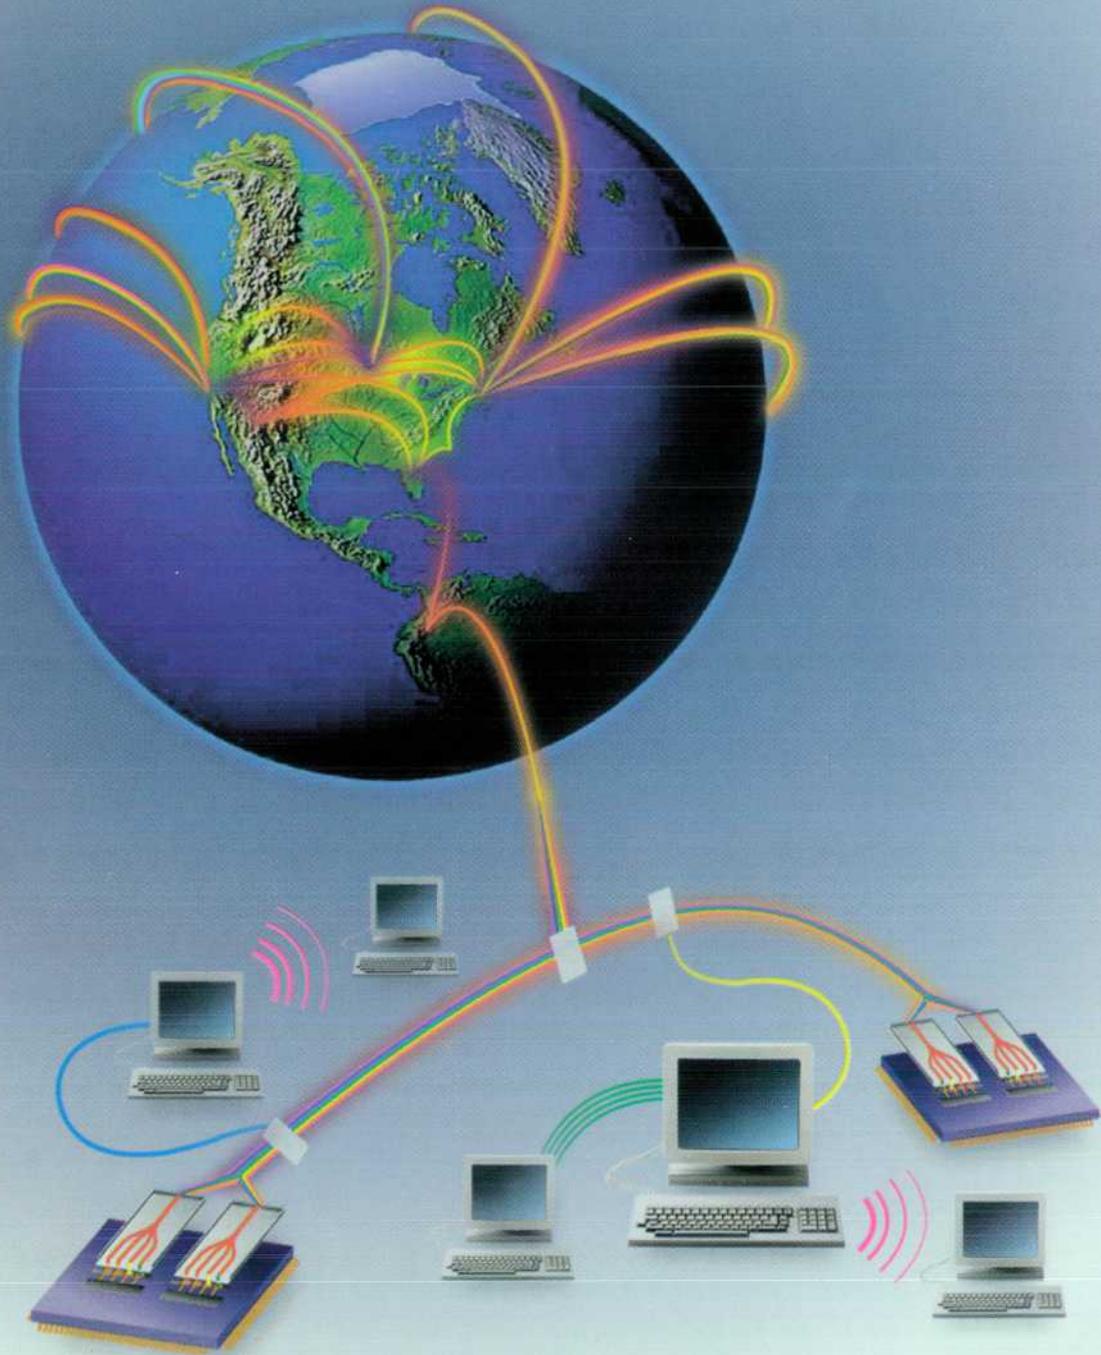


HEWLETT-PACKARD

Journal

DECEMBER 1997

Technical Information from the Hewlett-Packard Company



 HEWLETT®
PACKARD

DECEMBER 1997

THE COVER



An artistic rendition of global internet communications, showing fiber-optic technology as the backbone for universal connectivity.

Tera Era

We are approaching an era in which people will need 1-Gbit/s communications ports in their offices, their homes, and even on the road. These high-speed communications ports will enable telecommuting, telemedicine, tele-education, and a variety of multimedia applications for entertainment and computing. These demands for high-speed communications will require new telecommunications and data communications infrastructure with terabit/s data rates. Additionally, these communications networks will require very high-speed computers (Tflops), very high-speed instrumentation (THz), and large information storage (Tbytes). The technologies needed to reach these rates are being worked on at many R&D organizations around the world. In fact, many demonstrations have been completed in 1996-1997 showing 1-Tbit/s communications links over more than 100 kilometers, 1-Tflops computers, and 1-THz instrumentation. Thus, we can safely say that the Tera Era is on the horizon.

Volume 48 • Number 5

The Hewlett-Packard Journal Staff

Steve Beitler, Executive Editor
Charles L. Leath, Managing Editor
Richard P. Dolan, Senior Editor
Susan E. Wright, Publication Production Manager
Renée D. Wright, Distribution Program Coordinator
John Nicoara, Layout/Illustration
John Hughes, Webmaster

Printed on recycled paper.

©Hewlett-Packard Company 1997 Printed in U.S.A.

Advisory Board

Rajeev Badyal, Integrated Circuit Business Division, Fort Collins, Colorado
William W. Brown, Integrated Circuit Business Division, Santa Clara, California
Rajesh Desai, Commercial Systems Division, Cupertino, California
Kevin G. Ewert, Integrated Systems Division, Sunnyvale, California
Bernhard Fischer, Böblingen Medical Division, Böblingen, Germany
Douglas Gennetten, Greeley Hardcopy Division, Greeley, Colorado
Gary Gordon, HP Laboratories, Palo Alto, California
Mark Gorzynski, Inkjet Supplies Business Unit, Corvallis, Oregon
Matt J. Hartline, Systems Technology Division, Roseville, California
Kiyoyasu Hiwada, Hachioji Semiconductor Test Division, Tokyo, Japan

Bryan Hoog, Lake Stevens Instrument Division, Everett, Washington

C. Steven Joiner, Optical Communication Division, San Jose, California

Roger L. Jungerman, Microwave Technology Division, Santa Rosa, California

Forrest Kellert, Microwave Technology Division, Santa Rosa, California

Ruby B. Lee, Networked Systems Group, Cupertino, California

Swee Kwang Lim, Asia Peripherals Division, Singapore

Alfred Maute, Waldbronn Analytical Division, Waldbronn, Germany

Andrew McLean, Enterprise Messaging Operation, Pinewood, England

Dona L. Miller, Worldwide Customer Support Division, Mountain View, California

Mitchell J. Mlinar, HP-EESof Division, Westlake Village, California

WHAT'S AHEAD

This special issue of the HP Journal contains papers on technologies needed for high-speed communications. The papers address the long-haul, access, and local area networks using fiber-optic, copper wire, and wireless communications links. This series of articles is split into two parts: part 1 (this month) will address the area of communications technologies for fiber-optic and copper wire links. Part 2 (next issue) will address the area of wireless communications.

Besides six articles on wireless communications, the February issue will also contain an article on an option in the HP 8648A signal generator for pager testing, and an article on HP CaLan, a cable system tester.

*Waguhih Ishak
Director
Communications and Optics
Research Laboratory
HP Laboratories*

M. Shahid Mujtaba, HP Laboratories, Palo Alto, California
Steven J. Narciso, VXI Systems Division, Loveland, Colorado
Danny J. Oldfield, Electronic Measurements Division, Colorado Springs, Colorado
Garry Orsolini, Software Technology Division, Roseville, California
Ken Poulton, HP Laboratories, Palo Alto, California
Günter Riebesell, Böblingen Instruments Division, Böblingen, Germany
Michael B. Saunders, Integrated Circuit Business Division, Corvallis, Oregon
Philip Stenton, HP Laboratories Bristol, Bristol, England
Stephen R. Undy, Systems Technology Division, Fort Collins, Colorado
Jim Willits, Network and System Management Division, Fort Collins, Colorado
Koichi Yanagawa, Kobe Instrument Division, Kobe, Japan
Barbara Zimmer, Corporate Engineering, Palo Alto, California



The Hewlett-Packard Journal Online

The Hewlett-Packard Journal is available online at:

<http://www.hp.com/hpj/journal.html>

A quick tour of our site:

Current Issue—see it before it reaches your mailbox.

Past Issues—review back to February 1994, complete with links to other HP sites.

Index—search back to our first issue in 1949, complete with an **Order** button for issues or articles you'd like for your library.

