switch distributes each service type to the appropriate switch or router. The details of the distribution network are not shown.

A second drawback is that the number of line cards has to be over-provisioned to ensure that the blocking probability is low enough, e.g., there has to be enough OC48 line cards to satisfy customer demand at any given time even though for the majority of time some of them will be unused.

Another solution is to assign a customer a data rate that is at least as high as it would ever need. Thus, if a customer normally transmits at OC3, but occasionally needs to transmit at OC12, that customer would be provided with an OC12 connection, even though 75% of the OC12 frame would usually be unused. This solution eliminates the need for a crossconnect to distribute the connections to line cards in the access node. A possible drawback to this approach relates to the size of the electronic switch fabric. If the switch fabric needs to be sized to meet the sum of the line card rates, then using connection rates that are higher than necessary results in switch fabrics that are unnecessarily large. If the switch fabric is sized to the sum of the actual *used* portion of the incoming frames, then this is not an issue. This should be one of the important requirements for switch manufacturers.

Overall, we favor the latter approach, where the connection rate always equals the maximum that a customer may need on a given wavelength.

2) *Aggregation of a Range of Service Types:* An ORMA-Net needs to provide access to a variety of electronic networks, e.g., ATM, IP, and Frame Relay. Furthermore, an individual customer may request multiple electronic service types. In this section, we discuss several options for dealing with the diversity of services. The options described expand upon those presented in [1].

In one solution, the various electronic types requested by a customer are multiplexed in a SONET Mux at the customer premises as shown in Fig. 3. For example, at Customer Premises 1, IP and ATM are multiplexed onto a single wavelength. This wavelength is eventually delivered to a port on a SONET switch within the access node. The SONET switch multiplexes traffic together from various customers (e.g., the ATM traffic from Customer 1 and Customer 2 can be multiplexed together), and directs the various traffic types to the appropriate switch or router (e.g., the ATM traffic is directed to the ATM switch). Statistical multiplexing of traffic typically occurs in the switch or router, and the traffic is sent out on the feeder ring to the corresponding egress node (e.g.,

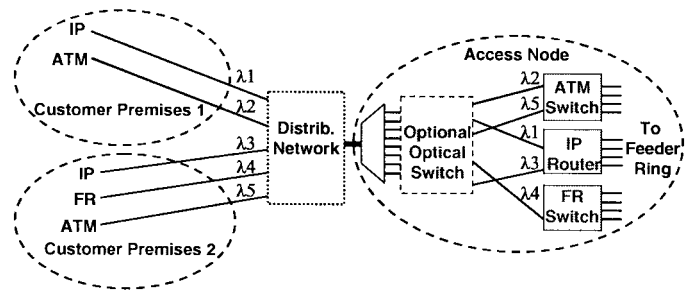


Fig. 4. In a variation of Fig. 3, each customer service type is assigned to a separate wavelength, so that no muxing of service types takes place at the customer premises. In the access node, each wavelength is either permanently tied to a particular switch, or to add flexibility, an optical switch can be used.




Fig. 5. In this solution, electronic services are mapped into a single format, for example, ATM, and multiplexed in an ATM Mux at the customer premises. In the access node, an ATM switch groups traffic by service type, and multiplexes traffic from multiple customers. The output of the ATM switch is sent to the appropriate egress node corresponding to the service type.

the ATM traffic is delivered to the egress node that interfaces with the ATM backbone network).

A powerful variation of this architecture, shown in Fig. 4, is to dedicate a wavelength per electronic service type requested by the customer. Thus, Customer 1 would require two wavelengths and Customer 2 three wavelengths. The wavelengths can be directly tied to a switch in the access node, or, an optical switch can be used in the access node so that a customer may change its service types.

In another solution, shown in Fig. 5, all electronic service types are mapped at the customer premises to a single data type, which we assume in the figure to be ATM. The traffic is multiplexed together in an ATM multiplexer at the customer premises, and delivered to an ATM switch in the access node. The ATM switch multiplexes traffic together that is destined for the same backbone network, e.g., it groups together all the IP traffic from all of the customers, and directs it to the appropriate egress node.

There are advantages and disadvantages to each of the above three options. The great advantage of the architecture shown in Fig. 4, where a wavelength is dedicated to each service type, is that it eliminates the need for multiplexing equipment at the customer premises. In addition, this option provides a great deal of flexibility in where switches are placed, as is discussed in Section IV-D. The disadvantage is the potentially inefficient use of wavelengths for customers with relatively small data rates such as OC-3 or below. For such customers, it may be

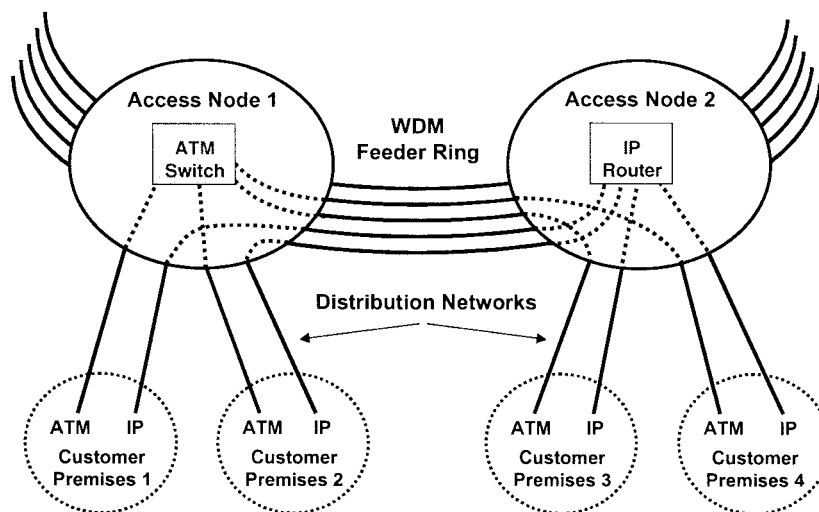


Fig. 6. Ability to transparently route wavelengths from the customer onto the feeder ring allows the option of deploying switches in only a subset of the access nodes. In the figure, which assumes bidirectional routing on the feeder ring, an IP router is deployed in Access Node 2, but not Access Node 1. The IP traffic from Access Node 1 is routed optically directly from the distribution network of Access Node 1 to the router in Access Node 2.

advantageous to have small SONET rings that aggregate the traffic from multiple customers before feeding the traffic onto the distribution network, as was shown in Fig. 1. For example, the ATM traffic of multiple customers can be multiplexed onto a single wavelength, which is then carried as shown in Fig. 4. In general, although the architecture of Fig. 4 requires more wavelengths, overall, it still could be a more economical solution.

The advantage of the option shown in Fig. 5 is that mapping the traffic to one data type leads to fewer required electronic boxes in the ORMA-Net. This results in access nodes that are easier to maintain, easier to upgrade, lower cost, and that have smaller space and power requirements. However, as will be shown in Section IV-D, these same benefits can be realized with the architecture of Fig. 4. One disadvantage of mapping the traffic is the inefficiency of mapping one protocol type to another. For example, there is additional overhead in breaking an IP packet into cells and adding an ATM header. Also, less statistical multiplexing may be achieved when protocols are mapped, and there may be added latency with this architecture. Another important disadvantage is the need for protocol translators and muxes at the customer premises.

The architecture of Fig. 3 is not advantageous due to its requiring a mux in the customer premises, and requiring numerous electronic boxes in the access node.

Note that it is possible to combine aspects of these architectures.

D. Optimal Placement of Electronic Switches

The ability to route a wavelength transparently from the distribution network onto the feeder ring, combined with the architecture shown in Fig. 4, affords the opportunity to place electronic switches in only a subset of the access nodes. Consider the two access nodes shown in Fig. 6, where we assume bidirectional routing is supported on the ring. Assume that customers served by the two access nodes have IP and ATM services. One option would be to place both an IP router

and an ATM switch in both access nodes. A better option, as shown in Fig. 6, is to have the wavelengths carrying one traffic type, say IP, not be terminated at all in Access Node 1, and instead, be routed to Access Node 2, where an IP router is located. Thus, the IP router in Access Node 2 is shared between customers in both Nodes 1 and 2. Similarly, the ATM switch in Access Node 1 processes the ATM traffic entering at either of the nodes. (Alternatively, both the ATM switch and the IP router could be placed in only one of the access nodes.) Note that this architecture requires that customers place their ATM and IP traffic on separate wavelengths as was shown in Fig. 4.

Eliminating switches from some of the nodes results in less equipment, and correspondingly lower maintenance and upgrade costs, and less power and space needs. However, more feeder ring wavelengths are needed because traffic is carried on the feeder without first being multiplexed in a switch. Nevertheless, the quantitative analysis below demonstrates that significant cost savings can potentially be achieved.

