

Optical Regional Access Network (ORAN) Project

Jane M. Simmons and Adel A. M. Saleh, *AT&T Labs*

100 Schulz Dr., Red Bank, NJ, 07701, jane@research.att.com, (732) 345-3220, (732) 345-3040 (fax)
Ondria J. Wasem, *Bellcore, Red Bank, NJ*, Elisa A. Caridi, *Bell Labs, Lucent Technologies, Holmdel, NJ*
and Richard A. Barry, *Sycamore Networks, Tewksbury, MA (work performed at MIT Lincoln Lab)*

1. Introduction

Wavelength Division Multiplexing (WDM) technology has emerged as the leading solution for high-speed transmission and is beginning to dramatically change the underlying characteristics of backbone networks. While WDM is likely to soon dominate the backbone, its role in an access environment is just beginning to take shape. The Optical Regional Access Network (ORAN) Project, an offshoot of the DARPA-sponsored Multiwavelength Optical Networking (MONET) Consortium [1], was a one-year architectural study that investigated introducing WDM into the access environment. The resulting ORAN architecture is a scalable, flexible, cost-effective access network that delivers both huge bandwidth and a high degree of upgradability to high-end customers.

An access network differs from a long distance or interoffice backbone network in several key areas. 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 be finer grain and more flexible; aggregation is a more significant function than transport. Also, the costs of an access network are shared by much fewer customers than that of a backbone network; thus, designing a cost-effective network is extremely important.

Transparency, the ability to directly carry signals independent of data rate and format, plays a more significant role in the access environment than in a long distance backbone. Many of the impairments that hinder the deployment of a transparent long-distance backbone are not as significant in an access network, primarily due to the short distances involved. Also, the ability to route a range of traffic types in their native format is growing in importance as customers request that the access network provide virtual LAN functionality across a range of data protocols and data rates. Furthermore, one of the prime benefits of transparency is that it can be used to simplify the delivery of electronic services, as exemplified in the ORAN architecture.

ORAN provides several novel features in addition to serving the traditional role of aggregating electronic traffic and delivering the traffic to backbone networks. The ORAN architecture can deliver WDM all the way to the end-user, which greatly enhances the upgradability and flexibility of the network. Wavelengths can be provided on a fixed basis, or they can be provided on-demand. The ORAN architecture is capable of exploiting the huge bandwidth afforded by optical technology. For example, it can easily support customers with very large demand, e.g., multiple wavelengths worth of traffic. In addition, portions of the architecture are designed to provide the option of using only a fraction of the transport capabilities of a wavelength in order to simplify the technology. The architecture also uses optical technology to provide more flexible placement of electronic switches.

2. Overview of the ORAN Architecture

An access network architecture can be functionally partitioned into 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.

A high-level view of the ORAN architecture is shown in Figure 1. The feeder network has a ring topology on which are located a set of Access Nodes and Egress Nodes. Most of the aggregation and switching in ORAN 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. It is expected that a single ORAN area would provide access to multiple network backbones, such as Internet Protocol (IP), Asynchronous Transfer Mode (ATM), and Frame Relay (FR). The distribution network in ORAN can be one of several topologies depending on the required redundancy, e.g., tree, bus, or ring. A key feature of the distribution network is that it is totally *passive*.

An ORAN area is approximately 100 to 500 square miles, serving about 500 to 2000 customers. The feeder ring circumference is expected to be in the 25 to 50 mile range, with customers located no more than a few miles from an access node. In densely populated areas, it may be necessary to deploy more than one ring. ORAN customers are high-end users such as businesses, government facilities, cable head-ends, and campuses.




Figure 1. High-level view of the ORAN architecture. 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. The customers are high-end, such as businesses or campuses. Access is provided to deliver a range of electronic services.




Figure 2. Multiple feeder fibers enter an access node. Both fiber and wavelength bypass, and configurable add/drop are supported. Some wavelengths are terminated on an electronic switch, while others are routed transparently to the distribution network.

2.1 Feeder Network

A ring topology was chosen for the ORAN feeder due to the inherent reliability of a ring. WDM is deployed around the ring, with roughly 64 wavelengths per fiber. The number of fibers in the ring is likely to be in the 2 to 30 range (including protection fiber). An example of an access node architecture is shown in Figure 2. Optical bypass of an access node is supported on both a wavelength level and a fiber level. In addition, the access nodes are highly configurable; this is represented in the figure by the 2x2 switches that optionally drop or pass a wavelength. Some of the wavelengths may terminate on an electronic switch while others pass transparently between the distribution network and the feeder ring.

The ORAN architecture includes several options for aggregating electronic traffic, two of which are presented here. In one solution, the electronic service types requested by a customer are multiplexed in a SONET Mux at the customer premises, as shown in Figure 3. A SONET switch in the access node multiplexes traffic together from various customers 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 from various customers occurs in the switch/router, and the traffic is sent out on the feeder ring to the appropriate egress node.

A powerful variation of this architecture, shown in Figure 4, is to dedicate a wavelength per electronic service type requested by the customer. Thus, Customer 1 would require two 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. Although the architecture requires more wavelengths, overall, it still could be more economical. One advantage is that it eliminates the need for multiplexing equipment at the customer premises. Also, it allows a range of data formats to be transported in their native format, e.g., to provide extended LAN capabilities.

