SpectraLAN: A Low-Cost Multiwavelength Local Area Network

Brian E. Lemoff

Lewis B. Aronson

Lisa A. Buckman

The first-generation SpectraLAN module will allow existing 62.5-μm multimode fiber-optic links to carry four times higher data rates than is possible with conventional methods. Four-channel error-free operation at aggregate data rates of 2.5 and 4.0 Gbits/s has been demonstrated over distances of 500 m and 300 m, respectively. The module is compact and potentially low-cost.

With desktop computers requiring ever more communications bandwidth, fiber optics has become an essential part of the local area network (LAN). While electrical cables still run to the desktop, fiber-optic links now form the network backbone in a rapidly increasing number of office buildings. Hewlett-Packard is currently the leading supplier of optical components for these fiber-optic LANs. The fastest components for LAN backbones currently on the market send data at a rate of 1 billion bits per second (written 1 Gbit/s).[1] As computers become faster and applications become more sophisticated, even higher data rates will be required. However, the type of fiber currently installed in most buildings may not be capable, using conventional methods, of carrying data faster than 1 Gbit/s over the required distances.

HP Laboratories has begun a project called *SpectraLAN* to use wavelength-division multiplexing (WDM) to enable the currently installed fiber to support data rates many times higher than the conventional limit. In a WDM system, light from several lasers of different wavelengths is combined into a single fiber. Each wavelength carries an independent signal, which can be as fast as the conventional data rate limit for the fiber. At the receiving end of the fiber, the different wavelengths are separated and detected separately. In this way, the total capacity of the fiber can be increased by a factor equal to the number of wavelength channels.

WDM has already become quite important in long-distance telecommunications applications (***Article 3***). However, size and cost are much more critical in LAN applications than in long-distance telecommunications. The goal of the SpectraLAN project is to develop a compact, low-cost WDM transmitter/receiver module with comparable size and cost to current HP transceiver modules, but with several times higher bandwidth.

Need for WDM in the LAN

To understand the details of the SpectraLAN project, it is necessary to understand the LAN issues that motivate it. The term LAN can be used to describe many different types and sizes of networks, from a small office network connecting three or four computers to a server and a printer, to a campus-wide network connecting dozens of buildings. Most LANs have several types of data links, as shown in **Figure 1**. Workgroup links connect desktop computers, printers, and servers to one another and to the building backbone. Building backbone links connect different workgroups within a building to one another, and campus backbone links run between nearby buildings. A telecommunications link will typically connect the LAN to the rest of the world.

Figure 1

Diagram of a typical LAN, showing typical implementations for workgroup, building backbone, and campus backbone links.

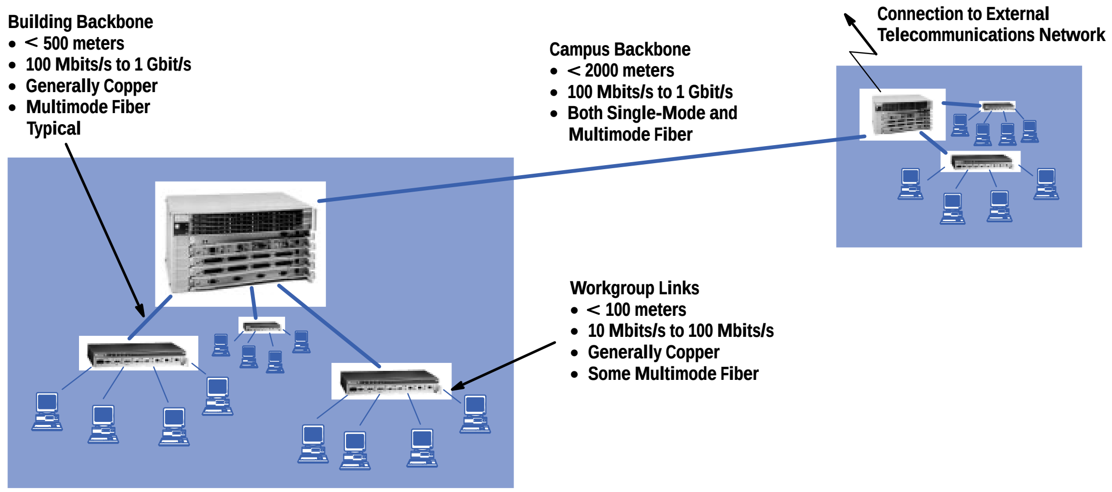

Physically, the various link types can be quite different. Workgroup links are generally shorter than 100 m in length and have data rates of 10 Mbits/s to 100 Mbits/s. Copper wire, usually in the form of twisted-pair, is well-suited for this, although optical fiber is occasionally used. Wireless communication may also become useful in the workgroup environment. Because of the short distance and relatively low bandwidth required in the workgroup, existing technologies will likely be sufficient for many years. Another area in which existing technology will likely suffice for some time is the telecommunications link connecting the LAN to the outside world. Such links typically operate at lower speeds than the LAN backbone and are carried on single-mode optical fiber, which has very high bandwidth capacity.