Subscription Information—guidelines for U.S. and international subscriptions and a form you can fill out to receive an e-mail message about upcoming issues and changes to our Website.

Previews—contains an early look at future articles.

WWW

Articles

6 Communications Challenges of the Digital Information Utility

Joel S. Birnbaum

The director of HP Labs shows how the Internet and World Wide Web are the harbingers of an emerging era of pervasive computing.

11 Residential Communications

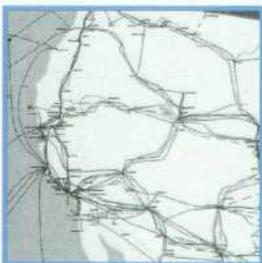
Daniel A. Pitt

The demand for Internet access in the home is driving advances of many kinds in residential communications.

19 Optical Networks: Backbones for Universal Connectivity

Robert C. Bray and Douglas M. Baney

The authors discuss the rapid growth in communications traffic on the world's fiber-optic backbone and the new technologies that address it.

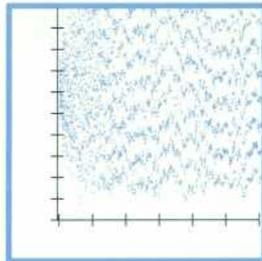


Map courtesy of KMI Inc., Newport, RI, U.S.A.

32 Data Transmission for Higher-Speed IEEE 802 LANs Using Twisted-Pair Copper Cabling

Steven G. Methley, Alistair N. Coles, and Eric Deliot

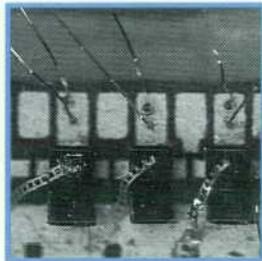
Can twisted-pair copper cabling support LAN speeds above 100 Mbits/second?



42 SpectraLAN: A Low-Cost Multi-wavelength Local Area Network

Brian E. Lemoff, Lewis B. Aronson, and Lisa A. Buckman

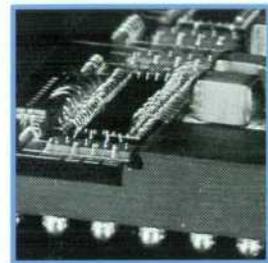
HP's SpectraLAN project is an approach to multiwavelength LANs with multigigabit-per-second data rates.



53 Gigabyte-per-Second Optical Interconnection Modules for Data Communications

Kenneth H. Hahn, Kirk S. Giboney, Robert E. Wilson, and Joseph Straznicky

Optical interconnects are essential to meeting the bandwidth needs of evolving communications technologies.



62 Developing Leading-Edge Fiber-Optic Network Link Standards

David G. Cunningham, Delon C. Hanson, Mark C. Nowell, and C. Steven Joiner

The authors discuss HP's role in developing and promoting open standards for the global fiber-optic LAN market.

74 1300-nm Strained Quantum Well Lasers For Fiber-Optic Communications

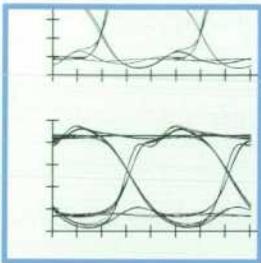
William S. Ring, Simon J. Wrathall, and Adrian J. Taylor

Uncooled semiconductor lasers offer compelling price/performance advantages.

86 Modeled Optimization and Experimental Verification of a Low-Dispersion Source for Long-Haul 2.488-Gbit/s Systems

Stephen M. Gee, Herbert Lage, Chris Park, Kevin A. Williams, Richard V. Penty, Ian H. White, and Joseph A. Barnard

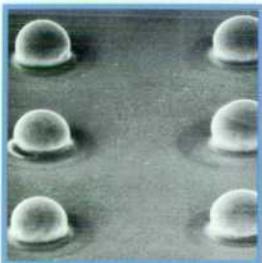
Low-dispersion-penalty diodes are crucial in today's world of faster digital modulation and longer transmission length.



102 Flip-Chip Photodetector for High-Speed Communications Instrumentation

Tun S. Tan, David M. Braun, Tim L. Bagwell, Chris Kocot, Joseph Straznicki, and Susan R. Sloan

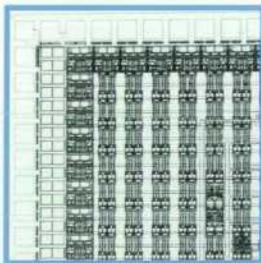
The authors describe a promising new micro-flip-chip process for packaging photodetector and electronic amplifier ICs.



111 A 2.488-Gbit/s Silicon Bipolar Clock and Data Recovery Circuit for SONET Fiber-Optic Communications Networks

Richard Walker, Cheryl Stout, Chu-Sun Yen, and Lewis R. Dove

A new clock and data recovery circuit from HP can regenerate clean data from corrupted signals.



120 Testing Erbium-Doped Fiber Amplifiers

Jim Stimple

Increasingly sophisticated testing of erbium-doped fiber amplifiers is crucial to their ability to deliver significant performance gains consistently.

127 Hewlett-Packard Professional Books

Recently published HP Press books and how-to-order information.

The Hewlett-Packard Journal Online

<http://www.hp.com/hpj/journal.html>

What's new?

- The Previews section, contains an early look at our February 1998 issue featuring:

Wireless Communications: A Spectrum of Opportunities

The IrDA Standards for High-Speed Infrared Communications

RF Technology Trade-offs for Wireless Data Applications

An Enhancement-Mode PHEMT for Single-Supply Power Amplifiers

Direct Broadcasting Satellite Applications

0.1- μ m Gate-Length AlInAs/GaInAs/GaAs MODFET NMIC Process for Applications in High-Speed Wireless Communication

Pager Testing with a Specially Equipped Signal Generator

HP CaLan: A Cable System Tester that is Accurate Even in the Presence of Ingress

In addition, you will find:

- A nostalgic look at our June 1972 issue on the HP 35 handheld calculator.
- A small example showing how animation can be used to illustrate TDR waveforms.

The HP Journal Website is evolving rapidly—we invite you to keep up!

WWW

Communications Challenges of the Digital Information Utility

Joel S. Birnbaum

The Internet and World Wide Web are forerunners of a digital information utility that in time will provide computing as well as information to society, just as other utilities provide water and electric power.



Joel S. Birnbaum

Joel S. Birnbaum is Hewlett-Packard's Senior Vice President of Research & Develop-

ment and Director of HP Laboratories (the company's central research and development organization), a position he assumed in 1991. Joel joined HP in November 1980 as founding director of the Computer Research Center of HP Laboratories in Palo Alto, California. Before that, he spent 15 years at IBM Corporation's Thomas J. Watson Research Laboratory in Yorktown Heights, New York, where he last served as Director of Computer Sciences. His personal contributions are in the areas of distributed computer system architecture, real-time data acquisition, analysis and control, and RISC processor architecture. He holds a Bachelor's degree in Engineering Physics from Cornell University (1960), and MS (1961) and PhD (1965) degrees in Nuclear Physics from Yale University. He is a member of several professional science, engineering, and computer organizations. He also serves on engineering advisory councils of Carnegie Mellon University, Stanford University, and the University of California at Berkeley.

This is a turbulent, uncertain time in communications and computing because major changes are taking place at the same time in related measurement, computing, and communications technologies. This is a time when suppliers of technology, operators, and users alike must attempt to predict the future, yet that's never been harder. As Yogi Berra, the former manager of the New York Mets, once said, "Prophecy is hard, particularly about the future."

In this article, I will explore the idea that the Internet and the World Wide Web built upon it are the earliest form of what will become a digital information utility that eventually will bring information and computing to most homes, schools, and businesses. Because important breakthroughs are needed in high-speed communications, network

management, security, and many other issues that result from the enormous increases in the scale of these emerging global nets, we might not see this utility for a decade or more. I believe, though, that market forces will ultimately cause these barriers to be overcome and that the resultant digital information utility will create a new style of computing that will both demand and enable a new type of high-speed communications infrastructure.

Pervasive Information Systems

For about 20 years, I have pursued a dream that one day computing would become a pervasive technology—that is, more noticeable by its absence than its presence, just as automobiles, television sets, telephones, and many other technologies have become part of everyday life for most people.

Figure 1 shows the progression of computing technology over several decades. It shows computing elements below the S-curves and shows how they are interconnected above the curves.

A key transition occurred at the beginning of the 1990s. At that point, our ability to make devices small enough let us begin specializing them according to their function and producing them in sufficient volume that they could be priced for mass markets. The architected interface between the requesters and providers of services—the so-called clients and servers that produced open systems—is the enabling transition for computing technology to become pervasive within the next decade. Because only people born after a technology has been invented think of it as part of the environment, and not as a technology, tomorrow's children will think of computers the way we think of telephones and TV today.

For a technology to become truly pervasive, it must transcend being merely manufacturable and commonplace. It must become intuitively

accessible to ordinary people and deliver enough value to justify the large investment required to build the supporting infrastructure. Think for a moment how easily most of us can use a telephone or drive a car with an automatic transmission; think too of the investment needed to build the infrastructure behind the telephone and highway systems.