Furthermore, this latter option provides additional flexibility in where electronic switches are placed. Rather than deploying each type of switch in every access node, the ORAN architecture allows switches to be placed in only a subset of the nodes. For example, the access node shown in Figure 4 could include just an IP router. Any non-IP traffic is routed transparently from the distribution network onto the feeder ring and terminated in an appropriate switch in another access node. The ability to eliminate switches from some nodes is extremely important for scalability, maintainability, and cost effectiveness.

Protection of the ORAN feeder network occurs in both the optical and electronic domains. One architecture is a unidirectional line-switched ring, where the line switching function is performed through the use of optical 2x2 switches that operate on service and protection fiber pairs. In addition, the transceivers in the access and egress node switches are 1xN protected. Thus, the feeder is protected against a simultaneous fiber cut and transceiver failure.

2.2 Distribution Network

The ORAN distribution network contains only passive devices; there are no active switches and no amplifiers deployed in the field; this is highly desirable from a maintenance standpoint.

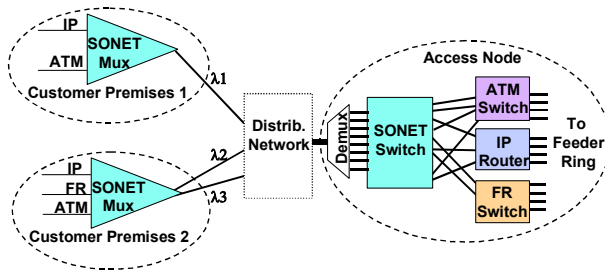


Figure 3. At the customer premises, a SONET Mux is used to multiplex the various electronic service types. In the access node, a SONET switch distributes each service type to the appropriate switch/router. The details of the distribution network are not shown.

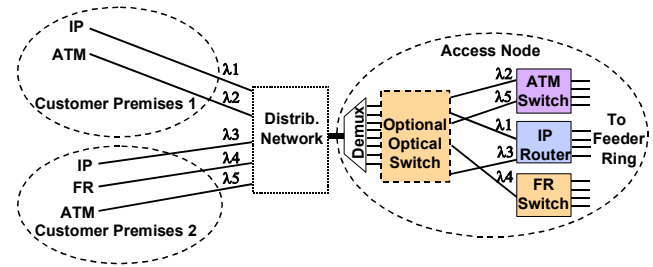


Figure 4. In this option, each customer service type is assigned to a separate wavelength, so that no muxing of service types takes place at the customer premises. In the access node, each wavelength is either permanently tied to a particular switch, or to add flexibility, an optical switch can be used.

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. In a new installation, the non-WDM solution may be cheaper. The extra fiber required in the non-WDM solution should be more than compensated for by the cheap laser that can be used, and the absence of multiplexing and demultiplexing equipment. However, adding additional customers becomes very expensive when the installed fiber is exhausted. WDM allows much easier upgrades, affords much greater flexibility to each customer, and provides a greater degree of resource sharing. WDM also provides a degree of independence among customers, in contrast to a TDM solution; e.g., a customer can upgrade its data rate without affecting the equipment required at other customer premises.

In deploying WDM in the distribution network, one strategy is to deploy the same dense wavelength scheme as in the feeder ring, and include an optical amplifier where the distribution fiber enters the access node. 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. This allows the deployment of more tolerant components in the distribution network, although it requires optical translators (preferably digitally transparent) to be used to map distribution wavelengths to feeder wavelengths (a separate optical amplifier would then not be needed).

Wavelengths can be assigned to a customer on both a dedicated and a temporary basis. Dedicated wavelengths are delivered via routed technologies, e.g., passive wavelength add/drops (WADs) or passive wavelength band splitters, such that the wavelength of one customer does not pass through any other customer premises. In addition, a set of wavelengths in the distribution network can be set aside as ‘shared over time’, and be provided on-demand to a customer. One solution for distributing on-demand wavelengths is to deploy broadcast couplers at each customer premises junction point. One limitation of this approach is the loss that accumulates from passing the optical signal through a number of couplers. Another solution is to make use of tunable WADs, where a customer’s WAD can be tuned to drop any of the shared wavelengths. While remotely tunable passive WADs are not currently available, some interesting recent work [2] proposes devices with similar capabilities.

In addition, a particular wavelength can be shared at the same time by several users, for example, through the use of optical MAC protocols or electronic-based TDM (e.g., a SONET-based logical ring).

3. Summary

The ORAN project investigated many of the fundamental issues of deploying a high-end access network. Several architectural designs were investigated, with the emphasis on keeping the network as flexible and cost-effective as possible. The ORAN architecture accommodates a range of data types and data rates; it supports both permanent and on-demand connections. The large bandwidth, flexibility, and scalability of ORAN will allow it to evolve to meet growing customer demand.

References

- [1] Wagner, R., et. al., “MONET: Multiwavelength Optical Networking,” *Journal of Lightwave Technology*, Vol. 14, No. 6, June 1996, pp. 1349-1355.
- [2] Giles, R., et. al., “Highly Efficient Light-Actuated Micromechanical Photonic Switch for Enhanced Functionality at Remote Nodes,” Postdeadline Paper, *OFC'98*, San Jose, CA, February, 1998.