While single-mode fiber is widespread in telecommunications, it is somewhat less common in campus backbone links and almost nonexistent in building backbones. The type of fiber used in building backbones in the U.S.A. and Europe is almost exclusively 62.5-µm-core multimode fiber. Because of its larger core diameter, components and connectors used in multimode fiber links can have looser alignment tolerances, and hence lower cost than those used in single-mode fiber links. This larger core also imposes a limit—a result of modal dispersion, which causes portions of the signal to travel at different velocities—on the product of bandwidth and distance over multimode fiber. At a wavelength of 850 nm, where inexpensive lasers are available, the bandwidth-distance limit is approximately 300 Mbits/s · km. Thus, a building backbone link 500 m long can carry a maximum data rate of approximately 600 Mbits/s, and a 2-km-long multimode fiber campus backbone link can carry a maximum data rate of approximately 150 Mbits/s. High-end LAN products are already pushing these limits in building and campus backbones.

Figure 2 shows a plot of distance and bandwidth for existing LAN communications standards compared with the 850-nm dispersion limit of 62.5-µm multimode fiber. As workgroups demand even higher data rates, the backbones that connect them will become bottlenecks unless a solution is found to push well beyond the dispersion limit. This issue has already been raised in the ongoing standards development work for gigabit Ethernet (IEEE 802.3z). It has been recognized that 850-nm components will not be capable of meeting the 500-m length for building backbones.[2]

Figure 2

Distance versus data rate for some fiber-optic LAN standards and some possible WDM systems. Also shown are estimated LAN requirements and the 850-nm bandwidth-distance limit for 62.5-µm multimode fiber.

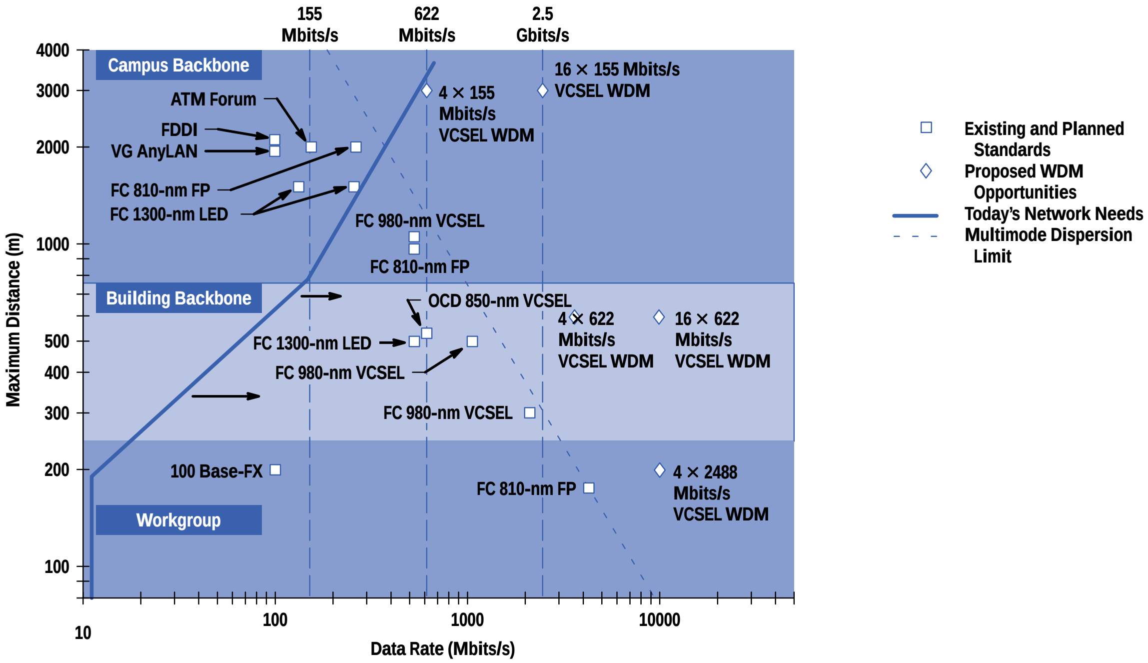

There are several possible solutions to the impending backbone bandwidth crisis. One solution is simply to replace all of the multimode fiber currently installed with single-mode fiber or with a new type of multimode fiber that allows higher bandwidth. While this would clearly work, it would be a costly and time-consuming endeavor. Most customers would prefer to avoid this procedure if at all possible. A second solution would be to use 1300-nm-wavelength lasers rather than shorter-wavelength lasers. The bandwidth-distance limit of 62.5-µm multimode fiber is three times higher at this wavelength than at 850 nm, making this solution viable as a short-term fix.

To reach the 500-m length at 1 Gbit/s, the gigabit Ethernet standard proposes to use 1300-nm components originally designed for single-mode applications. There are some disadvantages, however. Current 1300-nm lasers are much more expensive than their shorter-wavelength cousins, and when used with multimode fiber their performance can be limited by a phenomenon known as modal noise. The development of vertical-cavity surface emitting lasers (VCSELs) at

1300-nm wavelength would overcome these disadvantages. However, this research may still be several years from producing a commercially viable product.