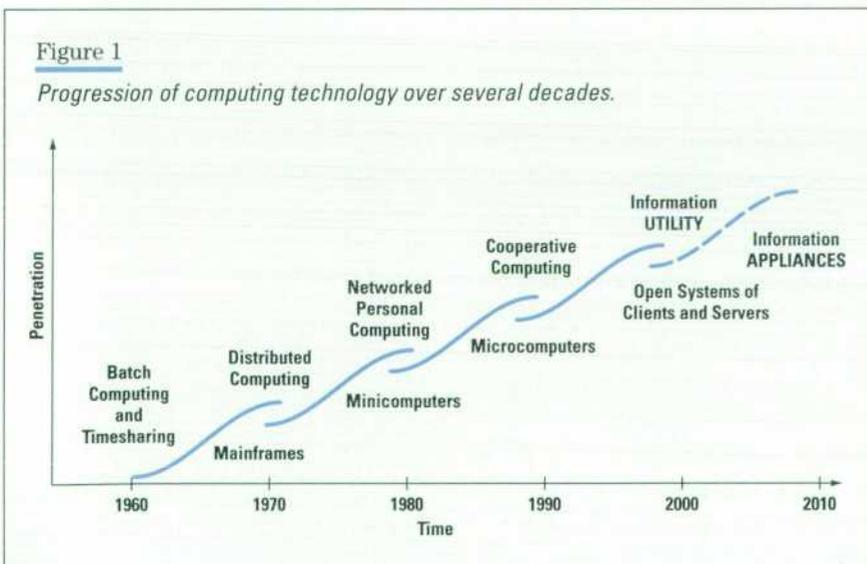
In the next century, computers will often be embedded in devices dedicated to performing a particular task, and they will let people think in terms of the task being performed instead of the underlying technology. A comparison with electric motors and consumer appliances helps convey this idea. We think of the function of a washing machine, electric toothbrush, or VCR, not about the motors within. In fact, the motors are very different and have required sophisticated engineering development. However, we do not think of this when we use these appliances (note that they are almost always named by what they do). We do not know, nor do we care, how many motors we own; the same will be

true for computers when they too are pervasive.

I believe that we will see the emergence of information appliances, which I view as devices that will offer a means to a particular end and whose use will seem quite natural to most people. They will differ greatly from portable general-purpose computers because of what these appliances will do and because they will be easier to use. We expect appliances to evolve to hide their own complexity, just as today's one-button automatically tuned television set has replaced its less-complex but harder-to-adjust ancestor. Because of continued advances in semiconductors and software, tomorrow's information appliances will be able to do the same and likely will conform to a simple model of use, like the "Neutral-Drive-Reverse" of an automatic transmission. The Internet and the World Wide Web could not have achieved such dramatic growth without the concurrent invention of the intuitive point-and-click-style browsers, which let users think about what they want to access without concern for where or how it is stored.

Figure 1

Progression of computing technology over several decades.



The Digital Information Utility

I mentioned earlier that the Internet and World Wide Web are the forerunners of a digital information utility that in time will provide computing as well as information to society, just as other utilities provide water and electric power. User expectations certainly will change once we begin thinking of information flowing over a digital utility. Like all successful infrastructures, the digital information infrastructure will have to be so dependable that we notice it only when it fails. It will also have to be

secure, endure over generations, be found everywhere, and serve a purpose important to almost all of the population.

Like other pervasive infrastructures, it will foster new industries. Think for a moment of how the public highway infrastructure, which made the personal automobile a pervasive technology, spawned many new industries such as auto insurance, car rentals, driver training, car washes, muffler shops, and gasoline stations, to name just a few.

In a world in which clients—whether general-purpose computers or information appliances—connect to this utility, most people will pay for their computing by usage. This will change what is now a capital investment into a competitive service such as that provided by the electric power and water utilities. This will not be the same as timesharing, which was proprietary, closed, and often location-dependent. In client-utility computing, open resources, located arbitrarily, will be combined as needed. In fact, one day, when the communications bandwidth is great enough and the costs low enough, it will no longer matter where the computers are located or which manufacturer makes them. Today's suicidal obsolescence schedule will be replaced by the capacity requirements of the service provider. Just as users don't notice when their electric utility replaces a generator at the power plant, so information utility users shouldn't detect upgrades.

Quality of service will become a crucial competitive differentiator at the systems level because users will expect the information utility to be available, ready and waiting, just

as today they pick up a phone and expect a dial tone.

A key aspect of this information utility is that it will be digital. This means that devices that send and receive information can be independent of each other instead of having to be designed in matched pairs, as analog fax machines or TV and radio transmitters and receivers must be today. Given the appropriate interchange standards, your digital fax machine, for example, could begin communicating with your television, which could also receive your newspaper as well as diagnostic information from your car.

This change also means that many of the decisions now made at the transmission end instead can be done at the receiver. Once information interchange standards are in place, appliance and peripheral families will emerge, and many of these will be able to communicate directly without the intervention or invocation of a general-purpose computer and the attendant, cumbersome general-purpose operating system. In fact, many, if not most, of the computers of the next generation will be enormously powerful nonreprogrammable invisible processors. I'm thinking here, for example, of the several dozen embedded processors hidden in a modern car that control the car's ignition, suspension, braking, steering, engine-management, and climate-control systems and that provide diagnostic information to the driver and the mechanic. No one has ever asked a car dealer whether these computers run UNIX[®] or Windows[®] NT—the interface is the steering wheel or brake pedal, and function is augmented without the need to introduce an

unfamiliar interface or force users to pay attention to the inner workings of the software. The embedded processors of the future will enable users to invoke powerful functions at a much higher level of abstraction than is common today.

Everything Could Have a Web Page

It is already practical to embed a web server, which can be a small amount of code in a very inexpensive microprocessor, in individual devices. This means that everyday appliances can have a web page at negligible cost. With a conventional browser, you could easily check things like home security and heating or cooling systems and control individual appliances like a hot tub or toaster from anywhere that you can click on a web page. We may even make it possible for ordinary adults to program a VCR! In fact, any instrument's front panel, or a subset or simpler replacement of it, could be viewed from anywhere by people with access rights. This would let engineers and scientists inexpensively collaborate across vast distances and make practical remote maintenance for low-cost devices. Virtual instruments, which are networked combinations of real ones, could be created. Notification capability can be incorporated; imagine, for example, a printer that automatically signals a supplier to ship a replacement toner or ink cartridge when some usage threshold has been crossed.

If the communications pipe is fast enough and cheap enough, this also means that you can do things far away that appear to be done locally. Distributed computing power, remote distributed measurements, very rich user interfaces in small, inexpensive

devices, and remote printing and scanning all become practical and will enable applications at a scale and cost unthinkable today.

I believe that client-utility computing will provoke as great a change in the computing industry as open systems did in this decade. The new paradigm will do for computation what the Web did for data and will produce such dramatic decreases in cost of ownership, with concomitant increases in uptime and accessibility, that those companies that do not react to the opportunity will not be able to satisfy customer expectations. History tells us that this can have dire consequences.

Much technology—consisting mostly of middleware to address issues of scale, interoperability, security, robustness, and end-to-end systems management—is needed. Developing this technology won't be easy, but HP and many of the world's computer and software companies and research laboratories are today busy developing practical engineering solutions to these problems. It is just a matter of time before the greatly decreased cost of ownership of client-utility systems, coupled with their functional advantages, makes this the way that most people will access multimedia information and solve problems requiring computation.

Needed: Communications Breakthroughs

A number of telecommunications technology breakthroughs are needed for this dream of pervasive information systems to come true. We are entering a period that some at HP Laboratories have begun calling the "Tera Era" because the demanding technical requirements needed to support inexpensive high-bandwidth

networks are measured in quantities of trillions—that is, trillion-bit-per-second transmission, trillion-byte memories, and trillion-instructions-per-second computers. The viability of the digital infrastructure for multimedia documents, appliance-utility computing, and distributed remote measurements depends on a number of key technologies coming together. Particularly critical is having high bandwidth at low cost, and many people are working to bring this about.

Other articles in this issue discuss some of these technologies, but I would like to focus for a minute on important developments taking place in optics technology that could have a huge effect on the telecommunications system. The theoretical capacity of the fiber is vast—something on the order of 25,000 GHz. In fact, though, we typically don't take advantage of anything but a tiny fraction of that capability. If we could find a way to send and switch more signals through a fiber-optic cable, we could increase the system's capacity by at least two or three orders of magnitude. It would mark a radical change in such systems because for the first time the electronic switching components would become the bottleneck instead of the transmission lines connected to them. This suggests that an all-optical system is needed, and the most promising approach is called wavelength-division multiplexing (WDM). This is not a new idea, but it is now becoming realistic to think of such systems being widely deployed in the not-too-distant future.

Fiber technology has advanced over the last 20 years to the point that the distance across which a usable light

pulse can be sent has grown from a fraction of a kilometer to hundreds of kilometers, and cost has plummeted concurrently. If we could transmit and then select all the theoretically possible frequencies, a wavelength-division multiplexed system could work the way that a radio does. That is, at one end, a particular station chooses its frequency; at the other end, the user has what amounts to a big dial. Depending on whom you want to be connected to, you turn the dial and change the frequency of the receiver. Sometimes a movie comes over the pipe, sometimes a newspaper, sometimes the result of an economic model from a distant supercomputer, sometimes your child's voice.

This is, of course, a lot more easily said than done. Early WDM systems used mechanical electrooptical devices for frequency selection. Essentially, movable gratings would preferentially select a single frequency. However, this is not very practical because such interferometers are slow, expensive, and limited in frequency range. Laboratories today are investigating promising low-cost high-performance optoelectronic transceivers and other electronic and photonic devices that could replace or augment existing terminal equipment.

We also need other new or dramatically improved technology, including multigigabit-per-second semiconductor lasers and photodetectors, multigigabit-per-second integrated electronics for laser drivers and clock and data recovery, large input/output cross-connect switches, and optical circuits such as add/drop filters and wavelength converters.

If these technology challenges can be met, the results will be fabulous. Just one strand of fiber, in principle, could carry a billion separate audio telephone calls—all the telephone calls for a large city. Many users of the digital information utility will want enough bandwidth, say, a 100-MHz channel, for high-resolution real-time multimedia. That's ten times an entire Ethernet for each user. With WDM, a single optical fiber could carry a quarter million of these superbandwidth channels.