1) *Cost Analysis:* In this section, we present a first-order cost analysis of the optimal electronic switch location in the feeder ring. In general, placing switches in only a subset of the nodes reduces the cost of the electronics (fewer switches) while increasing the cost of the optics (more wavelengths). These tradeoffs are quantified below.

The concept of statistical multiplexing is very important in the analysis. Statistical multiplexing occurs when multiple packet-switched, bursty traffic sources are multiplexed together in a switch to produce a smoother stream. For example, assume there are five streams of bursty data, each with a peak rate of R but an average rate lower than R . Assume that when the five streams are multiplexed together in the switch, they generate a steady traffic stream of rate R . The statistical multiplexing gain (SMG) would be 5 in this example. SMG is difficult to predict in a data network and varies depending on the traffic being carried. Typically, however, greater SMG can be realized when more traffic streams are multiplexed together.

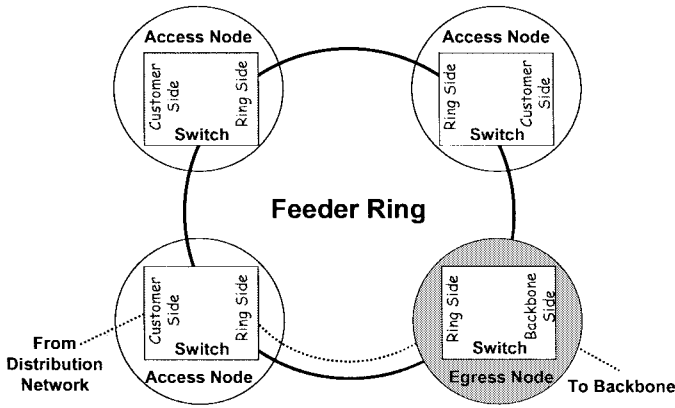


Fig. 7. In this architecture, a switch is placed in each access node and the egress node. Traffic is multiplexed as soon as it enters the access node, and is further multiplexed at the egress node.

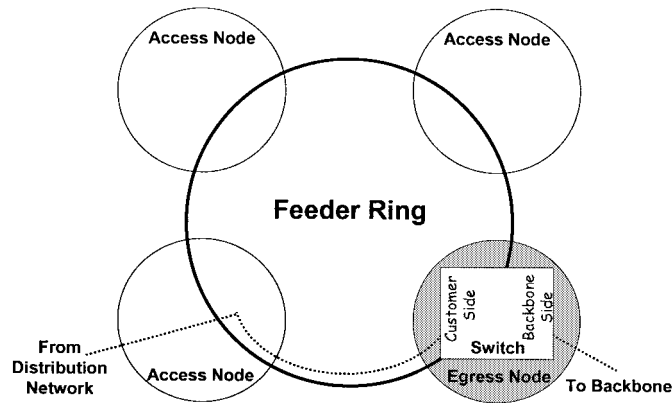


Fig. 8. In this architecture, there is a switch only in the egress node. Traffic is routed transparently on the feeder ring to the egress node. Fewer electronic switches are needed, at the expense of more wavelengths.

For simplicity, we consider the placement of just one type of electronic switch within the feeder network (e.g., an IP router), however, the results generalize to the more realistic scenario of multiple switch types. We consider two extreme architectures as shown in Figs. 7 and 8. In Fig. 7, a switch is placed in each access node, such that traffic is multiplexed as soon as it enters the node. From there, the traffic is sent to the egress node, where it enters a second switch. The purpose of being processed by a second switch is to achieve greater SMG, because the switch in the egress node processes the traffic from *all* access nodes, as opposed to just one access node.

In the architecture of Fig. 8, a switch is present in only the egress node. Thus, traffic is routed from the distribution networks directly onto the feeder ring, and delivered to the switch in the egress node. We assume that the *overall* SMG achieved in this scenario is the same as in Fig. 7, i.e., the same number of full wavelengths exit the egress node switch in either architecture.

In the analysis below, we assume that traffic is symmetric to and from the backbone, and that all customer traffic is destined for the backbone. We assume that the total traffic in the feeder ring is large, so that there are multiple wavelengths worth of traffic, and that switches can perfectly pack the wavelengths with traffic. For concreteness, we assume the maximum rate

carried by a wavelength is OC48, however, the results are very general.

2) *Cost of Electronics*: First, we focus on the cost of the electronic switches, which is comprised of two components: the line card cost and the switching fabric cost. Initially, we assume that these costs are linear with respect to data rate (e.g., the OC12 line card cost is four times that of an OC3 line card). We assume time-based multiplexing occurs before traffic enters the switch fabric (e.g., if just one OC3 is carried in an OC48 frame, the fabric processes just the OC3, not an OC48), but that statistical multiplexing is accomplished as part of the switch fabric. The switch fabric must be sized to accommodate the traffic on the input ports in both directions (i.e., to and from the backbone).

The following notation is used.

- N_A Number of access nodes in the feeder ring.
- T Total amount of traffic entering an Access Node, in Gb/s; initially, we assume all nodes have equal traffic, however, this assumption is later relaxed.
- \mathcal{L} Line card cost, per Gb/s, for an OC48 line card.
- \mathcal{F} Switch fabric cost, per Gb/s.
- α_A *Reciprocal* of the SMG achieved by a switch in an access node.
- α_E *Reciprocal* of the SMG achieved by a switch in the Egress Node, where $\alpha_E \leq \alpha_A$.

Case 1: Switch in Every Access Node and in the Egress Node (i.e., Fig. 7)

Total line card cost of all Access Nodes on

$$\text{customer side of switch} = N_A T \mathcal{L}$$

Total line card cost of all Access Nodes on ring

$$\text{side of switch} = N_A T \alpha_A \mathcal{L}$$

Line card cost of the Egress Node on ring

$$\text{side of switch} = N_A T \alpha_A \mathcal{L}$$

Line card cost of the Egress Node on backbone

$$\text{side of switch} = N_A T \alpha_E \mathcal{L}$$

$$\text{Total Line Card Cost} = N_A T (1 + 2\alpha_A + \alpha_E) \mathcal{L}. \quad (5)$$

Total Switch Fabric cost of all Access Nodes

$$= N_A T (1 + \alpha_A) \mathcal{F}$$

Switch Fabric cost of the Egress Node

$$= N_A T (\alpha_A + \alpha_E) \mathcal{F}$$

$$\text{Total Switch Fabric Cost} = N_A T (1 + 2\alpha_A + \alpha_E) \mathcal{F} \quad (6)$$

$$\text{Total Electronics Cost} = N_A T (1 + 2\alpha_A + \alpha_E) (\mathcal{L} + \mathcal{F}). \quad (7)$$

Case 2: Switch in Egress Node Only (i.e., Fig. 8)

Line card cost of the Egress Node on ring side of

$$\text{switch} = N_A T \mathcal{L}$$

Line card cost of the Egress Node on backbone

$$\text{side of switch} = N_A T \alpha_E \mathcal{L}$$

$$\text{Total Line Card Cost} = N_A T (1 + \alpha_E) \mathcal{L} \quad (8)$$

$$\text{Total Switch Fabric Cost} = N_A T (1 + \alpha_E) \mathcal{F} \quad (9)$$

$$\text{Total Electronics Cost} = N_A T (1 + \alpha_E) (\mathcal{L} + \mathcal{F}). \quad (10)$$

Comparing Cases 1 and 2, we see that the ratio of electronics costs is

$$\begin{aligned} & \frac{\text{Cost with Switches in Access and Egress Nodes}}{\text{Cost with Switch in Egress Nodes Only}} \\ &= \frac{1 + 2\alpha_A + \alpha_E}{1 + \alpha_E}. \end{aligned} \quad (11)$$

As indicated by (11), the electronics cost is always greater with a switch in both the access and egress nodes, as expected. Interestingly, it can be easily shown that (11) holds regardless of the total traffic in the network or how the traffic is allocated among access nodes; it is also independent of \mathcal{L} and \mathcal{F} . If there is no statistical multiplexing gain (e.g., all circuit-based traffic), then the cost ratio is two. More reasonable values are $\alpha_A = \frac{1}{2}$ and $\alpha_E = 1/3$, which yield a cost ratio of 1.75.

3) *Sublinear Line Card Costs*: As equipment technology matures for a particular data service, the pricing structure tends to change; for example, it is common that line card costs become sublinear. Assume that line card costs follow a square-root law such that for each quadrupling in line card speed there is only a doubling in line card cost. The analysis above can be modified accordingly. Let

- p_1 proportion of customer traffic at OC3 rate;
- p_2 proportion of customer traffic at OC12 rate;
- p_3 proportion of customer traffic at OC48 rate.