There have also been more novel approaches suggested to squeeze more bandwidth from existing multimode fiber. In one scheme, data would be transmitted with multiple logic levels, rather than the traditional binary logic. Each bit of data could have four or eight levels rather than two. In this way, more information could be transmitted in the same available bandwidth. This method requires much better signal-to-noise performance than the traditional on/off approach. Another technique that has been investigated is one in which the spatial properties of the laser beam are controlled as the light is launched into the fiber. If a spot much smaller than the 62.5-μm core of the fiber is aligned to a particular position on the face of the fiber, a factor of two or three higher bandwidth can be obtained. This technique is very sensitive to the quality of the fiber and to the alignment of the connectors in the link.[3]

Wavelength-division multiplexing is a promising technique for extending the capacity of existing multimode fiber. Each wavelength channel in a WDM link can use existing low-cost lasers and standard 622-Mbit/s electronics to transmit data within the conventional limits of the fiber. By combining multiple wavelengths in the fiber, an immediate increase in total capacity is obtained. In addition, should one of the alternate techniques mentioned above prove useful, it could be combined with WDM to multiply the bandwidth even further. WDM also offers a great deal of flexibility for future upgrades. While many of the other methods only offer a factor of two to four increase in fiber bandwidth, WDM is limited only by the number of wavelengths that can be combined in the fiber and separated. While initially a four-channel system might be sufficient, there is no fundamental reason why 8, 16, or more wavelength channels would not be feasible in the future. As we will discuss in the next section, we are looking at ways to implement WDM in an inexpensive and compact manner, using much of the same technology already found in HP multimode optical components.

SpectraLAN Project

The first goal of the SpectraLAN project has been to develop a WDM transceiver module capable of sending and receiving four parallel 622-Mbit/s signals over a single 62.5-μm multimode fiber, with link lengths of up to 500 meters. The project emphasizes small size and low cost, essential qualities for LAN applications. Much of the SpectraLAN technology is similar to that used in the POLO project, described in the ***Article 6***.

A conceptual picture of the SpectraLAN transceiver module is shown in **Figure 3**. Initially, a ceramic multichip module is being used, upon which the electronics, lasers, and detectors are mounted. Four VCSELs, each of a different wavelength, are driven by laser drivers fabricated in HP-25, a 25-GHz f_T silicon bipolar IC process. The light from the VCSELs is combined in a polymer-waveguide 4-to-1 combiner, which is coupled to an outgoing fiber. The incoming light passes from the fiber through a wavelength demultiplexer. In the figure, this demultiplexer is a 1-to-4 polymer waveguide splitter with dielectric interference filters on the four outputs, each passing only one wavelength. The light is detected on an array of four photodiodes, and HP-25 receiver electronics convert the signal to a digital electronic output. A close-up photograph of an MSM (metal-semiconductor-metal) detector array as it is packaged in the SpectraLAN module is shown in **Figure 4**.

One of the key technologies in the SpectraLAN module is the vertical-cavity surface emitting laser. The VCSEL is the ideal source for multimode fiber transmission for several reasons. The relatively low-divergence, round laser beam output of a VCSEL can be coupled into a multimode fiber with very high efficiency without the use of a lens. Also, the relatively broad optical spectrum (1-to-4-nm linewidth) possible with multimode devices reduces undesirable coherent effects such as modal noise, a major problem with edge emitting lasers. Because VCSELs are surface emitting, they are small in area, they can be tested on-wafer, and dicing is a noncritical step, giving this technology the potential to be extremely low-cost. In SpectraLAN, four VCSELs are used, with wavelengths of 820 nm, 835 nm, 850 nm, and 865 nm. In the first-generation modules, the four devices are individually mounted dice, selected from different parts of a single nonuniform wafer, or taken from separate wafers, each of a different wavelength. **Figure 5** shows a set of four discrete VCSELs packaged in a multichip module. Eventually, we hope to develop monolithic multiwavelength arrays in which all four devices will be on a single die. This will greatly simplify alignment in the packaging process.

Figure 3

Conceptual drawing of a SpectraLAN module. Key components are indicated in the drawing.

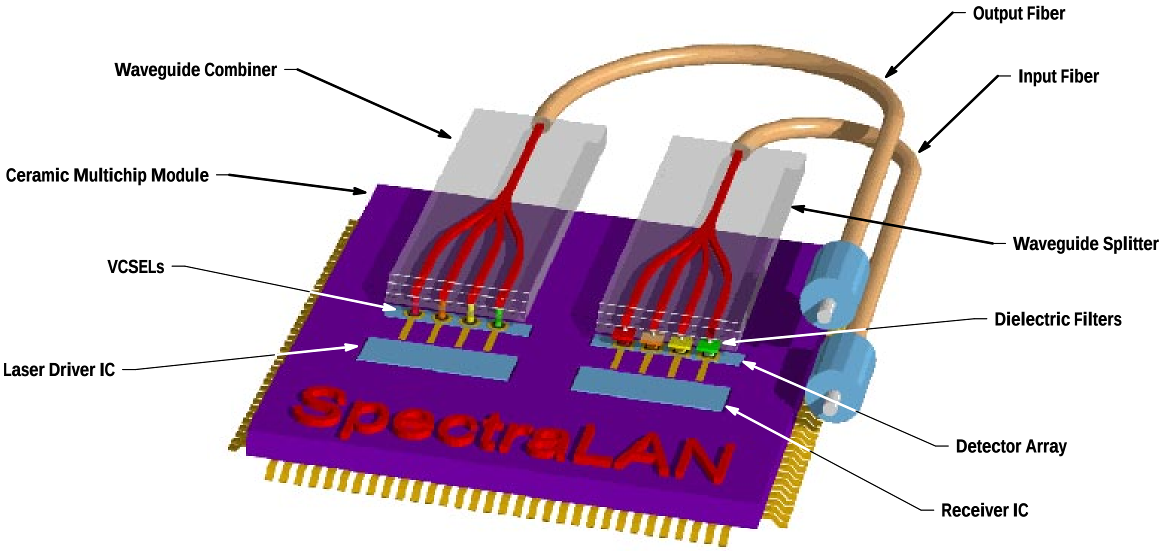

Figure 4

Photograph of a GaAs MSM detector array packaged in a ceramic module. Also shown is a Polyguide™ splitter demultiplexer, situated immediately above the detectors.