Network Measurement and Management

The digital utility will also require much more sophisticated network management because users will expect the quality of service they now get from their telephone system. The Internet as we know it today is nowhere close to this.

To manage applications end-to-end, there can be no alternative to making distributed, continuous measurements of performance, capacity, and usage statistics and relating them to a model of the system, or of part of it. By measuring what's actually happening across the entire system, you can, among other things, adjust its capacity, detect many types of fraud, predict where performance bottlenecks are likely to occur, locate outages, and identify unused parts of the system for reserve capacity. Once again, Internet technology will reduce the complexity and lower the cost. In fact, HP has built a successful prototype of such a system that is operating today in a London cellular telephony trial.

Network management and measurement are the Achilles heel of the robust, flexible infrastructure that operators and users want. I think

these will be the pacing core technologies of the Information Age, and they demand the immediate attention of telecommunications and computer manufacturers alike. It will be important to develop international standards; and while the technology will be similar in some ways to that now used by distributed computer systems and the telephone networks, it will have to solve problems of scale and speed that people have not had to deal with before. Systems with millions, or even tens of millions, of nodes will be commonplace, and the heterogeneity of the hardware and software will be unprecedented.

Once developed, this core capability of measuring and managing the evolved Internet itself will be extended to enable a vast range of distributed measurement applications that today would require specialized, expensive systems. The utility provides the infrastructure to link distributed sensors at low cost in a ubiquitous way. Many industries, such as healthcare, agriculture, and transportation, will be transformed by this ability to operate on a continental basis.

Security

An issue much in the news is security, which stands in the way of many commercial applications. Security is a hard problem because we're on a difficult tightrope. On one hand, we want maximum interoperability for authorized, authenticated users among all computers, all appliances, and all nations. Essentially, this means that we impose no barriers to interconnection. But this aim of unfettered interoperability conflicts with access control, privacy, and system integrity. Citizens' rights to privacy are often in direct conflict

with the needs of national security and criminal justice agencies. Building a truly open, global system is technically antithetical to a secure system with good performance and attractive cost, creating a difficult set of technical, social, and political trade-offs to be resolved, and compromises are not easily achieved.

I believe that the security issue will be solved more easily than the bandwidth and network-management problems. Most of the world's computing and communication companies are working feverishly on security-related issues because these issues stand in the way of the profits that can be reaped from electronic commerce. Although no perfect solution is likely to emerge, I think that an acceptable de facto standard, based on a sensible range of trade-offs and compromises, will emerge because of the overwhelming financial and market forces demanding it.

Conclusion

In closing, the telephone system that we know today and the Internet and Web technology built upon it are precursors of a global digital information utility. The Internet is delivering today at low bandwidth and relatively high cost entirely new classes of information and services. As the bandwidth and usage grow, costs inevitably will decline. The resulting information utility will dramatically change computing as well as telephony and the delivery of multimedia information. In time, we will think of today's systems as quaint.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

X/Open® is a registered trademark and the X device is a trademark of X/Open Company Limited in the UK and other countries.

Windows is a U.S. registered trademark of Microsoft Corporation.

Residential Communications

Daniel A. Pitt

Establishing a communications infrastructure to get information to, from, or around a residence is not a straightforward task today. However, in the future the equipment and wiring within a residence for Internet communications will be treated the same as the wiring and equipment for services such as telephone, electricity, and cable television are treated today.

Residential communications involves getting information to, from, and around where people live. What makes this different (from, say, enterprise communications), unique, or interesting? Well, for one thing, people do different things at home than they do at work (although they are working more at home today). Also, there are many more homes than there are places of work. Perhaps most important, residents themselves pay for most of the communications out of their discretionary budgets. What is clear is that residential communications is growing in magnitude, type, and extent. It is being driven by demand (as for Internet access) and competition (from deregulation) rather than by technological advances, which are being sought as a means of satisfying demand.

As we discuss residential communications, we must keep in mind two distinctions. The first is between communications to the residence and communications within the residence. The second is between different types of residences. All too often we fall into the trap of equating the residence with the owner-occupied single-family home, but the worldwide market includes huge numbers of potential customers in single-family rental and multiple-dwelling units. These can vary from duplexes to high-rises, and they differ from single-family homes in their construction materials, infrastructure, and ownership. Indeed, the fact that these dwelling units are not owned by their residents is significant.



Daniel A. Pitt

Dan Pitt came to HP after 13 years at IBM where he worked as a networking

architect responsible for token ring architecture and token ring standards. He joined HP in 1992 as manager of the multimedia systems department at HP Laboratories. He recently left HP to become vice president of Bay Networks. He holds a PhD degree in computer science from the University of Illinois. Dan was born in Madison, Wisconsin, is married and has two children. He is comfortable in several languages besides English, particularly French and Italian. He plays the mandolin and the trombonium, and is a prolific writer of profound limericks.

Communications to the Residence

Communications to a residence today comes in many forms and what is notable is that they are all service-specific. Telephone lines provide telephone service. Cellular networks provide telephone service. Cable TV networks provide video entertainment. Satellite networks provide video entertainment. Paging networks provide paging services. Where they exist, energy-management networks provide energy-monitoring and control services. Moreover, each type is owned and operated by a service provider whose business is to sell a particular service. To date, only the telephone and cellular companies view communications as their business, and the others view services as their business. While the telephone and cellular providers welcome the growth of revenue from fax and data applications on their networks, they are not adequately equipped to handle the growth. These *residential access networks* base their economics on the cost of bringing the network to each residence and the anticipated revenue from each residence, especially as regulation diminishes.

The effect of deregulation and competition is an attempt by network operators to offer multiple services on their networks and to increase the capacity of their networks. These efforts stimulate the advance of network technology. Let us now examine some of these technologies. In each case we will connect the technology to the services it most naturally supports and speculate on its future. Specific technical aspects of these technologies are discussed in other articles in this issue.

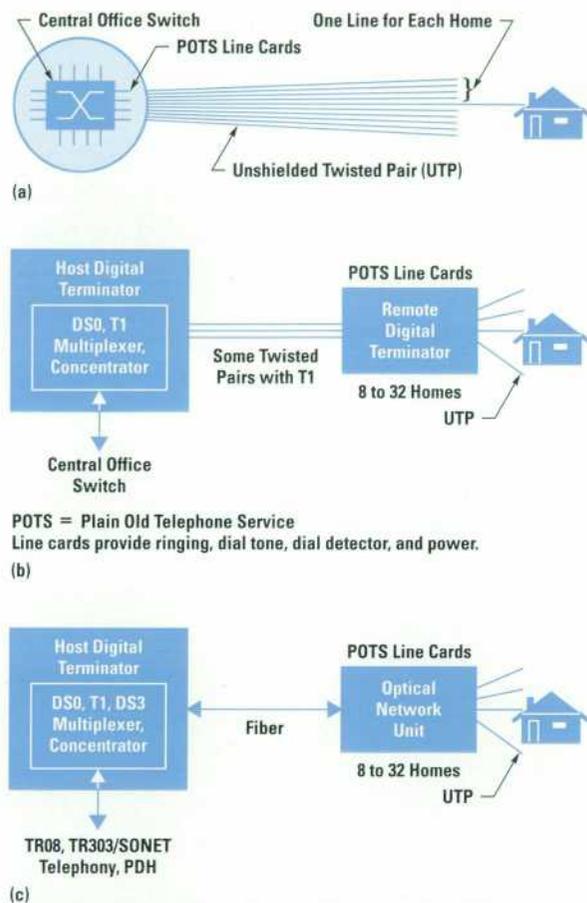
Wireline Telephony Networks

Over 80% of homes in the U.S.A., and a not dissimilar percentage elsewhere, connect to the central offices of their telephony providers with twisted-pair copper wires (see **Figure 1a**). In some cases a few dozen customers share a high-quality twisted pair partway from the central office to the home, but in all cases each home has 4 kHz of analog bandwidth dedicated to itself. A bidirectional telephone conversation occupies all of this band, or it can be modulated to carry bidirectional, half-duplex data. Advances in signal processing have enabled the data rate to rise recently from the common 14.4 kbits/s to 28.8 kbits/s, 33.6 kbits/s, and even 56 kbits/s.

Perhaps 12 to 15% of U.S. homes are served by digital loop carrier systems (**Figures 1b and 1c**). In these systems, the

Figure 1

Telephone network evolution. (a) All-metal beginnings in which twisted-pair copper wires go to each home. (b) Digital loop carriers. (c) Fiber-optic digital loop.



telephone signals are carried in digital form on an optical fiber from the central office to a point somewhere in the neighborhood, with all the signals for a group of homes multiplexed onto the same fiber. In a typical case, the fiber carries around 700 simultaneous calls (each at 64 kbits/s) giving an aggregate bit rate of 43 Mbits/s. Calling patterns suggest that this capacity can serve around 3500 homes. The remote node where the fiber terminates converts the digital signals to analog (for voice, dialing, ringing, and other signaling) and sends them on individual twisted pairs to each home. The user sees no difference between the all-copper system and the digital loop carrier system, but the operator's maintenance costs are lower with digital loop carriers. Telephone operators that wish

to offer higher-bandwidth services have three main alternatives: integrated services digital networks (ISDN), asymmetric digital subscriber line (ADSL), and fiber to the curb (FTTC).

