

# A Next-Generation Optical Regional Access Network

Mark Kuznetsov, Nan M. Froberg, Scott R. Henion, and Hemonth G. Rao, MIT Lincoln Laboratory

Jeff Korn, Axsun Technologies

Kristin A. Rauschenbach, Photonex Corp.

Eytan H. Modiano and Vincent W. S. Chan, Massachusetts Institute of Technology

## ABSTRACT

We describe an optical regional access network which combines electronic IP routing with intelligent networking functionality of the optical WDM layer. The optical WDM layer provides such networking functions as network logical topology reconfiguration, optical flow switching to offload traffic and bypass IP routers, wavelength routing of signals, protection switching and restoration in the optical domain, and flexible network service provisioning by reconfigurable wavelength connectivity. We discuss key enabling technologies for the WDM layer and describe their limitations. The symbiosis of electronic and optical WDM networking functions also allows support for heterogeneous format traffic and will enable efficient gigabit-per-second user access in next-generation Internet networks.

## OPTICAL NETWORKING: INTRODUCTION

The current explosion of communication traffic volume is driven by an insatiable appetite for Internet connectivity on geographic scales from intra-building to world-wide. This exponential growth of traffic volume and the demand for ever higher end-user data rates are expected to continue in the foreseeable future. Optical fiber communication technology has kept up with the growing traffic volume by expanding the use of wavelength-division multiplexing (WDM) technology, which allows multiple data channels at different optical wavelengths to be transmitted simultaneously over a single optical fiber. Communication system suppliers today are advertising WDM transmission systems with capacities greater than 1 Tb/s ( $10^{12}$  b/s) over a single fiber by means of more than a hundred channels at 10 Gb/s each. At the same time, communication service providers are laying fiber cables into the ground with more than 100 fibers per cable. Much of this capacity growth has been in point-to-point long-distance backbone transmission. Importantly, observation of the present growth indicates that the dominant future traffic on national and worldwide backbones will be in the form of Internet Protocol (IP)

packets; next-generation communication infrastructure should also support heterogeneous traffic formats, including synchronous optical network (SONET) and asynchronous transfer mode (ATM). High-capacity electronic IP routers are already used to switch the backbone network traffic, router capacities are expected to reach several terabits per second in the very near future.

While the growth of the backbone communication capacity has been tremendous, end-user access to this capacity is still expensive and limited to data rates of kilobits and megabits per second. Many research groups and developers in commercial companies are exploring ways to extend high-data-rate capability from the backbone to the user; gigabit-per-second data rates are desirable for some high-end users, such as business premise routers and high-speed workstations. The use of WDM technology in such metropolitan area access networks is a powerful approach that is currently being actively explored.

MIT Lincoln Laboratory is part of a consortium developing a next-generation optical WDM regional network architecture for high-speed user access to the IP Internet backbone. This work is conducted under the auspices of the Next Generation Internet-Optical Network for Regional Access with Multiwavelength Protocols (NGI-ONRAMP) program sponsored by the Defense Advanced Research Projects Agency (DARPA). This consortium includes partners AT&T, Cabletron, JDS Uniphase, MIT, and Nortel Networks. At present, optical long-distance backbone communication uses wavelength multiplexing for high-speed point-to-point data pipes; simple optical functionality, such as optical fiber cross-connects and fixed add/drop multiplexing, is being introduced into the backbone. Our optical access network design incorporates electronic IP routers interconnected by an intelligent WDM optical core in a symbiosis of electronic and optical networking elements and functions. We use WDM to provide important networking functionality in the optical layer of the network, such as network logical topology reconfiguration in response to changes in traffic demand, switching of high-speed data flows directly in the optical wavelength domain to offload traffic and bypass IP routers, wavelength routing of optical data signals to their

---

*This work is supported by the Defense Advanced Research Projects Agency (DARPA).*