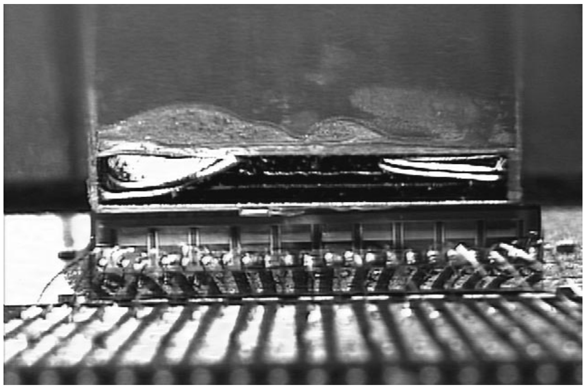

Figure 5

Photograph of a discrete multiple wavelength VCSEL array packaged in a ceramic module.

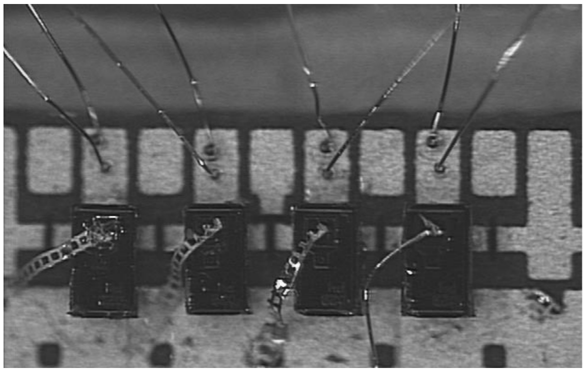

The four detectors in the SpectraLAN module are in a monolithic array. Both silicon and GaAs photodiodes are being evaluated. (The detectors shown in **Figure 4** are GaAs MSM detectors.) By using large-area detectors, excellent alignment tolerance can be achieved. The electronics for driving the VCSELs and amplifying the detected signals uses the HP-25 silicon bipolar technology. This is currently needed to meet speed and power requirements. Eventually, as newer, low-current VCSELs become available, some of the electronics may be replaced with high-speed CMOS or BiCMOS circuitry, resulting in even lower cost and power consumption.

One additional piece of electronics that may be required is a skew compensation circuit. Since the goal of the project is to transmit data in parallel, the four bits transmitted simultaneously must be received together at the other end of the link. Chromatic dispersion in the fiber results in longer wavelengths traveling at a slightly higher velocity than shorter wavelengths. Thus, over sufficient distance, different-wavelength bits may lag one another by more than one bit period. Since the dispersion of fiber is a known quantity, this effect can be compensated for once the length of the link is determined. The compensation circuit could determine the link length by looking at the delay between bits of two adjacent wavelength channels, and could then add appropriate delays to each channel to provide parallel output from the receiver.

Most of the technology described above is common to many high-speed fiber-optic communications applications. There are two key components in the SpectraLAN module that use technology unique to WDM. These are, not surprisingly, the wavelength multiplexer and demultiplexer. There are countless ways of combining and separating light of different wavelengths, and over the years, many of these have been applied to WDM systems.[4] For SpectraLAN, we must find the least expensive, most compact methods for wavelength multiplexing and demultiplexing, be they elegant or not. Polymer waveguide technology, similar to that used in the POLO project, can be useful for this application. Polymer waveguides are plastic structures in which channels of higher refractive index material are buried in a cladding of lower refractive index material. Light coupled into these channels is guided by total internal reflection. In most technologies, the waveguides are defined photolithographically, and have roughly square cross sections. The technology used in the POLO project, called Polyguide™, was developed by DuPont and licensed by HP. Polyguide™ material is supplied in rolls many yards long, and devices are fabricated in sheets, making this an inexpensive technology. Polymer waveguide technologies developed by Allied Signal and Mitsubishi also show promise for multimode applications.

Multiplexing is a somewhat simpler function than demultiplexing. In principle, if VCSELs could be made small enough and packaged close enough together, they could be coupled directly to the input face of the fiber, all four lasers emitting their light into the fiber core. When multiwavelength monolithic arrays become available, a scheme similar to this may be possible. For the present, however, we must be content with a center-to-center VCSEL spacing of 500 μm.

In SpectraLAN, we use a polymer waveguide device (shown in **Figure 3**) in which four input waveguides merge into one output waveguide. The light emitted from the lasers is coupled into the waveguides through reflection off a 45° mirror cut into the edge of the waveguides. By choosing the width of the input waveguides to be smaller than the width of the output waveguide, loss in the combining process can be minimized. The output waveguide is then coupled to the fiber either through permanent attachment (pigtailing) or through a detachable connector. All of the devices built to date have been pigtailed. We have observed total insertion losses, from VCSEL to fiber, of 2.1 dB in a Polyguide™ combiner. A 4-to-1 combiner using single-mode waveguides would have a minimum loss of 6 dB. It is the multimode nature of our combiner that makes much lower loss possible. The insertion loss of the combiner is very sensitive to the numerical aperture (divergence) of the VCSEL. VCSELs with high divergence see larger losses, while VCSELs with more collimated output see lower loss. Other combiner designs, using bulk or diffractive optics, may be investigated in the future.