ISDN. This service provides a circuit-switched service made up of two full-duplex 64-kbit/s data (or voice) channels and a 16-kbit/s signaling (or data) channel. Many ISDN modems allow the user to combine the data channels to gain 128 kbit/s full-duplex, which is markedly better than dial-up modems. Because ISDN is a switched service, the user dials the destination (which must also be equipped with ISDN), uses the connection, and hangs up. During the connection, a path through a telephone switch is dedicated to the connection, whether or not data is flowing. A switch connection is a costly resource for both the subscriber and the operator, and it has proven notoriously difficult to get an ISDN line up and running, so ISDN is not a good technology for bursty, packet data. Its speed is much too slow for digital video, which requires anywhere from 1.5 to 6 Mbit/s for acceptable quality.

ADSL. The asymmetric digital subscriber line provides a packet-stream overlay on analog twisted pairs at megabit rates, exploiting unused bandwidth on the twisted pair of up to 1 MHz. Voice is sent unchanged in its analog form in the 0-to-4-kHz band. Megabit rates are possible in only one direction, hence the asymmetry, which is typically ten times greater in the downstream direction, which is in the direction of the home. Rates today range from 1.5 to 8 Mbit/s in the downstream direction and depend on the length of the line. Some 95% of U.S. customers are within 18,000 ft of their central offices, and 1.5 Mbit/s should be able to reach all these subscribers. Those at 12,000 ft or less could receive up to 8 Mbit/s. The latest designs incorporate rate adaptation so that the ADSL modem adjusts its speed to the line quality (for which distance is one factor) in a static or dynamic manner. At the central office, the voice and data are separated, with the voice signal going through a circuit switch as usual and the data going into a separate packet network. A similar separator operates in the home, feeding the voice into the telephone wires and the data available to an Ethernet port.

Two competing technologies appear in ADSL products. One, endorsed by American and European standards organizations, employs discrete multitone (DMT) modulation, in which 1024 kHz of bandwidth is divided into 256

4-kHz segments, each modulated at a bit rate dictated by its noise characteristic at the moment. Thus, it is robust to narrowband interference. The other ADSL product, which has been on the market longer and is vying for standards approval, employs carrierless amplitude and phase (CAP) modulation. We believe it unlikely that either one will prove vastly superior to the other in real deployments.

The main advantage of ADSL is that it can be deployed without any modification to the infrastructure in the field (only endpoint modifications are necessary). It can be deployed one subscriber at a time when the subscriber agrees to pay for it and without the security risks of a shared medium. The barriers to ADSL's acceptance are (1) the cost of the equipment, which has been slow to drop below U.S.\$1000 per line, (2) the data line concentration task at the central office, (3) the continued maintenance of the copper plant, and (4) its inability to carry broadcast video. As a medium-term technology for Internet traffic, ADSL might be promoted where there is competing cable-modem service.

FTTC. Fiber to the curb provides very high data rates over an expensive plant that carries fiber very close to the home. Voice traffic for one to three dozen homes is conveyed in digital form using digital loop carrier technology to the end of the optical fiber. Data is conveyed on the same or a different fiber to the same point but is kept logically separate. At the end of the fiber, the voice is converted to analog form (with ringing and signaling converted as well) and frequency multiplexed with the data on a dedicated twisted pair or coax line to each home. Downstream data rates can reach 51 Mbit/s with upstream rates ranging from 2 to 20 Mbit/s. Like ADSL, these figures represent dedicated bit rates for each home, not shared rates. FTTC has not been deployed widely because of its high cost and its inability to carry analog broadcast video. However, its data capabilities are unsurpassed and it represents how most telephone operators would like to send voice in the future. A compromise on cost, data rate, and distance between ADSL and FTTC is called *very high bitrate digital subscriber line*, VDSL. VDSL takes fiber to within 3000 ft (instead of FTTC's <1000 ft), reduces the data rate to 25 Mbit/s, and lowers the system cost substantially. From the end of the fiber to the home the techniques of ADSL are applied.

Cable TV Networks

The networks used for those 30, 50, or 70 channels of broadcast analog video are being upgraded to carry digital video and two-way data services. This requires replacing the all-coax tree-and-branch network with a fiber link that leads to a much smaller coax tree to just a few hundred homes instead of thousands. This is called hybrid fiber/coax (HFC).

More than any other factor, the development of wideband linear AM lasers, allowing the optical fiber to be modulated exactly as if it were coaxial cable has made HFC possible. HFC expands the system bandwidth from 300 MHz (a typical value) to 750 MHz. The 6-MHz channelization is retained,* and the new channels each carry around 30 Mbits/s of digital traffic, be it video, Internet, or voice. All traffic reaches every home since the system was designed for broadcasting. At home, a subscriber tunes to a given channel to receive one analog TV program (by a cable converter box), or one of up to 10 digital TV programs carried in one channel (by a digital set-top box), or some Internet data packets (by a cable modem).

The cable modem and digital set-top box also transmit upstream to convey upstream data or program-selection commands. Upstream transmissions fall into the 5-to-42-MHz band, which is notoriously noisy and narrow. Downstream transmission typically employs 16-, 64-, or 256-level quadrature amplitude modulation (QAM) to squeeze as many bits into the 6-MHz channel as possible. The most common upstream modulation scheme is quadrature phase shift keying (QPSK). Some vendors advocate spread-spectrum techniques to combat high noise levels and some even claim satisfactory operation over all-cable plants. At this point proprietary media access control (MAC) protocols govern the shared upstream channel, but several organizations are working on standardizing the MAC.

The 30 Mbits/s of downstream digital traffic is shared among several hundred homes. Upstream capacity is in the range of 10 Mbits/s with individual homes limited in most designs to around 2 Mbits/s; this, too, is shared among these same homes. Telephony, which requires symmetric bandwidth, must share the same upstream capacity but uses its own dedicated downstream channel.

* The 6-MHz channelization divides the spectrum into channels 6 MHz wide with each channel used either for one analog television program or for digital traffic of any sort.

Eventually, Asynchronous Transfer Mode (ATM) is likely to multiplex all traffic types in each channel so that integrated receivers in the home can access nearly all of the home's traffic from a single 6-MHz channel and system capacity is optimized.

The strength of the HFC solution is that it carries on one network all types of traffic types including broadcast analog video, broadcast and interactive digital video, high-speed data, and telephony. Its weak points are its upstream capacity, the cost of upgrading the existing all-coax plants, and the lack of experience in highly reliable two-way communications. It is a natural upgrade choice for those cable operators that can afford it. However, it is not likely to become popular among telephone operators.

Wireless Networks

Cellular Telephony Networks. Cellular systems can be and are used for data, but their capacity is limited and their cost is high. Speeds are typically limited to 19.2 kbits/s and the charges for cellular digital packet data make it unsuitable for file transfer, web page downloading, or anything other than small message exchange. The main advantage of cellular networks is their accessibility from nearly any location, including while moving.

Terrestrial Microwave and Millimeter-Wave Networks. These networks are emerging as attractive alternatives to reaching the home without digging up roads or yards or stringing cable. They use a grid of base stations, often similar in siting to cellular stations, and fixed receivers on homes (so no mobility is provided). Multichannel multipoint distribution service (MMDS), also known as wireless cable, is a one-way scheme occupying 200 MHz of bandwidth in the 2-GHz band. So far only broadcast analog video has been sent on the 33 6-MHz channels assigned to these wireless networks. However, many of these channels are gradually being converted to digital using 16- or 64-quadrature AM. There is no upstream channel so interactive services are limited to dial-up return. A multichannel multipoint distribution base station can serve tens of thousands of subscribers (at distances of tens of miles) so the data capacity per subscriber is low.

Local multipoint distribution service (LMDS) offers much greater capability. In 1997 the U.S. Federal Communications Commission (FCC) allocated over 1000 MHz in the 28-to-30-GHz band for local multipoint distribution. Other

countries have allocations in that band and in the 40-GHz bands. Deployments are expected to start in 1998 with most of the bandwidth being used for downstream and broadcast transmission, partly because of the way the bandwidth was allocated.

Depending on the amount of broadcast traffic and because of the use of sectorized antennas, the downstream capacity ranges from 1 to 2 Gbits/s and the upstream capacity from 300 to 600 Mbits/s (aggregate). A typical cell radius is 1 km serving anywhere from 1,000 to 20,000 subscribers, so the capacity is considerable. The rooftop antenna need be no larger than one square foot and is not expensive, though it needs to be correctly aimed. The main problem with local multipoint distribution service is that a line of sight between the base station and the home is required, so the technology works poorly or not at all in heavily forested areas. Also, the entire transmission must be digital for adequate signal reception, so analog video is not supported.

Satellites. For quite a while satellite networks have been employed for broadcasting television programs. They also offer capabilities for other services. Geosynchronous satellites, which are used for TV today, can be and are being used for data by modulating a TV channel digitally as is done in cable networks. The difficulty is finding the right application for this. The satellite has no return channel so interactivity is not possible without using telephony return, which is awkward and undesirable. Moreover, today's satellites, designed for broadcasting, have footprints that reach tens of hundreds of millions of people, so the capacity per subscriber is minuscule. These features, along with the long transmission time, make the geosynchronous satellite a very suboptimal choice for interactive services.

As applications for data broadcasting grow, however, the satellite could become the ideal transmission medium. Low earth orbit satellites (LEOS) offer an altogether different service. Designed for telephony, these systems of dozens or hundreds of satellites offer ubiquitous access without the erection of base stations or the worry of being out of range of them. They also allow access by users in motion. However, if these satellites retain the data rates of terrestrial cellular services, they will be suitable only for telephony and control services and of little use for Internet services. Video, it seems, is out of the question, as is broadcasting of any sort.

There is also some activity in other wireless technologies, using various little places in the spectrum and speeds in the ISDN range. These are likely to remain niche services. The pursuit of an effective and affordable wireless data service with capacities of tens of Mbits/s continues mainly as a research activity.

Where Are these Network Technologies Heading?