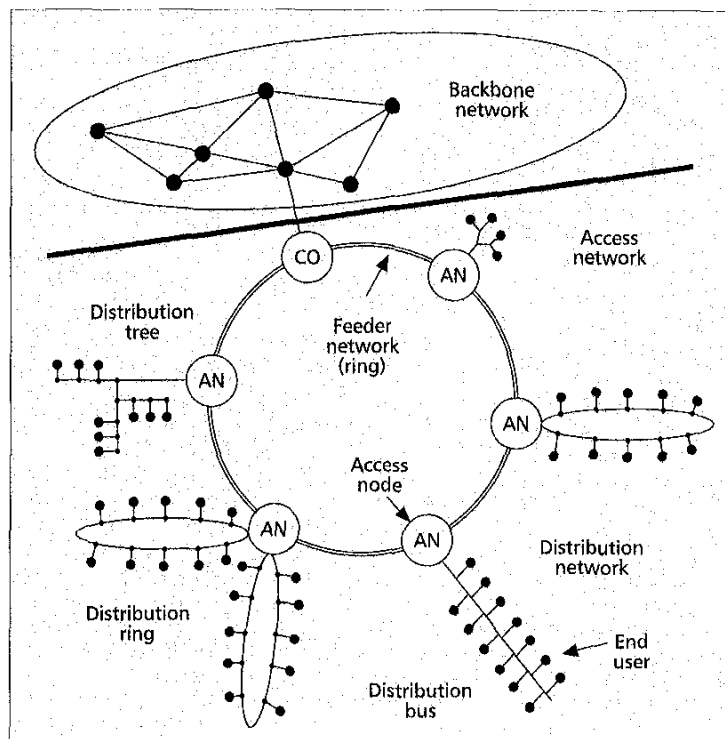
intended destinations, protection switching and service restoration in the optical domain in response to fiber cuts, and flexible and rapid network service provisioning using reconfigurable wavelength connectivity. By means of available transparent optical channels, this WDM multi-channel network also supports heterogeneous traffic formats, including SONET, ATM, Gigabit Ethernet, and analog-modulated optical signals. The major features of the proposed architecture will be implemented in an experimental testbed network. In this article we describe the main aspects of our physical network architecture, with the emphasis on WDM optical network configuration and elements required to provide the above optical networking functions.

Direct optical connections between network users also place demands on the optical transparency of the network, and we discuss issues related to signal degradation in such data connections. Other important issues of WDM optical networking which lie outside the scope of this article are the development of the network management and control systems to allow the IP layer of the network to efficiently utilize the optical WDM networking layer.

## PHYSICAL NETWORK ARCHITECTURE

The challenge in developing the physical architecture of the WDM network is to provide intelligent network functionality at the optical layer, as well as to allow optical and electronic switching layers to operate in synergy. Network scalability in terms of number of users and network capacity is also an important factor. In Fig. 1 we illustrate the key physical architecture features of our optical metropolitan access network. The access network is divided hierarchically into a feeder network and multiple distribution networks. End users are locally connected to the distribution networks, which in turn are connected to access nodes in the feeder network. In addition to connecting access nodes to one another, the feeder network has one or more connections to the communication backbone via backbone access nodes. In our implementation the feeder network has a ring physical topology, although other topologies are also possible, such as a mesh. Feeder access nodes can be interconnected by multiple fibers for greater capacity and route diversity; we use a dual-fiber ring feeder configuration for its simplicity while providing protection in case of fiber cuts.

The WDM feeder ring carries multiple wavelength channels  $\lambda_i$ , different wavelengths providing connectivity between different access nodes. Thus wavelength  $\lambda_1$  can connect nodes 1 and 3, while wavelength  $\lambda_5$  can connect nodes 2 and 6. We envision systems with 10 to 20 access nodes over a metropolitan area and 20 to 100 users on each distribution network. The feeder ring can carry 10 to 100 wavelength channels at data rates of 2.5 Gb/s (OC-48), 10 Gb/s (OC-192), and potentially higher; transparent optical connections, for example with analog-modulated signals, are also possible. Passive distribution networks can use various physical topologies, such as a tree, ring, or bus. Distribution networks carry both the feeder wavelength channels and the additional local distribution wavelength channels.




■ Figure 1. Physical architecture of the proposed WDM optical metropolitan access network.

## INTELLIGENT ACCESS NODES

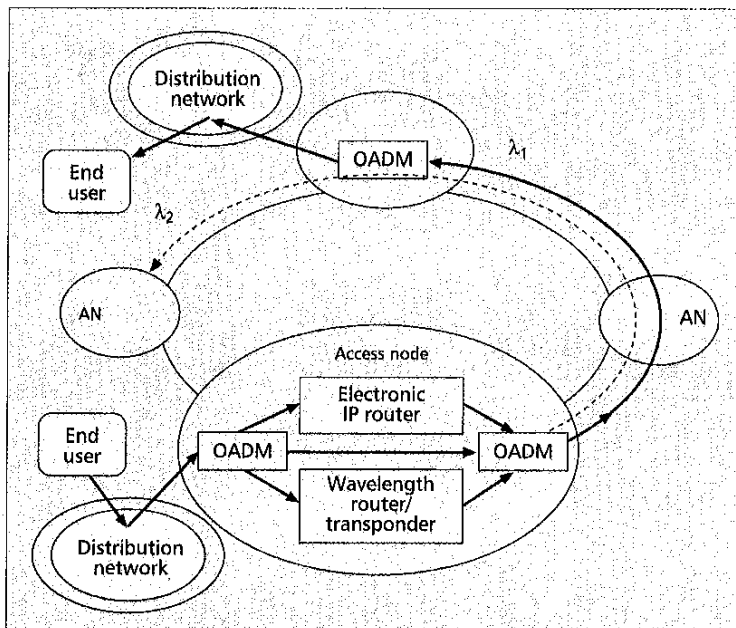
In this optical WDM network, both optical and electronic intelligent network functionality reside in the access nodes which interconnect data flows between the feeder and distribution networks. The role of the feeder access node is to route full optical wavelength channels and individual IP data packets inside wavelength channels toward their destinations; a functional diagram of the proposed intelligent access node is shown in Fig. 2. Two key features of our access node design are the electronic IP router and the optical wavelength router/transponder inside the node. A third important feature is the ability to direct optical signals to one of the two routers or to completely bypass the routers for transparent optical connections between users.