The wavelength demultiplexer must separate the four wavelengths, delivering each wavelength to a different detector. In addition to minimizing insertion loss, cross talk between different channels must be avoided. We have been investigating three techniques for demultiplexing. The first and simplest technique, shown in **Figure 6**, is to use a polymer waveguide 1-to-4 splitter, similar to the polymer combiner, to divide each wavelength equally among the four outputs. A tiny dielectric interference filter is then attached at each output. The filters are made by depositing a multilayer dielectric stack onto a fused quartz substrate using a PECVD (plasma enhanced chemical vapor deposition) process. Filters are lapped and diced into 200-μm squares, approximately 50 μm thick. Each filter passes only one of the four wavelengths. The advantage of such a device is its simplicity and ease of fabrication. Splitting the light before filtering, however, introduces a fundamental 6-dB loss for a four-channel system.

Figure 7 shows transmission as a function of wavelength for the four channels of a splitter demultiplexer that we have fabricated using Polyguide™. The measured insertion loss is between 8.4 dB and 8.9 dB per channel. Each channel has a useful range of >7.8 nm about the central channel wavelength. This allows for variations in VCSEL wavelength resulting

Figure 6

(a) Drawing of a polymer waveguide splitter demultiplexer. (b) Photograph of a splitter demultiplexer made using Polyguide™ technology.

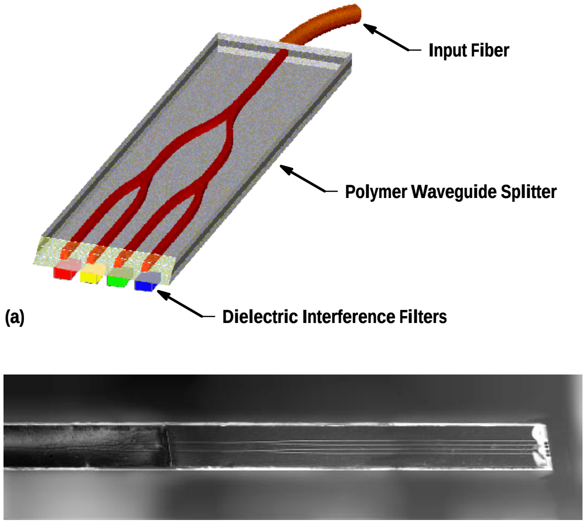

Figure 7

Measured transmission versus wavelength for the four outputs of the splitter demultiplexer shown in Fig. 7. Shaded regions indicate the useful wavelength range of each channel, corresponding to < – 15 dB relative cross talk and < 9.5 dB insertion loss.

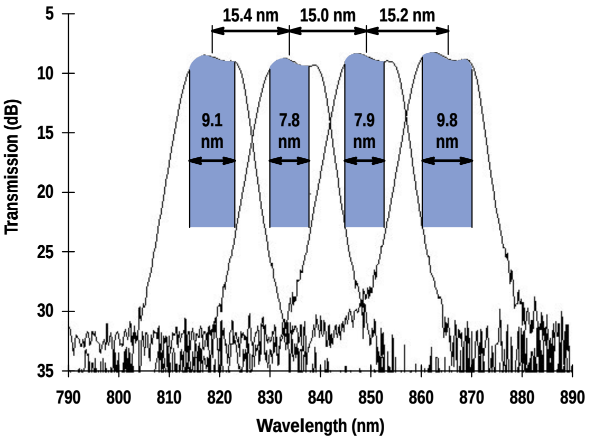

from nonuniform growth and drift in operating temperature. The large insertion loss of the splitter demultiplexer may still be small enough to allow a four-channel WDM link to function adequately. However, the loss of such a device scales directly with the number of channels, making it less likely that it can be successfully used in an 8-channel or 16-channel system.

A second type of wavelength demultiplexer, which should have a much lower insertion loss, is the zigzag demultiplexer, shown conceptually in **Figure 8**. This device uses the same inexpensive polymer waveguide and dielectric filter technologies found in the splitter demultiplexer. The geometry, however, is quite different. In the zigzag demultiplexer, the input

Figure 8

Conceptual drawing of zigzag demultiplexer.

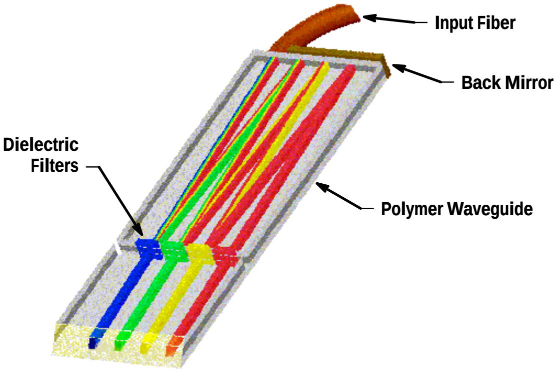

light impinges upon the first dielectric filter at an angle. One wavelength is transmitted through the filter, while the remaining wavelengths are reflected. The dielectric interference filters act as highly reflecting mirrors (99%) outside of their passband. A zigzag waveguide pattern guides the light to successive filters until each wavelength has exited the device through the appropriate output, where it can be detected. Unlike the splitter device, which throws away three quarters of the light from each channel, the zigzag demultiplexer uses all of the light, giving the device no geometrical loss factor. The insertion loss of the zigzag device should be dominated by the transmission loss of the filters and by bulk losses in the waveguides. Our initial experiments with zigzag demultiplexers have yielded encouraging results.