What we expect is a period of experimentation with these new technologies in which we will try to determine:

- How easy they are to install
- How well they support various services
- Which services are profitable
- Which technologies lead to the greatest profitability in the shortest time.

The deregulation of the telecommunications business in most countries will result in competing operators, usually with different technologies, in many regions and to many homes. Our expectation is that some services will migrate to certain technologies that offer incomparable economic advantage, such as satellites for nonlocal broadcast video, while most remaining services will be carried jointly on a few multiservice networks. The subscribers will have little say regarding which technologies reach their homes, but will increasingly have more than one multiservice network at the home. Homes in wealthy urban areas will have the greatest choice the soonest, since, for example, a telephone company is more likely to accelerate ADSL or FTTC deployment to homes that are getting cable modem service from a cable operator. In general, subscribers will get more bandwidth to the home, lower cost, and greater choice of service type.

Communications within the Residence

Application Areas

The deployment of communications technology within the home will depend on, more than anything else, the applications that the user wishes the communications to enable. We see the most important of these being: work at home, entertainment, and personal life and communications. Work at home includes home business, daytime telecommuting, and night and weekend access to corporate networks. Entertainment includes passive video and

audio programming as well as interactive games and information. Personal life and communications includes home finance, home photography and videography, home e-mail, and of course fax and telephony. Other areas of lesser importance are home automation, home health care, and electronic commerce. We have not mentioned the Internet explicitly because nearly every application area cited above will use the Internet. Particularly significant as motivators for home communications are the following factors:

- Increasing use of the home computer to interact with entertainment and automation equipment (especially using web technology),
- The growth of home imaging (embodying digital cameras, printers, and storage media)
- The multiplicity of home computers in many homes
- The need to share digitally delivered programming to cable modems and set-top boxes by multiple devices within the home.

Obstacles to Home Communications

Communications within the home will not become pervasive (nor will the applications that require it) until a number of obstacles are overcome.

Wiring. Although wiring is not a major obstacle in offices, it is a major one in homes. If the wires are not already in the walls or do not go where they are needed, the user needs to install wires or forgo communications between rooms. This obstacle, in terms of cost and delay, motivates the development of wireless and powerline communications.

User Skills. Knowledge about how to hook up, configure, and use a home communications system cannot be so great that it reduces the market to the technically adept few. Even the technically adept have a dedicated support staff at the office for maintaining their complex communications networks.

Consumer Pricing. Affordability for home communications systems will probably be reached only with mass markets, which suggests standard approaches. Standards, as well, promote interoperability among different vendors, another consumer necessity.

Embedding Communications Functions. This obstacle, perhaps better stated as the invisibility of communications functions, reflects a consequence of the previous

two obstacles. For low cost and ease of use, it will be helpful to embed the communications into appliances that perform some application function, rather than develop a bevy of communications widgets that users will have to tangle with.

Privacy. Users will demand privacy not only for financial records but for all communications that leave or enter the home. The broadcast media of HFC, LMDS, and satellite violate this, so explicit techniques to ensure privacy must be added.

External Access. While external access is not necessary for all home communications, which increasingly will originate and terminate in the home, it will be mandatory for many devices in the home to deliver their full value to the user. How the combination of internal and external communications is provided most simply and economically is a matter of no small import.

In-Home Communications Technologies

In-home communications technologies fall into three categories: wireline, power line, and wireless. Wires are suitable when devices to be connected are in the same room, when existing wiring suits the communications, or when new construction allows the installation of special wiring. Otherwise, power line or wireless communications is highly desirable. Other than telephone twisted pair or cable TV coax, there is no installed base of any magnitude for in-home communications.

Ethernet and AppleTalk. These two network technologies are used today to connect home computers to their peripherals. With cable modems and ADSL modems terminating in Ethernet ports, Ethernet will be used increasingly to connect multiple home computers. Its 10-Mbit/s data rate, inexpensive adapter cards, and operation over unshielded twisted pair (with a hub) make it attractive for computer interconnection. Ethernet versions with 100-Mbit/s and 1000-Mbit/s data rates will trickle into the home as their cost comes down because of increased enterprise penetration and as the need for more bandwidth justifies their higher cost.

Universal Serial Bus (USB). This is a new peripheral bus that supports up to 12 Mbits/s with a simple connector and daisy chaining or hub/star wiring. It has features for isochronous and asynchronous traffic and we expect it to appear on PC motherboards in 1997.

IEEE 1394. This standard, which is also called *firewire*, has been embraced by the consumer electronics manufacturers for connecting digital video cameras, VCRs, and televisions. It is being adopted by home computer and set-top box manufacturers as well. IEEE 1394 is a high-performance serial bus that operates at 100, 200, or 400 Mbits/s over its own special cable. The cable is currently limited to 4.5 meters but developers are working to extend it to whole-house distances. It is hooked up in daisy-chain fashion, which allows flexible configuration and requires no hubs.

Consumer Electronics Bus (CEBus). This bus is an emerging home networking standard defined for operation over power line carrier, radio frequency, infrared, coax, and twisted pair. It defines a 10-kbit/s control channel and frequency spectra for data on coax and twisted pair but no data protocol, allowing various MAC protocols to be used. CEBus is being deployed for home automation and control of low-bandwidth applications and not for high-speed data (or video, or voice). The CEBus specification also includes a common application language (CAL), which is being adopted in the IEEE 1394 world as well. A simple, object-oriented programming language, CAL will find use for control and communications among security systems, light switches, and consumer goods.

Power Line Carrier. This transmission medium allows communications over the home's electrical mains. It offers the advantage of using wires that are already in the walls and requires no separate plug. Data rates of tens of kbits/s are available now, and claims for Mbits/s are viewed skeptically. Powerline transmission is highly susceptible to noise and attenuation through transformers. The latter is an advantage for transmissions intended to stay within one home. However, the signals reach all homes that share the same transformer, resulting in loss of privacy and the need to coordinate transmissions.

Infrared Communications. Infrared offers regulation-free, radiation-free communications within a room. Interference with neighbors is never a worry, although high levels of ambient light can be a problem. Infrared links exhibit a high inverse relationship between speed and distance, with Gbits/s speeds possible over a few centimeters, a few Mbits/s possible over a meter, and kbits/s in diffuse room-wide use. Many in the industry are interested in seeing the speed-distance product increase, and no organization has been more effective at gaining industry-wide

adoption, pervasion, and interoperation of infrared than the Infrared Data Association (IrDA).

Plastic Optical Fiber. Optical fiber offers regulation-free, radiation-free, high-bit-rate communications at the expense of having to be strung around the home and using unfamiliar connectors. The diameter of plastic fiber is fifteen times that of glass fiber (980 μm versus 62.5 μm), so affixing connectors is easy, and 650-nm LED drivers and detectors are inexpensive. Its modal bandwidth of 10 MHz \cdot km* allows whole-house distances at 100-Mbit/s data rates. The cost of the medium and its connectors is comparable to that of data-grade twisted pair.

Radio Frequency (RF). RF communications covers a broad array of technologies, even within the context of the home. The attractions of RF are that it requires no wiring or connectors and that it allows mobility. The challenges are keeping the transmissions out of the neighbors' houses, finding an appropriate operating point in the speed-distance-cost triangle, dealing with the highly-regulated RF environment, defining suitable medium access controls, and keeping the power low enough for its use in small, lightweight, battery-powered appliances.

There is far too much material on RF for us to be anything but brief so here are only the barest of highlights. The most prominent RF technologies are for telephony (cellular and cordless) and these are only now moving from analog to digital and rising in frequency to the 900-, 1800-, and 1900-MHz bands. Their use for data transmission is limited. The highly regulated 1.9-to-2.4-GHz band is used for industrial, scientific, and medical purposes, for emerging 1-to-2-Mbit/s wireless local area networks in enterprises, and for unlicensed personal communications services. Modulation ranges from analog to QPSK digital to direct-sequence and frequency-hopping spread spectrum.

Most viable for the home is the 2.4-GHz band yielding around 1 Mbits/s. The real opportunity for the home lies at 5 GHz, where the European HIPERLAN and U.S. unlicensed national information infrastructure (UNII) allocations have opened up 150- and 300-MHz slots, respectively. Data rates of 10 to 20 Mbits/s per channel, with multiple channels, meet the needs of digital video, Internet communications, and home file and image transfer.

* The product of the number of megahertz times the number of kilometers is (at most) ten. Thus, you could transmit at 10 MHz over a distance of 1 km or at 1 MHz over a distance of 10 km, or anything in between as long as the product does not exceed 10.

The UNII band, especially, offers sufficient bandwidth to allow reasonable modulation efficiency with minimal interference.

All the frequencies mentioned thus far propagate through most home walls without difficulty, allowing whole-house coverage (and perhaps some of your neighbor's). The next frequency band, at 60 GHz, offers enormous bandwidth (5 GHz) but has difficulty propagating through walls. Products are still some years away and component costs are uncertain, especially with GaAs being necessary. Also, it is not clear if 60 GHz is best suited for point-to-point or omnidirectional use. At both 60 and 5 GHz, open questions remain regarding modulation, and especially whether multicarrier modulation can effectively be harnessed to combat narrowband noise or if distributed feedback equalizers are necessary.

Where Are these Technologies Heading?

The in-home communications technologies are heading in whatever direction they can that leads to just enough capability at the lowest possible cost for the applications that will actually be there. None of these factors is known yet so a great deal of investigation is still required. Everyone seems to think that consumers want more bandwidth than they have but no one can say exactly what for. We do not yet know what devices will need to communicate with other devices, what applications and demographics will tolerate the installation of new wiring, or how much money there is to be made in the communications versus the appliances that communicate. We feel that the pushes toward higher speed, mostly digital communications, and decreased reliance on wires are the right ones.