The effective line card cost per Gb/s on the customer side of the first encountered switch is then $\mathcal{L}(4p_1 + 2p_2 + p_3)$ as opposed to \mathcal{L} . (The remaining line card costs are not affected, as all other line cards are assumed to be OC48 rate and \mathcal{L} was defined as the unit cost per Gb/s for an OC48 line card.) The factor of “1” in (5) and (8) is replaced by $4p_1 + 2p_2 + p_3$. The overall electronics cost ratio given by (11) is modified to

$$\begin{aligned} & \frac{\text{Cost with Switches in Access and Egress Nodes}}{\text{Cost with Switch in Egress Nodes Only}} \\ &= \frac{\mathcal{L}(4p_1 + 2p_2 + p_3 + 2\alpha_A + \alpha_E) + \mathcal{F}(1 + 2\alpha_A + \alpha_E)}{\mathcal{L}(4p_1 + 2p_2 + p_3 + \alpha_E) + \mathcal{F}(1 + \alpha_E)}. \end{aligned} \quad (12)$$

This ratio is greater than 1 for all mixes of customer traffic rates, however, the greater the proportion of lower speed lines cards, the smaller (i.e., closer to 1) the cost ratio.

4) *Cost of Optics*: Placing a switch only in the egress node reduces the electronics costs, but results in higher optical component costs due to more required wavelengths. We use the additional notation:

- W number of wavelengths entering an Access Node from the distribution networks;
- U average utilization of a wavelength (e.g., an OC12 connection on one wavelength represents a 25% utilization);
- C_D cost of a wavelength in the distribution network;
- C_F cost of a wavelength in the feeder ring (all traffic is assumed to be hubbed so there is a one-to-one correspondence between connections and feeder wavelengths, regardless of whether bidirectional or unidirectional routing is used).

The wavelength costs include the cost of muxes/demuxes, amplifiers, optical translators, etc., as appropriate. We assume that all wavelengths are fully packed (e.g., carry an OC48 worth of traffic) after exiting a switch. Note that $T = (U)(W)$ (OC48), where “OC48” represents 2.5. If switches are only placed in the egress node, then the number of wavelengths in the feeder is the same as in the distribution networks.

Case 1: Switch in Every Access Node and in the Egress Node

$$\begin{aligned} \text{Total optics costs in Distribution Network} &= N_A W C_D \\ \text{Total optics costs in Feeder Ring} &= N_A (T \alpha_A / \text{OC48}) C_F = N_A (U W \alpha_A) C_F \\ \text{Total Electronics + Optics Cost (using (7),} & \\ \text{with linear line card costs)} &= N_A T (1 + 2\alpha_A + \alpha_E) (\mathcal{L} + \mathcal{F}) \\ &+ N_A W (C_D + U \alpha_A C_F) \\ &= N_A W [U (\text{OC48}) (1 + 2\alpha_A + \alpha_E) (\mathcal{L} + \mathcal{F}) \\ &+ (C_D + U \alpha_A C_F)]. \end{aligned} \quad (13)$$

Case 2: Switch in Egress Node Only

$$\begin{aligned} \text{Total optics costs in Distribution Network} &= N_A W C_D \\ \text{Total optics costs in Feeder Ring} &= N_A W C_F \\ \text{Total Electronics + Optics Cost (using (10),} & \\ \text{with linear line card costs)} &= N_A T (1 + \alpha_E) (\mathcal{L} + \mathcal{F}) + N_A W (C_D + C_F) \\ &= N_A W [U (\text{OC48}) (1 + \alpha_E) (\mathcal{L} + \mathcal{F}) \\ &+ (C_D + C_F)]. \end{aligned} \quad (14)$$

Thus, placing switches only in the egress node yields a cost difference of [see (14)–(13)]

$$N_A W [(1 - U \alpha_A) C_F - U (\text{OC48}) (2\alpha_A) (\mathcal{L} + \mathcal{F})]. \quad (15)$$

The cost criterion favoring the placement of a switch only in the egress node as opposed to in both access and egress nodes is that (15) be less than 0, or:

$$\frac{\mathcal{L} + \mathcal{F}}{C_F} > \frac{1 - U \alpha_A}{2U \alpha_A (\text{OC48})}. \quad (16)$$

The criterion of (16) is independent of the total traffic in the network and how the traffic is apportioned among access nodes. Furthermore, the precise same criterion holds for the case of nonlinear line card costs (in this cost model, the term representing the nonlinear costs appears in the cost equations for both architectures and is subtracted out when calculating the cost difference).

Lower average utilization of incoming wavelengths (i.e., smaller U) and larger SMG in the access node (i.e., smaller α_A) results in the right side of (16) being larger, i.e., the threshold for placing the switch only in the egress node becomes higher. To gain insight into (16), we use the following reasonable values for the parameters: $\mathcal{L} = \mathcal{F} = \10 K; $U = 30\%$; $\alpha_A = \frac{1}{2}$. With these values, as long as C_F (the cost of a wavelength in the feeder) is smaller than approximately \$17 K, there is a cost benefit to placing switches only in




Fig. 9. In this example of wavelength banding, the 32 wavelengths in the spectral comb are partitioned into four bands, each comprised of eight contiguous wavelengths. In general, only a subset of the bands needs to be dropped at each node.

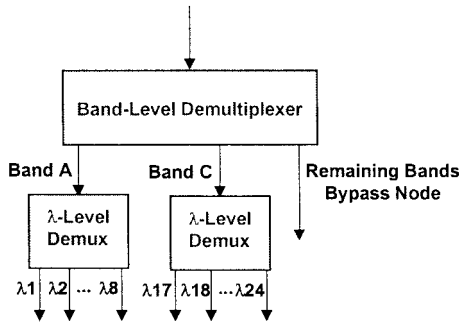


Fig. 10. Functional diagram indicating a hierarchical approach to demultiplexing. At the coarsest level, wavelength bands A and C are dropped in this figure, while all other bands pass through. Bands A and C are further demultiplexed into their constituent wavelengths.

the egress node. It is very likely that C_F will be below this threshold, especially as WDM prices mature.

E. Optical Bypass with Wavelength Banding

As illustrated in the previous section by wavelengths that are routed directly from a distribution network to an egress node, it is not necessary that every wavelength in the feeder be terminated in every node. Rather, a wavelength can optically bypass a node if it carries no traffic to or from that node. Implementing optical bypass is critical to the economics of a WDM access network. With possibly tens of fibers in the feeder and up to 100 wavelengths per fiber, demultiplexing every wavelength of every fiber would be prohibitively expensive. The benefits of optical bypass are well documented; e.g., see [18], [19], [13], [20], [21].

To simplify the optics of a node, it is best to implement bypass at the coarsest granularity possible. In this section, we use the technique of *wavelength banding* for purposes of efficiently implementing optical bypass, i.e., entire *bands of wavelengths* optically bypass a node. In general, a band is a subset of the wavelengths on a fiber, typically contiguous, as illustrated in Fig. 9 by a spectral comb of 32 wavelengths partitioned into four bands. (Note that, in general, a small number of wavelengths (e.g., one or two) could be sacrificed in a “guard band” between the wavelength bands. This is not shown in the figure, but will be discussed in Sections V-C and -D.) While here we focus on banding for the purposes of efficiently implementing optical bypass, banding is also used in some amplified systems to extend the optical spectrum of the amplified signal [22].

Wavelength banding allows an efficient hierarchical demultiplexing methodology to be used, as shown by the functional diagram of Fig. 10. The first level of demultiplexing drops only those bands that carry traffic to that node, while the




Fig. 11. Example of a band map to carry hubbed and all-to-all traffic. The H bands are used to carry hubbed traffic and the A bands are used for all-to-all traffic.

remaining bands optically bypass the node. The dropped bands are further demultiplexed at the wavelength level. The hierarchical structure is expanded upon in the context of configurability in Section IV-F.

1) *Fixed Traffic*: To better illustrate optical bypass with wavelength banding, we first consider banding of fixed traffic; in the next section, we extend the results to networks supporting dynamic traffic. Assume the feeder ring contains N nodes of which one is designated as a hub. Also assume the traffic on the ring is fixed, such that there are H wavelengths of hubbed traffic between every node and the hub, and A wavelengths of all-to-all traffic between each node pair (i.e., the same traffic pattern analyzed in Section IV-A). We assume that bidirectional routing is supported on the ring, however, wavelength banding is just as easily implemented with unidirectional routing. We only consider the service wavelengths in the discussion below. For simplicity, we assume N is odd, however the results are readily extended to N even. Fig. 11, where the number of nodes is seven, is used to illustrate the banding structure.