In the proposed architecture the wavelength channels selected for feeder add/drop can pass directly between the feeder and the distribution network or can be first processed in the electronic IP router or the wavelength router/transponder, which are part of the access node. The electronic IP router receives the selected wavelength channels at its input ports, converts the data from optical to electronic form, processes the IP packets in electronic form, and finally routes the packets by retransmitting them through its wavelength-tunable or fixed-wavelength optical output ports. Besides its routing function, the IP router can also aggregate the lower-data-rate traffic from the distribution network onto higher-speed feeder channels. The wavelength router/transponder routes an optical signal by receiving it on one wavelength channel and



■ **Figure 2.** A functional diagram of the network access node.

retransmitting it on another; control signaling is used to set the desired output port wavelength. The optical wavelength router/transponder allows end users with distribution-wavelength or fixed-wavelength transmitters to route their signals to the desired destinations by using transponder transmitters tunable over the feeder wavelengths. Wavelength transponder transmitters, through their use of long-reach external modulation lasers, can also possess the more demanding capability to transmit signals over longer distances of several hundred kilometers.



■ **Figure 3.** Network connection scenarios for electronic IP routing, wavelength routing, and optical flow switching.

Each fiber in our proposed feeder network is designed to carry multiple wavelength channels which we divide into WDM bands. This wavelength banding approach will be especially convenient and important when tens and hundreds of wavelength channels are carried over a single fiber. Grouping wavelengths into bands of 10 to 20 channels leaves a relatively small and manageable number of bands on the feeder, allowing the selection of channels for add/drop one band at a time. In this network architecture, the number of network users can be scaled conveniently by adding additional access nodes and wavelength bands on the feeder network.

An access node selects one or more bands to add/drop from the feeder to the distribution network; the nonselected wavelength bands bypass the access node transparently. In the present design the add/drop band selection is fixed and not switchable. The selected bands in the access node are demultiplexed into individual wavelength channels in the feeder optical add/drop multiplexer (OADM). We use a switched feeder OADM which can switch the individual channels to be added/dropped to the distribution network or to continue undisturbed on the feeder network. In the present design the distribution OADM is a passive, nonswitchable (de)multiplexer that separates all wavelength channels from the distribution network for processing in the access node.

## NETWORK CONNECTION SCENARIOS

User connection scenarios supported by this network are illustrated in Fig. 3. The end user first uses his distribution network connection to reach the entry access node; a passive distribution network is shared by many users by means of a WDM medium access control protocol. At the entry access node the data is routed toward its destination using the appropriate feeder-add wavelength channel. The channel is dropped at the egress access node and directed via the distribution network to the endpoint destination. Several connection scenarios are supported, with an appropriate scenario selected using control signaling. For IP traffic the user connection can be electronically terminated at the access node electronic IP router, which then routes the IP packets toward their destination. For large data flows, the user may bypass the IP router to send an optical data flow directly to its destination. Alternatively, the user may use the optical wavelength router/transponder to access different wavelength channels and data paths. We expect that a significant fraction of the data traffic can be handled using optical routing and bypassing the electronic IP routers. Therefore, for a given IP router capacity, our intelligent WDM network can handle a much larger traffic volume than a network with purely electronic routing.

Data path routing in this WDM network is accomplished by switching the access nodes' feeder OADMs to add/drop the desired feeder wavelength channels and by tuning the transmitters to the appropriate wavelength channels. Such wavelength-tunable transmitters are used in the end-user equipment, the electronic IP router output ports, and the wavelength router/transponder output ports. In principle, fast electronic wavelength

tuning of transmitters combined with a fast control channel could allow IP packet routing on a packet-by-packet basis. OADM switching is slower and also involves reconfiguration of the logical network topology; therefore, OADM switching is expected to be done infrequently in response to changes in traffic load patterns or for very large data file transfers.