Interconnection Technologies

One other important aspect of residential communications is the means of allowing different devices to actually find and communicate with one another, regardless of where they are located. Thus the picture of residential communications is incomplete without the switches, routers, gateways, and networking software that devices and users need to identify, find, and reach one another.

These technologies exist for public and enterprise networks but their translation to residential use is tenuous for several reasons. One is that they are presently far too expensive for consumers. Another is that they are too complex to set up and use, especially considering that the home has no dedicated information-technology staff. A third is that they do not deliver end-user applications or services but are only a means to that end, so vendors would have to overcome great consumer resistance to purchasing them. Nonetheless, consumers need ways to interconnect their automation network (CEBus?) with their computer (Ethernet?), their digital VCR (IEEE 1394?) to their printer (USB?), and lots of their in-home gear with the outside world.

People (and not just engineers) using these devices will need to know what other devices there are and how to address and name them. The interconnection of the in-home technologies with the access-network technologies is a particularly intriguing problem because the two domains have such different technologies, economics, demographics, and performance characteristics. A body of study has emerged on the *residential gateways* that allow devices within the home to access and exploit the external networks without understanding their technical specifics.

Conclusion

Residential communications offers enormous potential for communications technology development, pervasion, and revenue. The key problems to be solved are matching communications capabilities to the needs of the applications, appliances, and users, driving down costs to consumer levels, and making these systems usable for the nontechnical mass market. Technical innovation is needed in wireline, optical, powerline, satellite, terrestrial microwave, and millimeter-wave technologies and in their transceivers, media, connections, interconnections, configurations, and more. The other papers in this issue describe a wealth of basic-technology work touching most of these areas.

Optical Networks: Backbones for Universal Connectivity

Robert C. Bray

Douglas M. Baney

Communications traffic in the world's fiber-optic backbone network is growing more than 10% per year and the growth rate is accelerating. The ever-increasing bandwidth demands are being met by an array of technological innovations including higher time-division multiplex (TDM) transmission rates combined with wavelength-division multiplex (WDM) overlays.

We are living in a revolutionary age. Communications traffic is growing explosively. In this article we attempt to characterize and quantify this growth, and then comment on technologies to upgrade the usable bandwidth of the fiber-optic backbones of the world's communication networks.

Voice, data, fax, video—these are the forms of electronic communication that are growing at unprecedented rates. A large and increasing proportion of the messages are digital. This is because digital signals can be made practically error-free, and the computers that process them are cheap and getting cheaper and more powerful.

How can we quantify all this? Since all these signals are digital, the universal measure is the bit rate: how many bits per second leave the signal source. The aggregate bit rate of all the sources is the total communications traffic.



Robert C. Bray

A department scientist with HP's Lightwave Division, Bob Bray manages technology and metrology groups. He received his PhD degree in electrical engineering from Stanford University in 1981 and joined HP the same year. He is married, has two sons, and enjoys music and reading.



Douglas M. Baney

Doug Baney is a project manager in the optical communications and measurements department of HP Laboratories. With HP since 1981, he has developed numerous millimeter-wave and lightwave techniques and instrumentation and led a team that developed red, green, and blue doped-fiber lasers. He received his PhD degree in 1990 from the École Nationale Supérieure des Télécommunications in Paris, France. He is married, has three children, and enjoys travel.

A single telephone call is sampled at 8 kbits/s with 8-bit resolution in each direction. The two-way conversation thus represents traffic of 128 kbits/s. One million conversations thus represent traffic of 128 Gbits/s. In one month's time this traffic will have delivered a payload of (30 days) \times (86,400 s/day) \times (128 Gbits/s) = 331,776,000 Gbits = 331.776 Pbits. Typical voice telephone calls last five minutes. Average traffic is less than the daily peak traffic by a factor we will assume to be three.

Twisted pairs, coaxial cables, optical fibers, radio-frequency and microwave wireless, satellites, infrared—these are the media that carry the messages. The special role of single-mode optical fiber is as the high-capacity backbone of global connectivity. No other medium comes close to being able to fill this role.

Network Overview

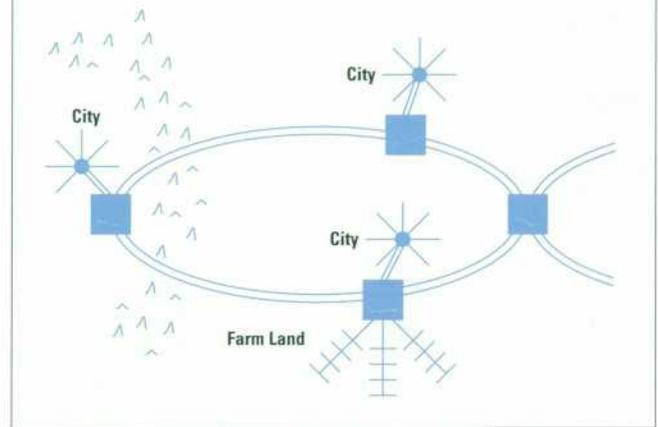
Taking a look at a hypothetical network, we might see various network shapes such as rings linking various sites, as shown in **Figure 1**. Star and tree structures may branch from the network depending on the locality, which could be dense metropolitan, suburban, or rural. Geographic features also influence the local network topology.

The *core network* consists of all trunk lines between major central offices in the larger cities, including trans-oceanic submarine links, festoon links, and long-haul terrestrial links.

Metropolitan networks consist of rings connecting major customers—that is, business locations—in the regions near large cities. An advantage of the ring architecture is that data can be sent in the opposite direction around the link should a break in the fiber cable occur. Both the core network and the metropolitan networks consist of single-mode optical fiber.

The *local access network* consists of the distribution network from a telephone company central office switch to and from customers. These networks are typically a star configuration of twisted-pair lines. The signals are analog until they undergo analog-to-digital conversion, either at the central office or at a remote station in the local loop. At that point multiple calls are multiplexed to a higher bit rate. This description applies to traditional wireline communications, which is supplemented by newer wireless networks. Cellular telephone systems are wireless networks that offer local access to mobile customers.

Figure 1
Typical network.

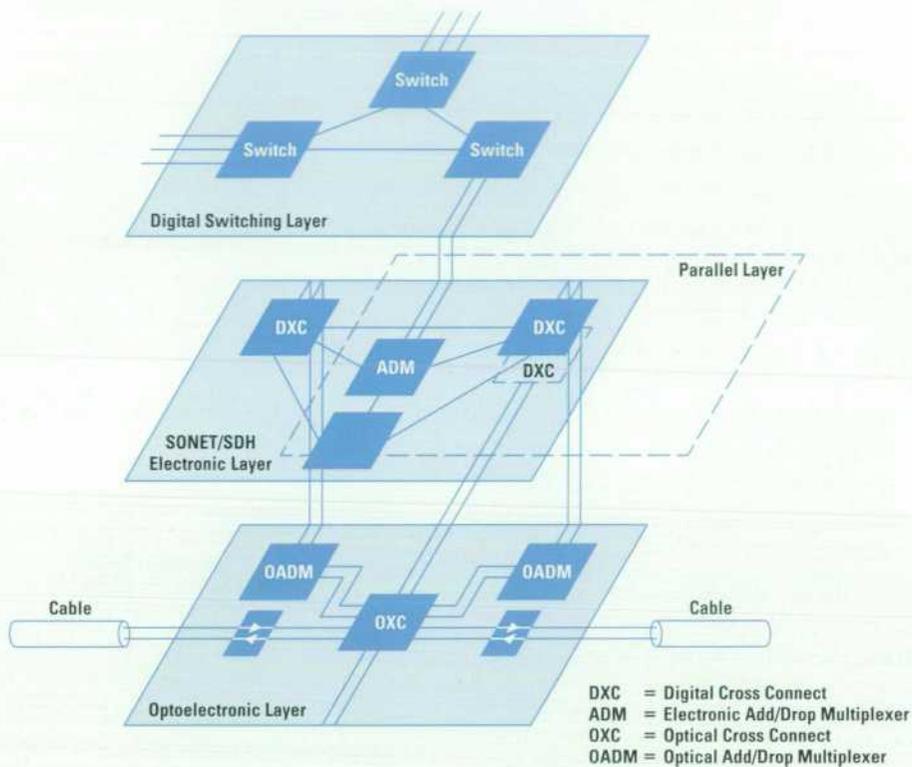


At present many local access networks contain optical fiber lines carrying multiplexed traffic from central offices to local substations, where a local star network of twisted-pair lines branches out to individual customers. Fiber to the home (FTTH) and fiber to the curb (FTTC) networks have had numerous trials but no large-scale deployments. This may change in the near future as one alternative to enable broadband services such as high-speed Internet access to homes.¹

At the network nodes, a number of operations are performed to provide for information routing and multiplexing. A layered model for a network node is shown in **Figure 2**. It provides a useful way of understanding the roles of the various network elements. In the optical layer, incoming signals in the form of modulated light are re-directed to other optical network paths or converted to electrical signals and sent to an electronic layer. In the SONET/SDH layer, digital switches and electronic add/drop multiplexers route and partition the data streams. These data streams are sent to the electronic switching layer where a finer level of partitioning occurs such as ATM (Asynchronous Transfer Mode) switching. Alternatively, the SONET/SDH layer can send data directly back to the optical layer after performing switching at a coarse level. The activities described above are performed on a large scale in the central office and at a smaller level at remote nodes.

Figure 2

Layered model of a network node.



Traffic on the Core Network

Core network telecommunications traffic within the U.S.A. has been estimated to be in the range of 200 to 400 Gbits/s, or the equivalent of 1.5 to 3 million telephone conversations.² This estimate includes the combined voice, data, fax, and video traffic on all U.S. long-distance carriers during midday daily peak traffic. The average annual growth rate of this traffic in 1995 was 10%.