First, consider the hubbed traffic. With bidirectional routing, a wavelength that carries hubbed traffic can be “reused” on the other side of the ring. For example, a clockwise wavelength carrying traffic from Node 7 to the hub can also be used to carry traffic from the hub to Node 4. Overall, the number of bands needed to carry the hubbed traffic is simply $(N - 1)$, where each band is dropped at two nodes. Half of the bands carry traffic in the clockwise direction, the other half in the counterclockwise directions. Each band is comprised of H wavelengths. In Fig. 11, which depicts the bands for only one direction, the hubbed bands are represented by H1, H2, and H3.

The number of bands required to carry the all-to-all traffic is $(N^2 - 1)/4$, with half of the bands travelling in each direction, and each band comprised of A wavelengths. These bands are represented by A1 through A6 in Fig. 11, where again, the bands in only one direction are shown. For any sized ring, it is theoretically possible to implement banding for all-to-all traffic such that no band is dropped in any more than 4 nodes [21]. Limiting the number of nodes in which a particular band is dropped may be important in minimizing the effects of accumulated impairments such as crosstalk [23].

To further illustrate optical bypass with wavelength banding, assume there is a total of 6 fibers worth of traffic, three in each

TABLE I
NUMBER OF WAVELENGTHS DEMULTIPLEXED AT EACH NODE (PER DIRECTION)

Fiber 1 = H1, A1, A2 Fiber 2 = H2, A3, A4 Fiber 3 = H3, A5, A6 (32 Wavelengths Per Fiber)			
Node 1 (Hub):	Fiber 1: 32	Node 5:	Fiber 1: 22
	Fiber 2: 22		Fiber 2: 10
	Fiber 3: 12		Fiber 3: 10
Node 2:	Fiber 1: 22	Node 6:	Fiber 1: 0
	Fiber 2: 10		Fiber 2: 32
	Fiber 3: 10		Fiber 3: 10
Node 3:	Fiber 1: 0	Node 7:	Fiber 1: 10
	Fiber 2: 32		Fiber 2: 0
	Fiber 3: 10		Fiber 3: 32
Node 4:	Fiber 1: 10		
	Fiber 2: 0		
	Fiber 3: 32		

direction, where each fiber carries 32 service wavelengths. Assume that H is 12 and A is ten. In Table I, we detail one particular band allocation scheme. For each of the three fibers at each node in a given direction, the table indicates how many of the 32 wavelengths need to be dropped. In most of the nodes, the architecture enables optical bypass at both the fiber and the band level. In this example, less than half of the total wavelengths need to be demultiplexed at any nonhub node. In general, the fraction of wavelengths that needs to be demultiplexed at a nonhub node is

$$\frac{8H + 4A(N - 1)}{4H(N - 1) + A(N^2 - 1)}. \quad (17)$$

2) *Dynamic Traffic*: It is very likely that in addition to fixed internodal connections, there will also be dynamic connections, for example to provide on-demand virtual-LAN connectivity. A banded architecture is suitable as well for a dynamic environment, although the precise banding pattern will depend on expected traffic patterns and the desired level of blocking. (Even without banding, dynamic traffic potentially results in blocking.)

One strategy for accommodating dynamic traffic is to lay out bands such that there is at least one band in common between any two nodes; thus, any desired internode connection request is potentially satisfied (although depending on the connections already established, a connection may be blocked). Assume that there is an additional requirement that no band be dropped in more than T nodes to minimize the impact of impairments. For an N -node ring, there are $N(N - 1)/2$ possible connections, and each band can provide at most $T(T - 1)/2$ connections. Thus, a lower bound on the number of required bands is

$$\left\lceil \frac{N(N - 1)}{T(T - 1)} \right\rceil. \quad (18)$$

Another strategy for accommodating dynamic traffic is to drop at least one band of wavelengths in all nodes. Such an

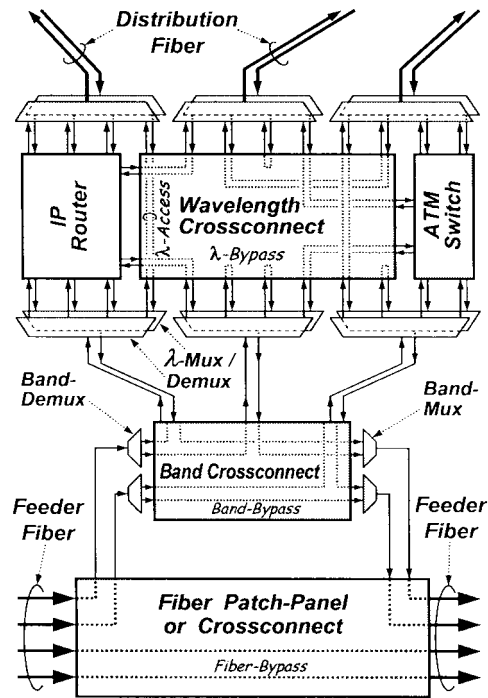


Fig. 12. Functional diagram of the levels of configurability that might be desirable in an access node. Optical bypass of the node can occur at a fiber, band, or wavelength level. Drop and bypass lines at each level can be either fixed or configurable. Some wavelengths are terminated on an electronic switch, while others flow transparently from feeder ring to distribution network, with the wavelength crossconnect providing flexibility in routing wavelengths to the various distribution trees. The wavelength crossconnect also is capable of routing wavelengths within a distribution tree or from one distribution tree to another.

architecture is also advantageous for supporting broadcast traffic. The drawback is that if the number of access nodes is too large, the accumulated impairments may degrade performance.

Ideally, the access node architecture allows the banding pattern to be configurable to accommodate changing traffic patterns and traffic growth. Configurability in the access nodes is the subject of the next section.

F. Access Node Configurability

Configurability in the access nodes provides efficient utilization of resources while accommodating dynamic traffic patterns. By allowing resources to be shared, it enables a given blocking probability to be achieved with fewer deployed muxes/demuxes, electronic switch ports, distribution wavelengths, etc. Configurability can be provided at the fiber, band, and wavelength level, as represented by the hierarchical functional diagram shown in Fig. 12. This figure includes a tremendous amount of functionality; only a subset of this might be present in any particular access node, depending on cost and expected demand.

A fiber-level crossconnect provides configurability at the coarsest level, allowing an entire fiber from the feeder ring to optionally drop at or bypass the node. The crossconnect may be realized via a bank of 2×2 switches (e.g., one switch per fiber), as the functionality of a fully connected crossconnect is likely not needed. Configurability may not be required for all ring fibers; some fibers may drop or bypass the node on a

fixed basis. In the extreme case, a fixed patch panel is deployed rather than a fiber-level crossconnect.

The dropped fibers are demultiplexed into their constituent bands (the figure depicts two bands per fiber), which are routed to the band-level crossconnect. The functionality of the band-level crossconnect is identical to that at the fiber level. Bands are either dropped, or shunted back to the fiber level (i.e., “band-bypass”). Again, a bank of 2×2 switches operating on individual bands may provide the necessary functionality. Some drop or bypass bands may be deployed on a fixed basis.

Bands that are dropped are demultiplexed into their constituent wavelengths, which in turn may be switched by the wavelength-level crossconnect. There is a wealth of functionality provided at the wavelength level, as indicated in the figure. Some wavelengths are immediately shunted back to the band level (i.e., “ λ -bypass”). Some wavelengths are directed to an electronic switch within the node (e.g., an IP router or an ATM switch); the wavelengths can either be tied to a particular electronic switch port, or through the use of the crossconnect, be directed to any of the switches. Other wavelengths flow directly to the distribution network (i.e., “ λ -access”); again, these wavelengths can be either dedicated to a particular distribution tree, or have the capability of being switched to any of the trees. In addition, the wavelength-level crossconnect provides flexibility for the distribution wavelengths. For example, it allows distribution wavelengths to be shunted back to the same tree or switched from one tree to another, to allow direct communication between ORMA-Net customers attached to the same access node.

The size of the wavelength-level crossconnect is no larger than $(F + D) \times (F + D)$, where F is the total number of demultiplexed wavelengths from the feeder and D is the total number of wavelengths in the distribution trees. However, full connectivity may not be needed in the crossconnect. The crossconnect may also be capable of multicast so that a particular wavelength is delivered to a number of the distribution trees.

Fig. 12 assumes that the same wavelength spectrum is deployed in the feeder and distribution networks; if this is not the case, then *optical translators* are needed, as will be discussed in Section V-B.

Note that if only a small number of the wavelengths from a particular fiber need to be dropped in the access node, then it may be more economical to use a wavelength add/drop rather than fiber- and band-level demultiplexers. Two interesting emerging optical muxing technologies capable of tunable multiple wavelength drops are the acoustooptical tunable filter (AOTF) [24], [25] and liquid crystal switch [26].