### DUAL-FIBER RING AND OPTICAL PROTECTION SWITCHING

Figure 4 illustrates the dual-fiber ring architecture of the feeder network. Each fiber is unidirectional; one carries traffic clockwise (CW) while the other carries traffic counter-clockwise (CCW). Thus, duplex connections on the network consist of two simplex connections. In our initial implementation we plan two wavelength bands on the ring, the A band on the CW fiber and the B band on the CCW fiber. In case of a fiber cut, which can be detected by a loss of light in the fiber, traffic in the cut fiber is redirected onto the second fiber going in the opposite direction around the ring. In this protection mode, fibers carry both their normal band and the protected band from the opposite fiber. After going around the ring in the protection band, the signal reaches its intended node on the other side of the fiber cut, where the signal is redirected back onto its normal band fiber. In protection mode the signals are dropped at access nodes only when they pass them in their normal band and not in the protection band; this significantly simplifies the protection control scheme and the access node structure.

### ACCESS NODE IMPLEMENTATION

For the access node functional schematic in Fig. 2, one possible physical implementation of the dual-fiber ring feeder access node is shown in Fig. 5. Here we show only the electronic IP router, but a similar scheme can be used with both IP and wavelength routers inside the access node. WDM multichannel signals from dropped feeder bands of both CW and CCW fibers are first demultiplexed into individual wavelength channels. The optical 2 x 2 switch for each channel then selects whether the channel is added/dropped from the stream or continues undisturbed back onto the feeder. If a particular wavelength channel is dropped from the feeder band onto the distribution network, this channel can also be used to add a signal from the distribution network onto the feeder. Wavelength channels destined for the feeder are then remultiplexed, with variable attenuators first equilibrating the channel power levels.

On the feeder drop side, the dropped feeder wavelength channels can go directly to the distribution network or be terminated in the IP router, with a set of 1 x N optical switches playing the role of a tunable channel selector. The IP router drop output ports connect to the distribution network using the distribution wavelengths. On the feeder add side, signals from the distribution network can be added directly to the




Figure 4. Feeder network architecture: dual-fiber ring with optical protection switching.