A third design that we have considered for wavelength separation, shown in **Figure 9**, uses a lens and a diffraction grating.[5] The lens collimates the light coming out of the fiber. This light diffracts off the grating, with different wavelengths diffracting at different angles. The light is then refocused by the same lens either directly onto the detectors or into a polymer waveguide array that guides the light to the detectors. **Figure 10** shows transmission as a function of wavelength for a bulk optical prototype of this device, using a 10-nm channel spacing. The advantages of such a design include relatively low loss, extremely low cross talk between channels, and excellent scaling to larger numbers of channels. The challenge is to make the grating and lens small enough and cheap enough to be economical for the LAN application. Recent advances in precision plastic molding may hold promise for accomplishing this goal.

Figure 9

Conceptual drawing of grating-based demultiplexer.

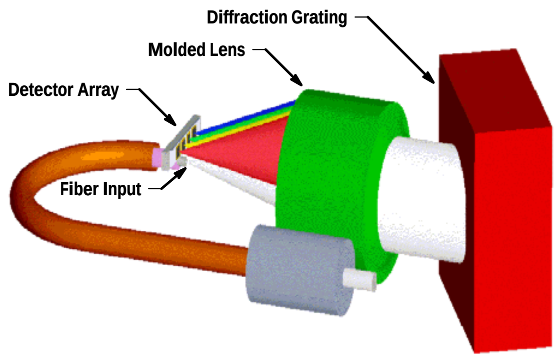

Figure 10

Measured transmission versus wavelength for the four outputs of a bulk grating-based demultiplexer. Shaded regions indicate the useful wavelength range of each channel, corresponding to < -15 dB relative cross talk and < 4.5 dB insertion loss.

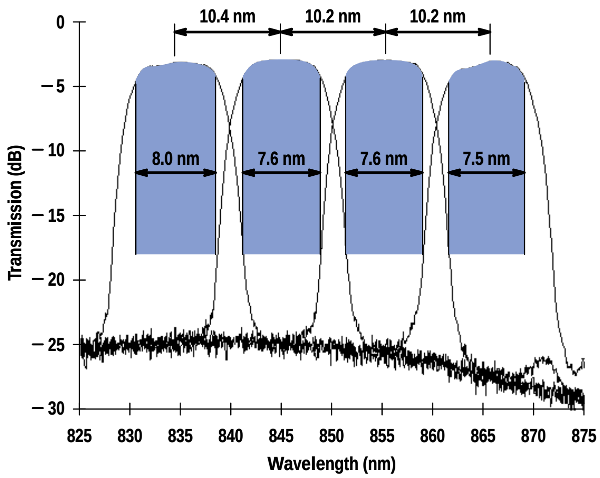

System Results

System measurements have been performed using two simplex SpectraLAN modules. Each of these modules fulfills half of the functions of the duplex module shown in **Figure 3**. The transmitter module has four VCSELs, a Polyguide™ multiplexer, and laser driver electronics as shown in **Figure 5**. The receiver module has a GaAs detector array, a Polyguide™ splitter demultiplexer, and receiver electronics as shown in **Figure 4**. **Figure 11** shows the measured spectrum of the light in the 62.5-µm multimode fiber. Superimposed is the splitter filter function shown in **Figure 7**. **Table I** shows the measured insertion loss and channel-to-channel cross talk of the system. The measurement setup is shown schematically in **Figure 12**. A bit error rate tester is used to drive the 835-nm channel of the transmitter module with a $2^{23}-1$ pseudorandom bit sequence. The other three channels are simultaneously driven by a parallel data generator that generates

Figure 11

Measured spectrum of wavelength-division multiplexed light using a simplex SpectraLAN transmitter module. Superimposed is the filter function of the splitter demultiplexer shown in Fig. 7.

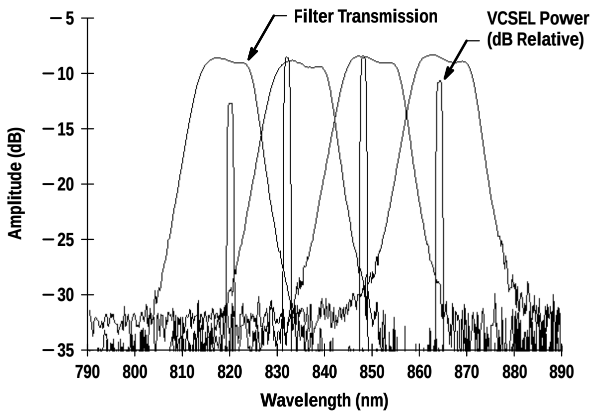

Figure 12

Experimental arrangement for four-channel WDM system measurements.