V. DISTRIBUTION NETWORK ARCHITECTURE

The distribution portion of the network is fundamentally very different from the feeder network. As discussed in Section III, it is highly desirable that the distribution network be passive and low-cost. While the use of only passive components presents several technological challenges, the passivity of the network also provides opportunities for alternative architectures. For example, the absence of amplifiers in the

distribution network allows the deployment of a wider or totally different (e.g., 1300 nm) spectral window.

There is a range of options for carrying traffic in the distribution network, including WDM, TDM (time division multiplexing), and simple single-channel-per-fiber. In Section V-A, the various options are compared, with WDM preferred in most scenarios due to advantages in upgradability, scalability, and flexibility in allocating resources. The choice of the WDM spectral comb density in the distribution network is also an interesting topic, and is discussed in Section V-B.

Passive muxing technology is critical to delivering wavelengths to customers; a high-level description of various possible technologies is presented in Section V-C. Wavelength delivery can be achieved via either routed or broadcast techniques. The advantages and disadvantages of both approaches are explored in Section V-D, along with architectures for both modes of wavelength delivery. Finally, Section V-E presents architectures that incorporate survivability in the distribution network.

For simplicity, we often use the term “distribution *tree*,” however, as was shown in Fig. 1, the distribution topology could be a bus, a ring, etc.

A. Multiplexing Options in the Distribution Network

One key architectural bifurcation is whether or not to deploy WDM on the distribution fibers. In a single-channel-per-fiber non-WDM solution, individual fibers are run from each customer to an access node, with a single, relatively cheap, e.g., 1300 nm, wavelength carried on the fiber. In a new installation, this solution may be cheaper than WDM. Customers are typically located within a few miles of an access node, thus the extra fiber required should be more than compensated for by the cheap transceivers that can be used, and the absence of multiplexing and demultiplexing equipment. However, adding additional customers may be expensive. An increase in customer base may result in all the fibers in a sheath being consumed, requiring that an entire new fiber sheath be laid, which is a very costly upgrade. In addition, for local access providers that rent fiber from another carrier, a one-fiber-per-customer solution may pose serious limitations if they are unable to rent additional fiber.

WDM allows customer growth without requiring that additional fiber be laid, which is significant given that predicting customer demand is very difficult. Apart from using fibers more efficiently, WDM also affords much greater flexibility to each customer. For example, assume a customer normally has one wavelength worth of traffic, but occasionally needs a second wavelength. WDM allows an additional wavelength to be temporarily provided on the same fiber. In the single-channel solution, a second fiber would need to be permanently tied up at the customer. As another example, assume that a customer has multiple data formats to transmit, such that carrying each traffic type transparently in a separate wavelength is the easiest transmission method (refer back to Fig. 4). Again, WDM easily provides a number of wavelengths as needed, whereas in the single-channel scenario, the customer must have multiple fiber runs. WDM also provides a degree of

independence among customers; e.g., a customer can upgrade its data rate without affecting others.

An alternative option for multiplexing traffic from multiple customers (which can optionally be used in conjunction with WDM) is electronic TDM. For example, the traffic from four OC3 customers can be multiplexed onto a single-wavelength SONET ring carrying an OC12 signal, with an OC12 add/drop multiplexer at each of the customer premises. However, if one of the customers wants to upgrade its connection rate from OC3 to OC12, the signal rate of the wavelength must be upgraded to OC48, which requires equipment upgrades at *all* of the customer premises. Thus, TDM does not afford the same degree of independence among users as does WDM. Another disadvantage of electronic-based TDM schemes is that the signals must be electronically processed, which is not amenable to carrying a range of data rates and types. In addition, the traffic of one customer passes through the premises of another, which may present security issues.

Another multiplexing option is to implement an optical Media Access Control (MAC) protocol. In such a scheme, several users share a single wavelength, with a scheduler controlling when each user can transmit. Several optical MAC protocols have been proposed in the literature, e.g., [27], [28]. The advantage of a MAC protocol is that the signals can remain in the optical domain. The disadvantages are the need for burst receivers, buffers, and synchronization, and the overhead of scheduling.

Overall, the more graceful upgradability and scalability and the increased flexibility of WDM favors its deployment in the distribution network. A non-WDM solution is favored only if the need for upgrades is rare.

B. Wavelength Density in the Distribution Network

The economics of the feeder ring justifies deploying very high quality, powered components and very dense WDM. The feeder will carry up to 100 wavelengths per fiber, with say 50 GHz wavelength spacing. One strategy is to deploy the same dense wavelength scheme in the distribution network as in the feeder ring. In this solution, depicted in Fig. 13(a), OT's do not have to be deployed in order to route the distribution wavelengths onto the feeder ring, although it may be necessary to include optical amplifiers on the feeder ring at some, or all, of the access nodes.

Such a dense wavelength spectral comb, however, may be too expensive for the distribution portion of the network. It may be preferable to use fewer, more coarsely spaced wavelengths, e.g., 16 wavelengths with 200 GHz spacing. This allows the deployment of cheaper, more tolerant components in the distribution network, although it requires OT's to be used in order to map distribution wavelengths to feeder wavelengths, as shown in Fig. 13(b). It is desirable that these OT's be at least digitally transparent. Note that the non-WDM solution is the most extreme extension of the coarse WDM architecture.

Fewer wavelengths per fiber will result in fewer customers per distribution tree, and hence more trees. This may be advantageous for distribution architectures with shared wavelengths

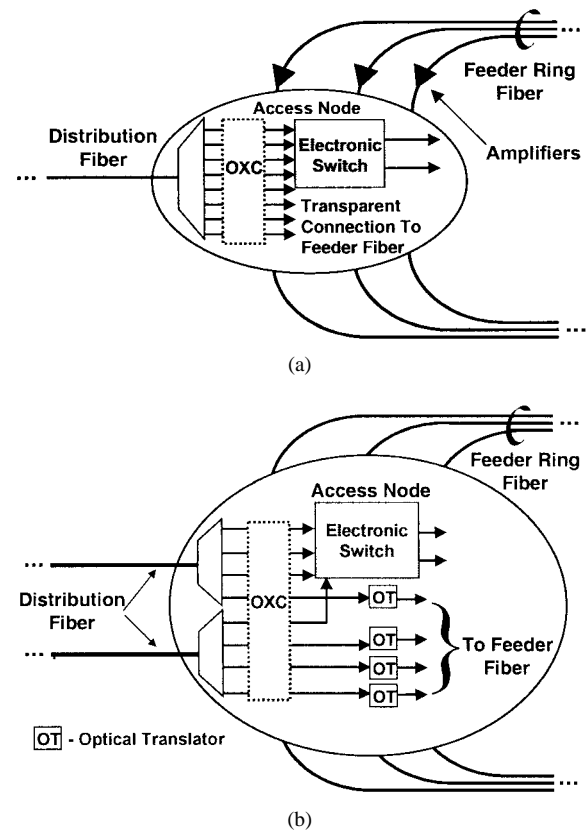


Fig. 13. The feeder fiber is likely to carry up to 100 wavelengths with 50 GHz spacing. (a) If the same spectral density is used in the distribution network, then wavelengths can be routed directly from distribution to feeder. Amplifiers are likely needed on the feeder ring at some, or all, of the access nodes. (b) It may be preferable to use fewer, more coarsely spaced wavelengths in the distribution network. This allows the deployment of cheaper, more tolerant components in the distribution network, although it requires *optical translators* (OT's) to be used in order to map distribution wavelengths to feeder wavelengths, and results in more distribution fibers.

where the technology favors fewer customers per tree due to loss considerations (see Section V-D2). Alternatively, the absence of amplifiers in the distribution network may allow a wider spectrum to be used, so that coarse spacing does not necessarily result in fewer wavelengths per fiber.

One advantage of deploying OT's is that more loss budget can be assigned to the distribution network, and perhaps fewer optical amplifiers may be needed on the feeder ring because the OT performs regeneration with reshaping [29]. Another advantage of OT's is that they may simplify network management and reduce the blocking probability of on-demand wavelength connections between customers on the same ORMA-Net, i.e., customers do not have to communicate via identical wavelengths.