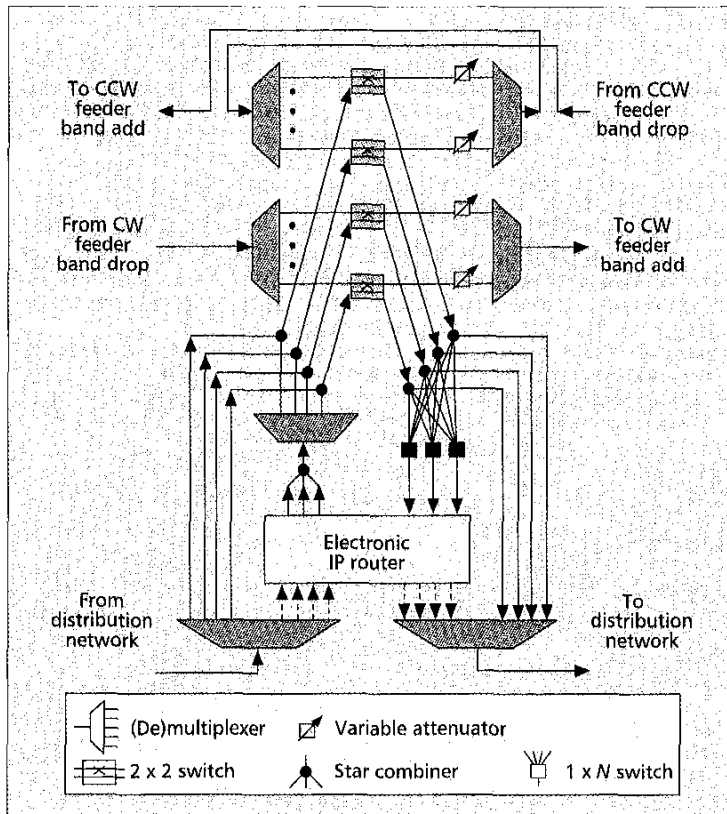


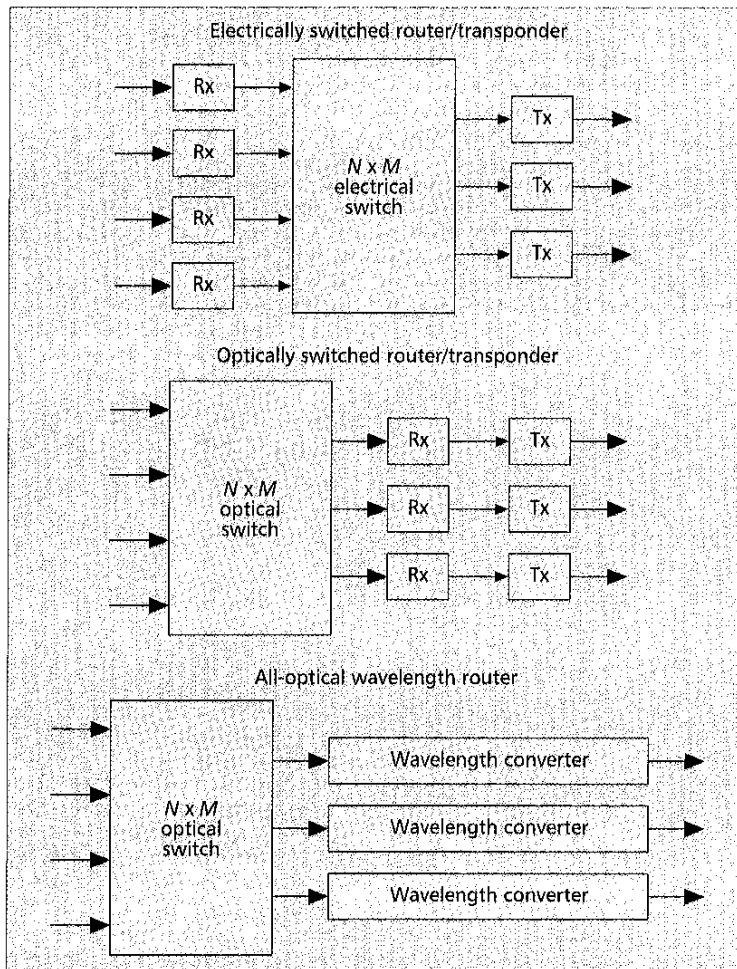
Figure 5. Physical implementation of the dual-fiber ring feeder access node.

feeder if they are on the feeder wavelengths. Signals on the distribution wavelengths, shown in dashed lines in Fig. 5, are electronically terminated in the IP router, where they are electronically processed and routed through the feeder using the routers output ports tunable over the feeder wavelengths. Not shown in the figure are optical

## OPTICAL WDM ENABLING TECHNOLOGIES

amplifiers which are used to compensate for propagation loss between access nodes and the various splitting and excess losses inside the access node.

Several schemes for implementing the access node wavelength router/transponder are shown in Fig. 6. In all cases, the use of  $N \times M$  signal switches allows many input signals to time share a smaller number of output ports, while wavelength routing is done by tuning the output ports' wavelengths. Control signaling is used for signal switching and output port tuning in these wavelength routers. In the two router/transponder configurations, the signals undergo optical-electrical-optical conversion with either electrical or optical switching of the signals; output transmitters are wavelength-tunable. It is easier to make optical-electrical-optical transpondering for specific bit rate signals, but bit-rate-independent transponders are also possible. In the all-optical wavelength router, the signals are switched optically; bit-rate-independent tunable all-optical wavelength converters are used in the output ports. All-optical wavelength converters are research type devices at present.



■ **Figure 6.** Tunable wavelength router/transponder implementations: electrically switched transponder, optically switched transponder, and optically switched all-optical wavelength converter. Rx: optoelectronic receiver; Tx: optoelectronic transmitter.

The optical WDM revolution in transmission and emerging applications in optical regional networking are enabled by a range of key optical technologies which are also used in our optical network design. At the foundation is the low loss, 0.25 dB/km, single-mode optical fiber, which allows long-distance transmission with a bandwidth window of about 25 THz. Erbium-doped fiber amplifiers (EDFAs) provide optical amplification to compensate power loss in optical signal transmission and processing (splitting, multiplexing etc.). Conventional C-band EDFAs cover the 1530–1565 nm wavelength range and extended L-band EDFAs cover the 1565–1605 nm range for a total available gain bandwidth of 9 THz. Commercial systems have already capitalized on this bandwidth by using 100 GHz and 50 GHz spaced channels at channel bit rates of up to 10 Gb/s and a total fiber capacity of greater than 1 Tb/s. Future trends are toward smaller channel spacing, higher bit rate per channel, and higher spectral efficiency (bits per second per hertz).

WDM filters are another key technology; they allow splitting and combining of the available wavelength band into more than 100 individual wavelength channels. A variety of WDM filter types are available, such as fiber Bragg grating filters, dielectric thin-film filters, and Mach-Zehnder and waveguide-grating-router devices. Semiconductor lasers are used in transmitters for the WDM multichannel systems. Fixed-wavelength distributed feedback (DFB) semiconductor lasers with wavelengths on an internationally standardized wavelength grid are now widely available. Semiconductor lasers tunable over as much as 40 nm (4.6 THz) are also becoming commercially available.