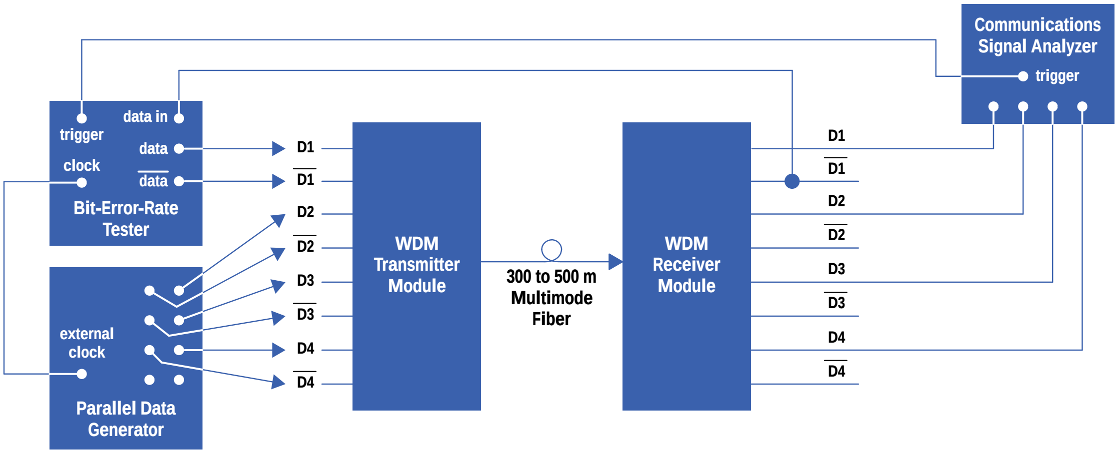

Table Ia

Absolute Insertion Loss of Four-Channel Splitter Demultiplexer (dB)

		Output Channel Wavelength (nm)			
		820	835	850	865
Input Channel Wavelength (nm)	820	8.83	26.9	38.3	31.8
	835	30.2	8.9	32.5	32.2
	850	41.0	29.3	8.4	29.0
	865	40.2	33.5	31.2	8.4

Table Ib

Relative Cross Talk of Four-Channel Splitter Demultiplexer (dB)

		Output Channel Wavelength (nm)			
		820	835	850	865
Input Channel Wavelength (nm)	820	0	−18.0	−29.9	−23.3
	835	−21.4	0	−24.1	−23.8
	850	−32.2	−20.4	0	−20.6
	865	−31.4	−24.6	−22.8	0

three independent $2^7 - 1$ pseudorandom bit sequences, which are synchronized to the bit error rate tester signal. A communications signal analyzer displays the four waveforms output from the receiver module. With high persistence, waveforms from many sweeps are superimposed, forming what is called an eye diagram. The bit error rate tester monitors the output of the 835-nm channel, comparing it with the input, to determine the bit error rate on that channel.

Figure 13a shows the four eye diagrams obtained when each channel is driven with a 622-Mbit/s signal, resulting in a 2.488-Gbit/s aggregate data rate. At this data rate, the maximum distance allowed by modal dispersion in a 62.5-μm multimode fiber is 500 m. **Figure 13b** shows the eye diagrams obtained with 500 m of fiber. **Figure 14a** shows the eye diagrams when each channel is driven at 1 Gbit/s, corresponding to a 4-Gbit/s aggregate data rate. At this rate, the maximum multimode fiber link length is 300 m. Eye diagrams with 300 m of fiber in the link are shown in **Figure14b**. The openness of these eyes is indicative of error-free operation. Measurements with the bit error rate tester indicate error rates of less than 10^{-11}, that is, for every 10^{11} bits transmitted on one channel of the data link, no more than one error is obtained.

Figure 13

Four-channel eye diagrams for a SpectraLAN data link driven at 622 Mbits/s per channel (2.488 Gbits/s aggregate data rate) through (a) ~ 1 m and (b) 500 m of 62.5-μm-core graded-index multimode fiber.

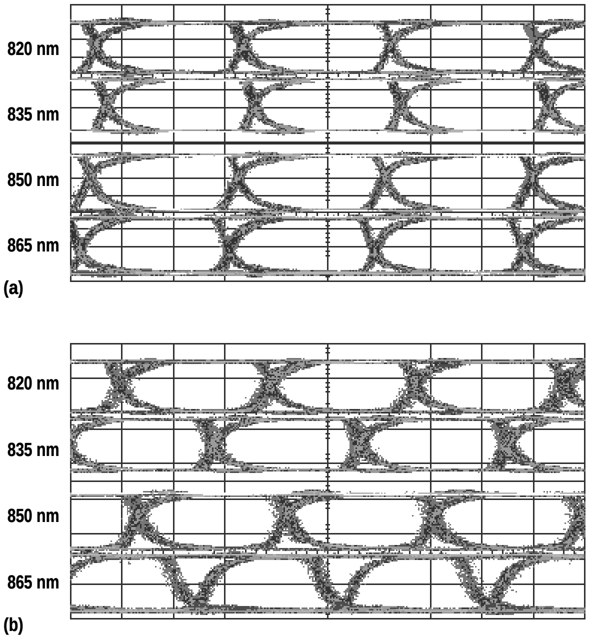

Figure 14

Four-channel eye diagrams for a SpectraLAN data link driven at 1.0 Gbits/s per channel (4.0 Gbits/s aggregate data rate) through (a) ~ 1 m and (b) 300 m of 62.5-μm-core graded-index multimode fiber.