C. Passive Distribution Technology

One of the most important components in a WDM distribution network is the passive wavelength add/drop (WAD). Such components will be deployed in the outside plant, necessitating that they function over a temperature range of, say, -20 to $+80$ °C. Passive WAD's designed for indoor use typically suffer from a relatively high temperature coefficient, where the dropped wavelength could vary as much as 1 nm over an

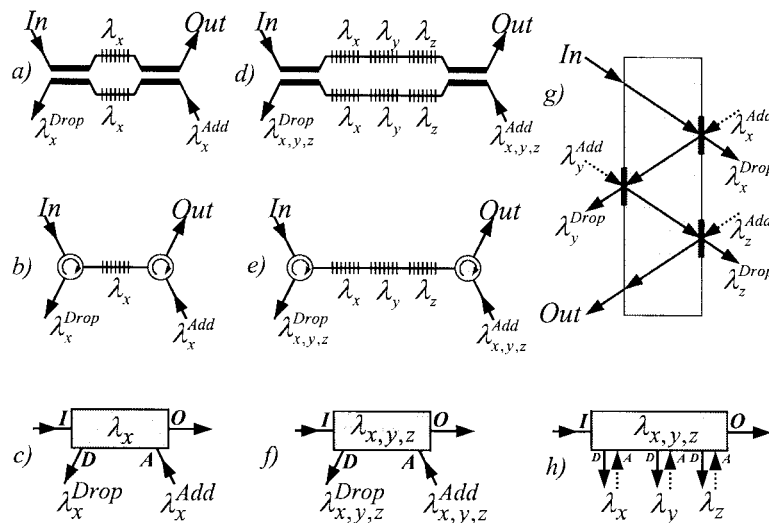


Fig. 14. Diagrams in the first column represent a single-wavelength WAD, with a grating filter used as illustration; the second column represents a multiple-wavelength WAD with a single-fiber-pair add/drop, also using grating technology as illustration; the third column represents a multiple-wavelength multiple-fiber-pair WAD, illustrated by thin-film technology.

operating temperature range of 100 °C. This wide variation could be very problematic depending on the density of the WDM scheme. However, recent advances in technology have greatly improved the performance of such devices. Temperature stabilization, through strictly passive means, has reduced the variation to less than 100 pm over a similar temperature range [30]. The two leading candidate technologies for passive temperature-stabilized WAD's are fiber Bragg grating (FBG) filters and thin-film filters (also known as dielectric coating filters).

Fig. 14 illustrates some examples of these WAD technologies in different configurations. The diagrams in the first column represent a single-wavelength WAD, with grating technology used as illustration; the second column represents a multiple-wavelength WAD with a single-fiber-pair add/drop, also using grating technology as illustration; the third column represents a multiple-wavelength multiple-fiber-pair WAD, illustrated by thin-film technology. Below, we present a high level description of the operation of these passive WAD's in order to clarify their needed functionality. For examples of alternative configurations and more operational details, see [13].

Fig. 14(a) illustrates an interferometer-based grating design, as described in [31]. Wavelength λ_x is reflected from the In port to the Drop port, while all other wavelengths pass from In to Out. The device is symmetric, so that if a signal enters on the Add port, λ_x is reflected to the Out port and the remaining wavelengths, if any, are sent to the Drop port. A variation of this device is shown in Fig. 14(b) where the couplers are replaced by circulators; the operation is similar. We use the functional diagram of Fig. 14(c) to represent any single-wavelength WAD device.

One can place multiple gratings in series, with either the coupler or circulator design, in order to construct multiwavelength single-fiber-pair WAD's. This is illustrated in the second column of Fig. 14, where λ_x, λ_y , and λ_z are reflected from the input signal to a single drop port. Any other wavelengths are

passed through from In to Out. As with the single grating filter, the device is symmetric. The functional schematic of Fig. 14(f) is used to represent multiwavelength, single-fiber-pair WAD's. (Note that the AOTF and liquid crystal switches mentioned in Section IV-F are examples of multiwavelength single-fiber-pair add/drop muxes, however, they are not passive devices.)

The WAD of Fig. 14(g) uses multiple dielectric thin-film stacks, where each stack passes a particular wavelength and reflects all others [32]. For example, the first dielectric stack passes λ_x to the drop port while reflecting all other wavelengths. Input wavelengths other than λ_x, λ_y , and λ_z are reflected by all three dielectric stacks to the Out port. An important difference between the thin-film device and the multiwavelength grating filter is that each dropped wavelength is output on a separate fiber. (This is functionally equivalent to placing in series multiple devices from the first column.)

To reduce coherent crosstalk, the add and drop functions depicted in Fig. 14 may need to be implemented with a cascade of two devices, one for drop, then one for add. Also, note that the technologies shown in Fig. 14 can also be used to provide add/drop functionality in the feeder.

The WAD's can be deployed in a variety of configurations depending on the required functionality. For example, assume two WAD's are used to add and drop wavelengths from two counter-directional fibers, as might be needed in a restorable architecture (see Section V-E). Three possible configurations are shown in Fig. 15, all of which are based on the functional diagram of Fig. 14(c). Fig. 15(a) depicts a four-fiber configuration, where the add and drop lines from the two WAD's remain distinct. Essentially, two independent WAD's are deployed in parallel. This configuration could be used, for example, if four fibers are run from the add/drop point to the customer premises. As shown in the figure, λ_x is added/dropped to/from the West-to-East fiber, and λ_y is added/dropped to/from the East-to-West fiber. λ_x may or may not equal λ_y .

Running four fibers to the customer premises may be too expensive, and hence, it may be desirable to run just two fibers.

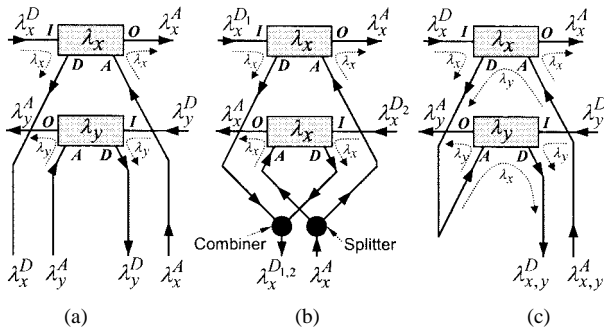


Fig. 15. Three different WAD configurations for adding/dropping wavelengths to/from two counter-directional fibers. The D superscript indicates Drop, and A indicates Add. (a) A four-fiber configuration, with separate fibers for each of the drop and add lines. (b) A two-fiber configuration, where the drop lines are combined onto a single drop fiber, and a single add fiber is split into two add lines. (c) Alternative two-fiber configuration, where the WAD's are judiciously interconnected to maintain distinct signals.

Fig. 15(b) illustrates a two-fiber configuration, where the two WAD drop ports are passively combined onto a single drop fiber, and a single add fiber feeds into two WAD add ports via a passive splitter. In this configuration, the two WAD's must be identical, i.e., both WAD's add/drop λ_x . Since the drop ports of the two WAD's are combined, λ_x^{D1} and λ_x^{D2} cannot both be present at the same time; the traffic source(s) must assure that this condition is met.

An alternative two-fiber scheme is shown in Fig. 15(c). Here, λ_y from the add fiber enters the Add port of the λ_x mux and passes through to the Drop port. It then feeds into the Add port of the λ_y mux, which sends it onto the East-to-West fiber. In a similar fashion, λ_x from the West-to-East fiber passes through both WAD's and ultimately ends up on the drop fiber. In this configuration, λ_x can not equal λ_y . Two distinct connections are maintained to and from the customer premises through the use of different wavelengths and the judicious interconnections of the WAD's shown in the figure; in Fig. 15(a), two distinct connections were provided through the use of diverse fibers.

Another important component based on filter technology is the passive band splitter. The input wavelength comb is split into multiple bands, each of which is directed to a separate output. Either thin-film or fiber-grating technology may be used to fabricate such devices. One simple implementation is a high/low band splitter, with one input and two outputs. Assuming wavelengths 1 through W are present at the input, then wavelengths 1 through B are directed to one output, and $B + T$ through W are directed to the other output. The value of B can be specified when the component is fabricated. Note that wavelengths $B+1$ through $B+T-1$, which fall within the splitter transition region, are 'wasted'; the value of T depends on the sharpness of the passband skirts.