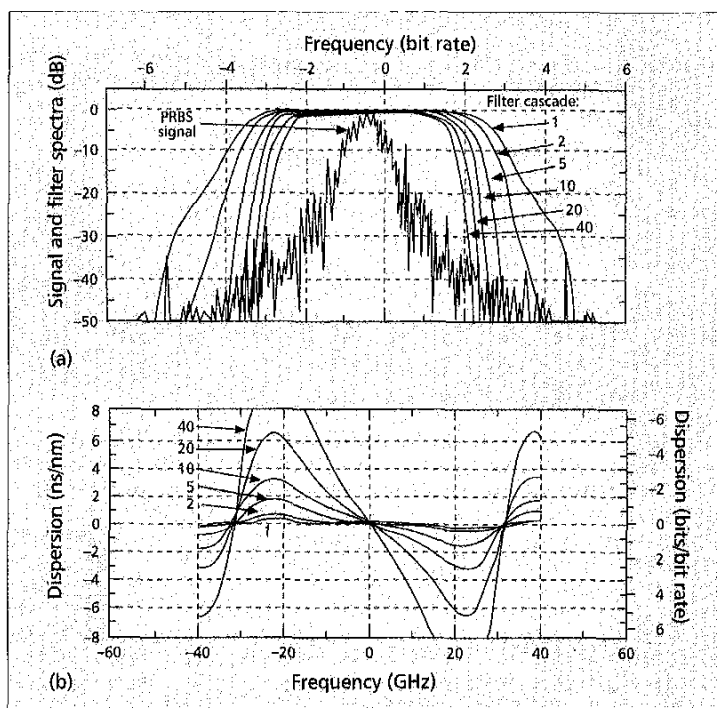
Electro-mechanical, and in the very near future micro-electromechanical system (MEMS), optical fiber switches switch full-band optical signals between multiple fibers; switching times are now typically between a few and tens of milliseconds. Signal splitters, combiners, and variable attenuators allow further manipulation of optical signals. High-speed, greater than 10 Gb/s, single-channel transmission is enabled by a variety of optical and electronic components. Among the key devices here are low-wavelength-chirp external lithium-niobate modulators, as well as integrated or copackaged semiconductor electroabsorption and Mach-Zehnder modulators, with modulation bandwidths in excess of 10 GHz. Commercial PIN photodetectors and avalanche photodetectors (APDs) provide receiver bit rates of 10 Gb/s and higher; EDFA-based optical preamplifiers enable high sensitivity of these optical receivers. Dispersion-compensating fiber allows compensation of the fiber link dispersion for high-speed long-distance signal propagation, as well as tailoring of the dispersion profile in the dispersion-managed optical signal path. These and other optical technologies are being further developed to provide richer functionality for WDM applications while improving their performance and reducing component-induced signal impairments.

## WDM NETWORK OPTICAL SIGNAL IMPAIRMENTS

Optical signals propagating over WDM networks suffer a variety of physical impairments that degrade the signals and limit system performance. Understanding these impairments and minimizing signal degradation are especially important for transparent optical connections across the networks. Some of the key impairments are dispersion of optical fibers and WDM multiplexing/demultiplexing filters, coherent and incoherent signal crosstalk from imperfect WDM channel filtering, WDM filter passband narrowing on filter cascading, polarization-dependent loss and polarization mode dispersion, finite optical signal extinction, wavelength chirp of modulated laser transmitters, and optical fiber nonlinearities. Optical amplifier impairments include additive spontaneous emission noise and nonflat gain profile. An important issue for reconfigurable WDM networks is transient gain dynamics in optical amplifiers due to WDM channel switching; these transients can be suppressed by electronic or all-optical gain control techniques. Optical amplifier impairments, inter-channel crosstalk, and signal dispersion are probably the most important signal degradation mechanisms in WDM networks. In the network context, where optical signals pass through multiple access nodes and potentially through an arbitrary path in a mesh-connected network, it is especially important to understand the cascading behavior of these physical impairments.

As an illustration of a physical impairment and its system implications, we show results of our calculations of receiver power penalty for optical signals propagating through a cascade of dispersive WDM filters. As shown in Fig. 5, WDM filters are encountered in network access nodes as channel demultiplexers and multiplexers; an optical signal traveling through a series of access nodes will traverse a cascade of such WDM filters. Signal distortion due to dispersive effects of the cascade results in intersymbol interference and introduces bit errors; this requires a compensating increase in received optical power, which is defined as the receiver power penalty, in order to restore the system-required bit error ratio.