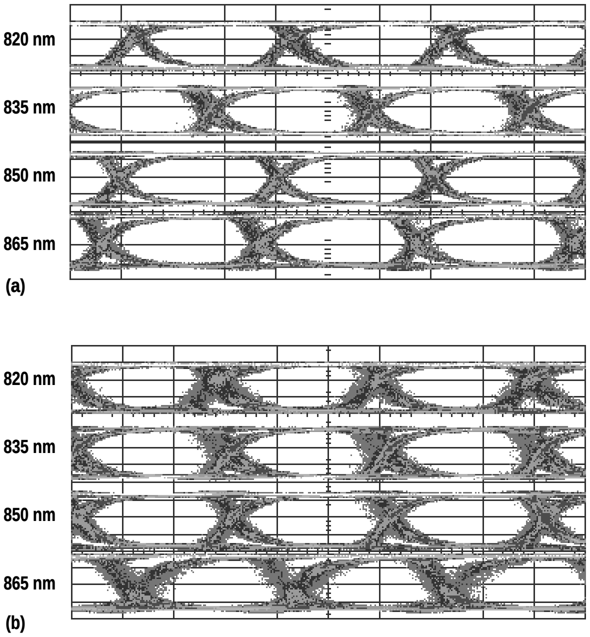

Conclusion

We have demonstrated four-channel error-free operation of the first-generation SpectraLAN module at aggregate data rates of 2.5 and 4.0 Gbits/s over distances of 500 m and 300 m, respectively. This will allow existing 62.5-μm multimode fiber-optic links to carry four times higher data rates than is possible with conventional methods. In addition, the module we have demonstrated is compact and potentially low-cost, giving promise that this technology can be practical for use in local area network backbones.

There is still a lot of work to be done in further reducing the size and cost of the module, and in investigating even higher data rates. The fiber-optics industry is moving towards a standard transceiver package that is only 0.5 inch wide and less than 1.5 inches long. To be competitive, future generations of our WDM module must fit into this small footprint. Costs must also be further reduced. The number of separate pieces that go into the module must be reduced whenever possible. One example of a current investigation towards this goal is in the area of VCSELs. The first-generation SpectraLAN module uses four discrete laser dies, each with a different wavelength. Work is under way to develop a single monolithic multiwavelength VCSEL array that can replace the four dies with a single die, requiring no separate alignment between lasers. Geometries that simplify the multiplexer and demultiplexer are also under investigation, as well as less-expensive technologies for fabricating and aligning the multichip module.

Data rates higher than 4 Gbits/s must also be addressed in future generations of SpectraLAN. The ATM forum, which sets standards for high-speed local area networks, is already considering standards for 10-Gbit/s links. The gigabit Ethernet committee will also most likely begin considering 10-Gbit/s standards within the next couple of years. WDM may be the only solution for multimode fiber-optic networks operating at such a high data rate. Short-wavelength lasers in the range of 850 nm might not be suitable for 10-Gbit/s networks, particularly if link lengths exceed 100 m. Lasers operating in the

1300-nm wavelength range have many advantages, including fiber bandwidth and eye safety. We plan to investigate WDM solutions using 1300-nm lasers as well as short-wavelength lasers.

Acknowledgments

The authors gratefully acknowledge the contributions of Elena Luiz, the process technician responsible for fabricating the polymer waveguides and dielectric filters, Jean Norman and Susan Davies, who performed the die attachment and wire-bonding, and David Dolfi, project manger. We also would like to acknowledge Chun Lei, Rick Schneider, and Sheila Steward for supplying VCSELs of the appropriate wavelength, Chris Kocot and Jeanne Glenn for supplying detectors, and Kenneth Hahn, Kirk Giboney, and Joseph Straznicky for their assistance.

References

1. *Hewlett-Packard Fiber-Optics Components Selection Guide*, 1996, p. 11.

2. D.C. Hanson, “IEEE 802.3 1.25 GBd MMF Link Specification Development Issues,” *Contribution to IEEE 802.3z*, Enschede, Netherlands, July 8-12, 1996.

3. K. Yamashita, Y. Koyamada, and Y. Hatano, “Launching condition dependence of bandwidth in graded-index multimode fibers fabricated by MCVD or VAD,” *IEEE Journal of Lightwave Technology*, Vol. 3, no. 3, June 1985, pp. 601-607.

4. J.P. Laude, *Wavelength Division Multiplexing*, Prentice Hall, 1993.

5. W. J. Tomlinson and G. D. Aumiller, “Optical multiplexer for multimode fiber transmission systems,” *Applied Physics Letters*, Vol. 31, no. 3, 1977.

Brian E. Lemoff

Brian Lemoff is a member of the technical staff at HP’s Communications and Optics Research Laboratory. His current responsibilities include development of optical components, such as wavelength multiplexers and demultiplexers, as well as WDM laser and detector arrays. Brian joined HP Labs in 1994 after completing his doctorate in physics at Stanford University. His nonwork interests include playing the folk harp and cooking.

Lewis B. Aronson

Lew Aronson’s work at HP Labs is focused on new fiber-optic link technologies for LANs. He is working on overall systems issues as well as the electronics, packaging, and link measurements related to these technologies. Lew joined HP Labs in 1992, the same year he received a doctorate in applied physics from Stanford University. He is married, has four children, and enjoys a variety of activities with his family as well as playing the piano and hiking.

Lisa A. Buckman

Lisa Buckman joined HP in 1996 after receiving her doctorate in electrical engineering from the University of California at Berkeley. She works in the Communications and Optics Research Laboratory of HP Labs, where she is helping to develop system applications for SpectraLAN and other projects. She is a native of northern California who enjoys softball, hiking, skiing and biking.

▶ Go to Next Article
▶ Go to Journal Home Page