D. Wavelength Distribution Architectures

Wavelengths can be delivered to a customer using various routed or broadcast technologies, with several important trade-offs in the two methods. Routed wavelengths provide better security and result in less loss because a given wavelength is only received by one customer on a distribution tree. Con-

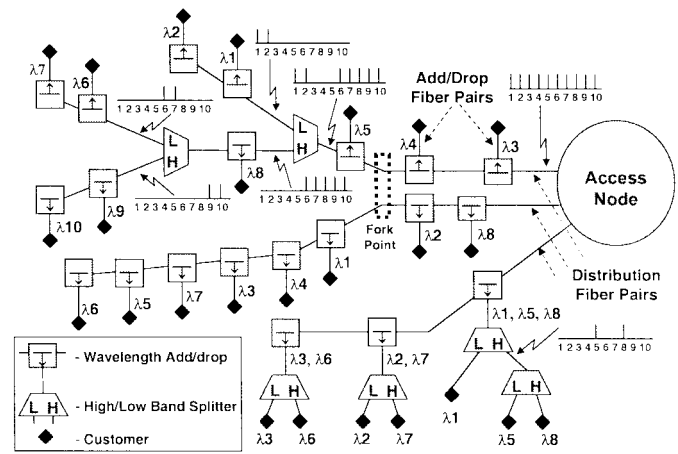


Fig. 16. Example of an architecture for routing wavelengths. Three fiber pairs (each drawn as a single line) emanate from the access node, one in a tree structure, one in a bus topology, and the other in a "bus of trees" topology. A combination of WAD's and high/low band splitters are used at the branch points. Spectral combs are shown at various points on the tree to indicate the wavelengths that are present at that point.

versely, shared wavelengths are less secure and more lossy, but they allow for more efficient use of wavelengths, they may result in lower blocking probability for establishing intra-ORMA-Net connections (especially if OT's are not used), and they are more amenable to delivering broadcast or multicast traffic. Ideally, the distribution network should be capable of delivering both routed and shared wavelengths.

1) *Routed Wavelength Architectures*: One example of a routed-wavelength architecture is shown in Fig. 16. In this figure, three distribution fiber pairs emanate from the access node (one fiber in each of the pairs carries traffic to the customers, and the other fiber carries traffic from the customers). The top fiber pair is deployed in a tree topology, the middle pair in a simple bus topology, and the bottom fiber pair in a "bus of trees" topology. In all three topologies, passive WAD's are used to pick out a particular wavelength from a fiber. The WAD's are deployed in a single-fiber-pair add/drop configuration, using the technologies represented by Fig. 14(c) (if single-wavelength drop) or Fig. 14(f) (if multiple-wavelength drop). The single add/drop fiber pair is run to the customer premises.

The tree topology shown in the figure also employs passive high/low splitters at the branch points of the tree. The first splitter encountered on the tree sends wavelengths 1 and 2 to the upper branch of the tree and wavelengths 6 through 10 to the lower branch. (Spectral combs are shown in the figure to indicate which wavelengths are present at a particular point of the tree.) We assume that wavelengths 3 through 5 fall in the transition region of the splitter due to nonideal roll-off. As many of these wavelengths as possible are assigned to the customers at the root of the tree so that they are not "wasted" (the tree in the figure represents the ideal case where no wavelengths are wasted).

The "bus of trees" also illustrates the principle of pulling out wavelengths prior to their being "lost" in a splitter transition band, however, the splitters and WAD's are used in different combinations. For example, at the first drop point on the

bus of trees, a WAD filters off wavelengths 1, 5, and 8. These are passed through a coarse high/low splitter that directs wavelength 1 to the left output, and 5 and 8 to the right output. A second coarse splitter then partitions wavelengths 5 and 8. These high/low splitters are not located on the main fiber pair, so that the transition region losses do not affect wavelengths downstream.

In addition to using WDM technology at branch points, bifurcation can also be achieved through spatial division. For example, the upper two fiber pairs initially lie along the same path, but spatially diverge after the fork point.

One of the important features of the architecture is that it eliminates the need to align two WAD's on the same fiber pair, i.e., two WAD's that both drop wavelength i are never deployed in series. This is highly desirable due to the drift that may occur in the WAD's. Band splitters have a much wider passband so can remain aligned with the WAD even if the WAD drifts. Of course, there has to be a guarantee that the sources be roughly locked to the drift of the WAD's; various techniques exist for this [33].

2) *Shared Wavelength Architectures*: In addition to a set of permanently assigned wavelengths, there should be a set of wavelengths in the distribution network that can be assigned on demand. These wavelengths are shared over time among customers (on the order of seconds to hours), resulting in more efficient utilization. For example, a company might need an extra wavelength for several hours per week. Rather than having the wavelength be idle the remainder of the time, it can be assigned dynamically to other users.

One solution for distributing on-demand wavelengths is to deploy broadcast couplers at each customer premises junction point. For example, the couplers can be 10 dB couplers where 1/10 of the signal power is dropped to the customer and the remaining 9/10 passes through. If the structure is a tree, then the coupler split ratios used at the tree branch points depend on how many customers are on the various branches of the tree. One limitation of this approach is the loss that accumulates from passing the optical signal through a number of couplers and the possibility of uneven power at the drop points. It is desirable that the distribution network be passive, which eliminates the possibility of adding amplifiers. Additionally, adding a new customer may result in having to adjust the couplers of the other customers along the path.

A variation of this approach is to have a set of shared wavelengths per distribution tree, where each customer is capable of receiving only a subset of these wavelengths. Thus, each wavelength is not dropped at every customer, resulting in less loss. The wavelengths must be distributed such that any two customers that potentially may communicate via an optical connection have at least one wavelength in common. However, if the distribution architecture uses a coarse WDM structure, such that optical translators (OT's) are used (see Section V-B), then this requirement is lifted. We assume that a customer-to-customer signal has to pass through an access node, even if the customers are on the same distribution tree.

Another solution is to make use of *tunable* WAD's, where a customer's WAD can be tuned to drop any of the shared wavelengths. At any given time, only the customer using a

particular shared wavelength has its WAD tuned to drop that wavelength; thus, the loss is approximately the same as for routed wavelengths. While this solution appears to be very elegant, remotely tunable *passive* WAD's are not currently available; however, some interesting recent work [34] proposes devices with similar capabilities.

Shared wavelengths can be delivered either using the same fiber pair as is used for distributing dedicated wavelengths (this is likely the case for the tunable WAD solution), or using a separate fiber pair that carries only the on-demand wavelengths.

E. Distribution Network Reliability

While survivability is extremely important in the feeder ring since a failure can impact the service of all customers, resiliency against failures in the distribution network may not be as critical depending on individual customer requirements. For example, some customers may be willing to accept a single point of failure if it results in significantly lower cost. Furthermore, implementing protection in the distribution network is somewhat more challenging given its passive architecture, as shown below.

We focus our attention on recovering from link failures in the distribution network since passive-component failures should be a rare event. (We do not consider protection of the add/drop lines that run from the customer junction point on the distribution fiber to the customer premises, though of course, this can be protected.) In order to recover from a distribution fiber cut, there must be a redundant path from the customer junction point to the feeder. This can be achieved through a ring topology, where the ring passes through one or more access nodes, as illustrated by the single- and double-homed distribution rings shown in Fig. 1. Alternatively, a customer could be attached to two disjoint distribution buses or trees that are rooted on the same or different access nodes.

Here, we consider a single-homed distribution ring, as shown in Fig. 17. Two fibers are deployed in the ring, carrying traffic in opposite directions. One 'stand-alone customer' and one 'active ring of customers' are shown attached to the ring. The active ring could represent, for example, a legacy SONET ring. For simplicity, the active ring is shown in the figure as supporting a single wavelength, however, it could also be a WDM ring. The recovery mechanisms are different for the stand-alone user and the active ring customers, as described below.

1) *Stand-Alone-Customer Protection*: The stand-alone customer, which is assigned λ_1 in the figure, makes use of a passive splitter and a passive combiner in order to achieve redundancy. The dashed-line box at the point of attachment to the distribution ring represents the WAD configuration shown in Fig. 15(b). First, we consider traffic sent from the customer premises to the access node. After passing through the two-way splitter, the traffic is sent in both the clockwise and counterclockwise directions to the access node. The access node, which has two λ_1 receivers corresponding to the two directions, selects the better of the two signals. Thus, if there is a fiber cut such that no signal is received on the clockwise fiber, the access node continues to receive data via




Fig. 17. An architecture for providing restorative capabilities in the distribution network. The stand-alone user relies on a splitter and a combiner in combination with redundant transmitters and receivers in the access node. The active ring can operate as if the access node were a node on the ring. The distribution ring is connected to one access node here, however, double-homed rings are also possible for additional redundancy.

the counterclockwise fiber. (This is essentially a path-based $1 + 1$ protection scheme.)

We cannot simply reverse the above operation for traffic sent from the access node to the customer. The combiner passively combines the signals arriving on the clockwise and counterclockwise fibers, without any selectivity; thus, the access node must transmit on only one fiber at a time to the stand-alone user. (Since it is desirable that the distribution ring be passive, a signal selector cannot be deployed in place of the combiner.) For example, under normal operation the access node may transmit on the clockwise fiber to the customer. If there is a failure such that the path to the customer is down, then the access node transmits on the counterclockwise fiber.

Thus, for the stand-alone customer configuration shown in Fig. 17, it is the responsibility of the access node to detect failures in the passive distribution ring and take appropriate action. No actions need to be taken by the customer, except perhaps, to notify the access node about the loss of signal.