Figure 7 shows the calculated magnitude and dispersion characteristics of cascaded fiber Bragg grating filters, a typical type of WDM filter. The power spectrum of a 10 Gb/s pseudo-random bit-stream (PRBS) is also shown. For these flat-top filters, the effective passband narrows only slightly on cascading as many as 40 filters. The dispersion varies linearly with optical frequency near the passband center. While the dispersion slope of the cascade increases linearly with the number of cascaded filters, the slope also varies as the inverse cube of the Bragg filter bandwidth. In Fig. 8 we show a contour plot of the cascade power penalty for optically preamplified receivers as a function of the accumulated dispersion slope and signal detuning from the cascade zero-dispersion point. The low-penalty operating region, defined as penalty of less than 6 dB, allows wide detuning for small dispersion slopes. This region narrows rapidly and eventually closes with increasing dis-

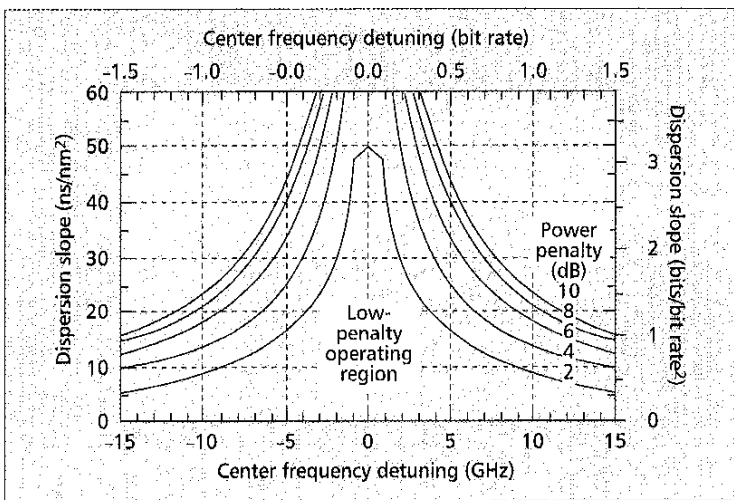


■ **Figure 7.** An example of the WDM network optical impairment: dispersion of cascaded WDM channel filters: a) cascade passband magnitude and 10 Gb/s PRBS spectrum; b) cascade dispersion.

persion slope. For 10 Gb/s signals and a tolerable penalty of 6 dB, we can tolerate an accumulated dispersion slope of 34 ns/nm<sup>2</sup> with allowed signal-filter detuning of  $\pm 5$  GHz, corresponding to 26 cascaded 60 GHz bandwidth fiber Bragg grating filters. The allowed signal-filter detuning implies a requirement on the laser optical frequency stability and the center frequency accuracy of the cascaded WDM filters. In a system trade-off, one can use fiber Bragg grating filters with a wider bandwidth and a smaller dispersion slope; the price is larger neighboring channel crosstalk of the wider filters. This impairment example illustrates how the dispersion slope of fiber Bragg grating filters can limit the number of cascaded WDM filters, and consequently the number of signal-traversed network access nodes.

## CONCLUSIONS

High-speed and high-capacity optical fiber communication technology already available on the backbone is beginning to penetrate into regional access networks. We describe a physical architecture of an optical wavelength-division multiplexed metropolitan access network which synergistically combines electronic IP routing with intelligent networking functionality of the WDM optical layer. Among the major networking functions of the optical layer are wavelength routing of signals, network logical topology reconfiguration, direct optical connections to offload traffic and bypass IP routers, and optical protection switching. Using its electronic and optical networking functions, the network offers the choice of IP routed service, wavelength routed optical service, and transparent



■ **Figure 8.** Receiver power penalty due to dispersion slope of the WDM channel filter cascade limits the number of cascaded filters and the allowed signal-filter detuning in WDM networks. Bottom and left scales are for 10 Gb/s signals, right and top scales are bit-rate-independent.

optical flow connections. The use of the WDM layer to perform networking functions also allows support of heterogeneous communication traffic and significantly expands network capacity and functionality beyond that of a purely electronically switched network. With 100 or more wavelength channels at a 10 Gb/s channel rate, the network capacity can be greater than 1 Tb/s, accommodating 1000 end users with user access rates of 1–10 Gb/s and higher. We also outline the key enabling optical technologies for such WDM networks and illustrate an optical signal impairment in the system. The WDM metropolitan access network technology outlined here will enable WDM networking revolution for gigabit-per-second user access in next-generation Internet networks.

#### ACKNOWLEDGMENTS

The authors acknowledge support for this work under the DARPA Next Generation Internet program. The authors also acknowledge Adel Saleh, Jane Simmons, R. Ian MacDonald, and Bruce Schofield for their contributions to this work.