Alternatively, one could remove the splitter and combiner, and replace the WAD configuration of Fig. 15(b) with that shown in Fig. 15(c). This architecture requires that the customer utilize two wavelengths to and from the access node, as was described in Section V-C. The two wavelengths allow two distinct copies of the traffic to travel from the access node to the customer premises. Thus, the access node can transmit simultaneously on both distribution fibers, relying on the customer to select the better signal. This solution requires active participation of the customer and more wavelengths, however, it is operationally simpler.

Another alternative is a four-line drop to the customer [i.e., Fig. 15(a)], where again, the customer must actively select the better signal, however, only one wavelength is required.

2) *Active-Ring Customer Protection*: Protection is somewhat more straightforward for the customers attached to the active ring, which itself has two counter-propagating fibers. The two distribution ring fibers can be treated as extensions of the active ring fibers, i.e., operation can occur as if there is a single ring extending from the customers up to, and including, the access node. The dashed-line box at the point of attachment to the distribution ring represents the four-fiber WAD configuration shown in Fig. 15(a). [In Fig. 17, λ_2 is used in both directions of the active ring, however, different wavelengths could be used in the different directions, as was represented by λ_x and λ_y in Fig. 15(a).]

Standard path-based or link-based protection can be implemented for any cut within the active ring or the distribution ring, assuming that the access node is properly equipped. (Two active rings extending from the same distribution network could implement different protection schemes as long as the access node equipment corresponding to each active ring's wavelength is appropriate.) Thus, restoration is achieved without any special functional requirements placed on the active ring customers.

VI. SUMMARY

We have investigated many of the fundamental issues of architecting a high-end WDM-based access network. An array of architectural designs were presented for both the feeder and distribution portions of the network, with an emphasis on keeping the network as flexible and scalable as possible.

The importance of transparency in delivering enhanced services as well as providing architectural flexibility is a theme that ran throughout the discussion. The crucial role of advanced add/drop technology in an access network was also emphasized. This is an area where more innovation is still needed. Configurability is clearly desirable to support on-demand connections and changing user demand. We presented a hierarchical approach to configurability in the feeder, and illustrated how configurability could enhance the flexible distribution of wavelengths to customer premises.

Overall, the large bandwidth, flexibility, and scalability of the architectures presented will allow an ORMA-Net to evolve to meet growing customer demand. Many of the design principles presented here can serve as guidelines for deploying ORMA-Nets in the near future.

REFERENCES

- [1] J. M. Simmons *et al.*, "Optical regional access network (ORAN) project," in *Proc. OFC'99*, San Diego, CA, Feb. 21–Feb. 26, 1999.
- [2] R. Wagner *et al.*, "MONET: Multiwavelength optical networking," *J. Lightwave Technol.*, vol. 14, pp. 1349–1355, June 1996.
- [3] S. B. Alexander, "A precompetitive consortium on wide-band all-optical networks," *J. Lightwave Technol.*, vol. 11, pp. 714–735, May/June 1993.
- [4] I. P. Kaminow *et al.*, "A wideband all-optical WDM network," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 780–799, June 1996.
- [5] N. R. Dono *et al.*, "A wavelength division multiple access network for computer communication," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 983–994, Aug. 1990.
- [6] E. Hall *et al.*, "The rainbow-II gigabit optical network," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 814–823, June 1996.
- [7] C. A. Brackett *et al.*, "A scalable multiwavelength multihop optical network: A proposal for research on all-optical networks," *J. Lightwave Technol.*, vol. 11, pp. 736–753, May/June 1993.

- [8] S. M. Gemelos, D. Wonglumsom, and L. G. Kazovsky, "LEARN: Stanford/sprint reconfigurable WDM metropolitan area networks," in *Proc. LEOS'98, 11th Annu. Meeting*, Orlando, FL, Dec. 1-4, 1998, pp. 245-246.
- [9] S. Johansson, "Metropolitan optical network," in *Proc. LEOS Summer Topicals: Broadband Optical Networks*, Aug. 5-9, 1996, pp. 61-62.
- [10] F. Arecco *et al.*, "A transparent, all-optical, metropolitan network experiment in a field environment: The 'PROMETEO' self-healing ring," *J. Lightwave Technol.*, vol. 15, pp. 2206-2213, Dec. 1997.
- [11] T.-H. Wu, *Fiber Network Service Survivability*. Norwood, MA: Artech House, 1992.
- [12] A. F. Elrefaie, "Multiwavelength survivable ring network architectures," in *Proc. ICC'93*, Geneva, Switzerland, May 23-26, 1993, paper 48.7, pp. 1245-1251.
- [13] R. Ramaswami and K. Sivarajan, *Optical Networks: A Practical Perspective*. San Francisco, CA: Morgan Kaufmann, 1998, pp. 83-119, pp. 359-389, pp. 444-446.
- [14] ANSI, *Synchronous Optical Network (SONET)—Automatic Protection Switching*, T1.105.01, 1995.
- [15] A. F. Elrefaie, "Self-healing WDM ring networks with all-optical protection path," in *Proc. OFC'92*, San Jose, CA, Feb. 2-7, 1992, pp. 255-256.
- [16] O. Gerstel, "Opportunities for optical protection and restoration," in *Proc. OFC'98*, San Jose, CA, Feb. 1998, pp. 269-270.
- [17] P. A. Bonefant, "Optical layer survivability: A comprehensive approach," in *Proc. OFC'98*, San Jose, CA, Feb. 1998, pp. 270-271.
- [18] A. A. M. Saleh, "Transparent optical networks for the next generation information infrastructure," in *Proc. OFC'95*, San Diego, CA, Feb. 26-Mar. 3, 1995, p. 241.
- [19] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, "On the value of wavelength-add/drop in WDM rings with uniform traffic," in *Proc. OFC'98*, San Jose, CA, Feb. 22-Feb. 27, 1998, pp. 361-362.
- [20] E. Modiano and A. Chiu, "Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming," in *Proc. GLOBECOM*, Sydney, Australia, Nov. 8-12, 1998.
- [21] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, "Quantifying the benefit of wavelength-add/drop in WDM rings with distance-independent and dependent traffic," *J. Lightwave Technol.*, vol. 17, pp. 48-57, Jan. 1999.
- [22] Y. Sun *et al.*, "Ultra wide band erbium-doped silica fiber amplifier with 80 nm of bandwidth," *Optic. Amplifiers and Their Applications*, Victoria, BC, Canada, July 21-23, 1997, PD2-1-PD2-5.
- [23] H. Toba *et al.*, "An optical FDM-based self-healing ring network employing arrayed waveguide grating filters and EDFA's with level equalizers," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 800-813, June 1996.
- [24] D. A. Smith *et al.*, "Integrated-optic acoustically-tunable filters for WDM networks," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 1151-1159, Aug. 1990.
- [25] D. A. Smith *et al.*, "Evolution of the acousto-optic wavelength routing switch," *J. Lightwave Technol.*, vol. 14, pp. 1005-1019, June 1996.
- [26] J. S. Patel and Y. Silverberg, "Liquid crystal and grating-based multiple-wavelength cross-connect switch," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 514-516, May 1995.
- [27] B. Mukherjee, *Optical Communication Networks*. New York: McGraw-Hill, 1997, pp. 118-142.
- [28] A. C. Kam *et al.*, "Toward best-effort services over WDM networks with fair access and minimum bandwidth guarantee," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1024-1039, Sept. 1998.
- [29] P. E. Green, F. J. Janniello, and R. Ramaswami, "WDM protocol-transparent distance extension using R2 remodulation," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 962-967, June 1996.
- [30] Highwave Optical Technologies, Add/drop Multiplexers Product Brochure, 1998.
- [31] F. Bilodeau *et al.*, "High-performance wavelength-division-multiplexing/demultiplexing using an all-fiber Mach-Zender interferometer and photoinduced Bragg grating," in *Proc. OFC'95*, San Diego, CA, Feb. 26-Mar. 3, 1995, pp. 130-132.
- [32] M. A. Scobey and D. E. Spock, "Passive DWDM components using microplasma optical interference filters," in *Proc. OFC'96*, San Jose, CA, Feb. 25-Mar. 1, 1996, pp. 242-243.
- [33] D. Mayweather *et al.*, "Wavelength tracking of a remote WDM router in a passive optical network," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 1238-1240, Sept. 1996.
- [34] R. Giles *et al.*, "Highly efficient light-actuated micromechanical photonic switch for enhanced functionality at remote nodes," in *Proc. OFC'98*, Postdeadline Paper, San Jose, CA, Feb. 1998.