#### BIBLIOGRAPHY

- [1] V. W. S. Chan, K. L. Hall, E. Modiano, and K. A. Rauschenbach, "Architectures and Technologies for High-Speed Optical Data Networks," *J. Lightwave Tech.*, vol. 16, 1998, pp. 2146–68.
- [2] E. Modiano, "WDM-Based Packet Networks," *IEEE Commun. Mag.*, Mar. 1999, pp. 130–35.
- [3] I. P. Kaminow et al., "A Wideband All-Optical WDM Network," *IEEE JSAC*, vol. 14, 1996, pp. 780–99.
- [4] R. E. Wagner et al., "MONET: Multiwavelength Optical Networking," *J. Lightwave Tech.*, vol. 14, 1996, pp. 1349–55.
- [5] "Special Issue on High-Capacity Optical Transport Networks," *IEEE JSAC*, vol. 16, Sept. 1998.
- [6] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, 1998.

#### BIOGRAPHIES

MARK KUZNETSOV (mkuznet@ll.mit.edu) received his B.S. degree in electrical engineering from the Ohio State University in 1979 and his S.M., E.E., and Sc.D. degrees in

electrical engineering from the Massachusetts Institute of Technology (MIT) in 1982 and 1986. From 1986 to 1993 he was with AT&T Bell Laboratories working on tunable semiconductor lasers and other devices for wavelength-division multiplexed optical networks. From 1993 to 1997 he was with Microron developing high-power semiconductor laser technology. In 1998 he joined MIT Lincoln Laboratory, where he is working on high-speed optical communication networks.

NAN MOORE FROBERG received a B.S.E. in mechanical engineering from Princeton University in 1985, and M.S. and Ph.D. degrees in electrical engineering from Columbia University in 1989 and 1992, respectively. After working at Lucent Technologies' Bell Laboratories, she joined MIT Lincoln Laboratory as a member of technical staff in 1998.

SCOTT HENION received his B.S. degree in physics from the State University of New York at Brockport in 1983 and an M.S. degree in physics from the University of Lowell in 1988. Since 1988 he has been working at MIT Lincoln Laboratory in the areas of solid-state lasers, RF photonics, and wavelength-division multiplexed optical networks.

HEMANTH RAO received a B.A. in physics from Cornell University in 1992, and M.S. and Ph.D. degrees in electrical engineering from the University of California, San Diego, in 1994 and 1997, respectively. He joined MIT Lincoln Laboratory in 1997.

JEFF KORN received his B.S. and M.S. degrees in electrical engineering from Polytechnic University of New York in 1988 and an M.A. degree in physics from the University of Rochester in 1994. In 1989 he joined the Quantum Electronics group at MIT Lincoln Laboratory. From 1996 to 1999 he was with the Advanced Networks Group of MIT Lincoln Laboratory conducting research on nonlinear pulse propagation, TDM and WDM optical networks, and optical amplifier technology. Presently, he is with AXSUN Technologies developing subsystems for optical networks.

KRISTIN RAUSCHENBACH received a B.S. degree in electrical engineering from the University of Minnesota, Minneapolis, in 1984, and M.S., E.E., and Ph.D. degrees in electrical engineering from MIT, Cambridge, in 1986, 1987, and 1989, respectively. In 1989, she joined Arthur D. Little, Inc., in Cambridge, where she worked on a variety of technical product development and business consulting projects. From 1991 to 1999 she worked at MIT Lincoln Laboratory, Lexington.

EYTAN MODIANO received his B.S. degree in electrical engineering and computer science from the University of Connecticut at Storrs in 1986, and his M.S. and Ph.D. degrees, both in electrical engineering, from the University of Maryland, College Park, in 1989 and 1992, respectively. He was a Naval Research Laboratory Fellow, 1987–1992, and a National Research Council Post Doctoral Fellow during 1992–1993 while he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he worked on communication protocols for satellite, wireless, and optical networks and was project leader for MIT Lincoln Laboratory's NGI project. In 1999 he joined the Aeronautics and Astronautics Department at MIT.

VINCENT W. S. CHAN [F] received his B.S., M.S., E.E., and Ph.D. degrees in electrical engineering from MIT in 1971, 1972, and 1974, respectively, in the area of communications. From 1974 to 1977 he was an assistant professor with the School of Electrical Engineering at Cornell University. He joined Lincoln Laboratory in 1977 as a staff member of the Satellite Communication System Engineering Group. In January 1981, he became assistant leader of the Communication Technology Group. In July 1983 he formed and became leader of the Optical Communication Technology Group and manager of the Laser Intersatellite Transmission Experiment Program. He was formerly head of the Communications and Information Technology Division and served as principal investigator of a new NGI consortium formed among AT&T, Cabletron, MIT, and JDS Fitel. On 1 January 1999, he became director of the Laboratory for Information and Decision Systems on the MIT campus and serves on the faculty as the Joan and Irwin Jacob Professor of Electrical Engineering and Computer Science and Aeronautics and Astronautics.