Core telecommunications network traffic outside the U.S.A. has been estimated to be comparable to the U.S. traffic. This is intracountry traffic, which does not cross international boundaries. The exact growth rate is not known, but we believe it to be less than the U.S. growth rate at the present time because of higher rate structures and less developed services outside the U.S.

International telecommunications traffic is well-documented by the International Telecommunication Union (ITU), with a three-year delay.³ As shown in **Table I**, calculations from ITU records estimate that U.S. outgoing

international traffic in 1993 was 2.8 Gbits/s on average, or a bit more than one OC-48 link. Total worldwide international traffic was 11.9 Gbits/s on average, or more than one OC-192 link. These figures include, but are not limited to, the undersea cable traffic.

Table I
International Telecommunications Traffic

Origin	1993 Calls/Day (millions)	1993 Average Traffic (Gbits/s)	1992 to 1993 Growth
U.S.A.	5.3	2.8	11%
Asia	4.6	2.2	21%
Europe	10.5	5.5	12%
World	21.7	11.9	18%

Market Drivers

The worldwide demand for communications drives the need for increased network capacity. Four factors in this capacity growth are:

- Upgrades of existing networks
- Increased global teleconnectivity
- Worldwide deregulation of the telecommunications industry
- Introduction of new broadband services.

Upgrades of Existing Networks. In the U.S.A., all carriers are presently upgrading their terrestrial networks, replacing repeater huts with optical amplifiers.⁴⁻⁷ This provides them with a significant cost savings, as well as a network that is much less data rate dependent. Increased reliability is achieved by configuring the network in a redundant, self-healing ring architecture. In this system, two redundant paths that are physically separated are provided between any two access points. If a fault occurs on one path, for example if an errant backhoe cuts a buried cable, network supervisory software automatically redirects traffic along the alternate path after the fault is detected.

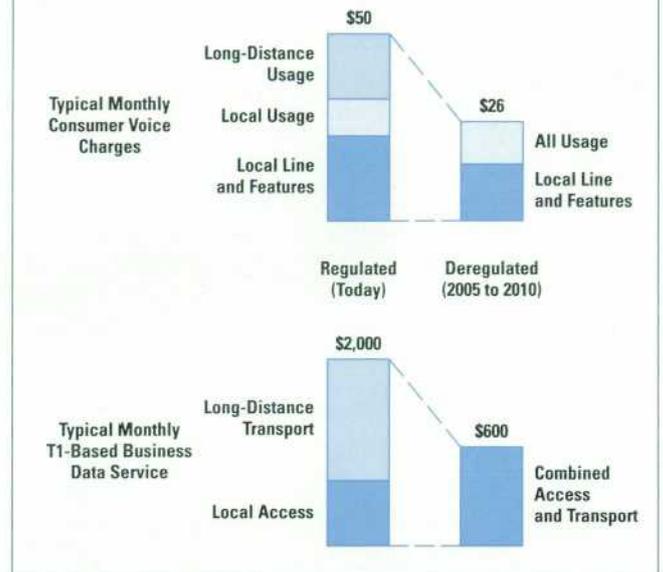
Increased Global Teleconnectivity. Around the world, backbone networks of optical fiber continue to be installed at a remarkable rate. 30-count single-mode optical fiber cable is installed at an average rate of 80 km/hr, around the clock, all year long.⁸ A reported 1.06 million fiber-km were installed in China alone in 1994. China plans to add between 75 million and 100 million new lines to homes and businesses by the year 2000.

So far, U.S.\$11.6 billion has been invested in undersea systems that connect 70 countries, with 17 more countries connected in 1996. In the next three years, it is expected that U.S.\$13.9 billion will be invested in new undersea systems, with less than 10 percent of the cable going into the Atlantic. With completion of the TAT-12/13 and TPC-5 cable systems, the transatlantic and transpacific routes each have installed redundant, self-healing capacity of 10 Gbits/s.^{9,10} This capacity was originally forecast to exceed the demand until 2006. However, by early 1997 the capacity on these cables was completely subscribed. New cables with capacity of 100 Gbits/s are planned.^{11,12}

Deregulation. Deregulation in the telecommunications industry will break down the present barriers to competition in all sectors of the market. Throughout the world,

Figure 3

Price changes before and after deregulation of the telecommunications industry. Courtesy Forrester Research, Inc., Cambridge, MA, U.S.A.

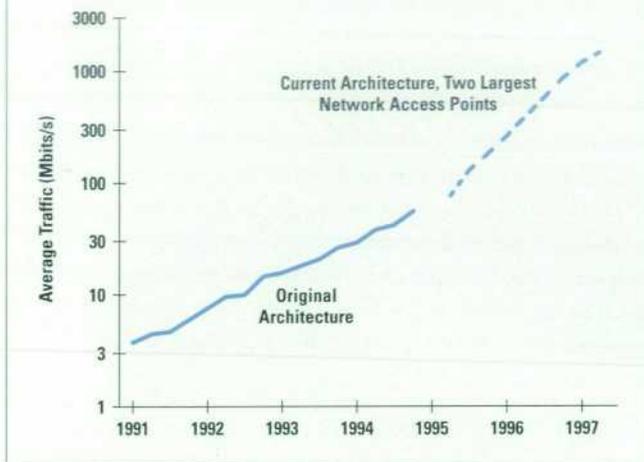


government owned or controlled telecommunications monopolies are undergoing deregulation. The U.S.A. is further along than most countries. The current grouping of service providers in the U.S.A.—regional Bell operating companies (RBOCs), interexchange carriers, and CATV providers—is quickly blurring through mergers, acquisitions, and collaborations between segments. The competitive nature of this market opening will likely stimulate installation of redundant capacity in different providers' networks. This should lead to lower prices (Figure 3) and an increased growth rate for telecommunications traffic. The consulting firm Forrester Research predicts a growth of 250% in traffic by the year 2005.

New Broadband Services. While telecommunications traffic growth rates of 10 to 17 percent are impressive, these growth rates are low compared to those for existing and new data services. The Internet has grown exponentially for at least seven years, as shown in Figure 4. Backbone traffic on the U.S. National Science Foundation's NSFNET exceeded 50 Mbits/s on average in late 1994, having doubled yearly for several years until then. The architecture changed in 1995 to the present network access point structure with multiple backbone services. Traffic is exchanged between services at network access points, the

Figure 4

Internet backbone traffic growth since 1991.¹³



largest of which—MAE-East (Washington, D.C.) and MAE-West (San Jose, California)—have combined traffic that has more than quadrupled in each of the two years since 1995—truly astounding growth! This combined traffic has reached about 1.5 Gbits/s on average in mid-1997, a number that still seems small, that is, the equivalent of about 12,000 telephone calls. The seemingly small traffic volume reflects the fact that the Internet is connectionless. It also reflects the fact that the Internet is insufferably slow! The connectionless nature of the Internet means that packets of data move from source to destination through a network of routers without a circuit being reserved as it would be for telephone service.¹⁴ This is true for an Internet exchange as it traverses the core telecommunications network. However, if that exchange originates in a local access network, a local call to the Internet service provider (ISP) does tie up a circuit through the telephone company's central office switch for the duration of the connection. Local telephone companies have based their service on an economic model that assumes that telephone calls last five minutes on average, and have not yet absorbed all of the calls to Internet service providers that last an hour or more.¹⁵

To predict Internet traffic growth to the year 2000 and beyond, we must recognize the following Internet traffic growth drivers:

- Connections worldwide double yearly.
- PCs will continue to proliferate.

- Faster PCs and LANs (1-GHz and 1-Gbit/s) will be widespread by 2000.
- Demand for faster access (ISDN lines, 50-Mbit/s cable modems) is rising.
- Demand for bandwidth-hungry services (3D graphics, video clips) will grow.
- Network software (Java) will increase network average bit rates.
- Lower telecommunications rates will unleash demand.
- Corporate intranets linking sites of the same firm, enabling efficiencies and savings, will proliferate.

Figure 5 shows the predictions of a model that accounts for the yearly doubling of Internet connections and the further, compounded explosive growth in Internet traffic that is likely to occur as the result of the factors listed above. In this model, the peak daily worldwide core network Internet traffic rises from 360 Mbits/s in 1996 to 110 Gbits/s in 2000.

To see the effect of the Internet traffic, we plotted the projected U.S. telecommunications peak long-haul traffic with and without Internet use in Figure 6. The lower curve shows the projected U.S. daily peak core telecommunications traffic through the year 2000 using a 10% growth rate. The upper curve adds the estimated Internet traffic. The model predicts that the Internet accounts for 10 to 20 percent of the total core telecommunications network traffic by 2000.

Figure 5

Internet traffic growth.

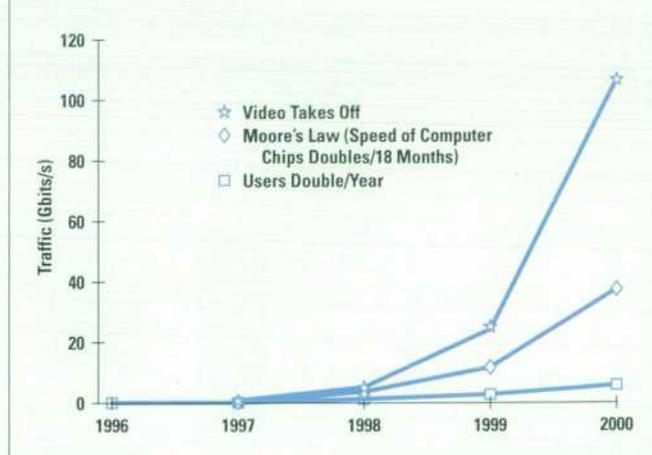
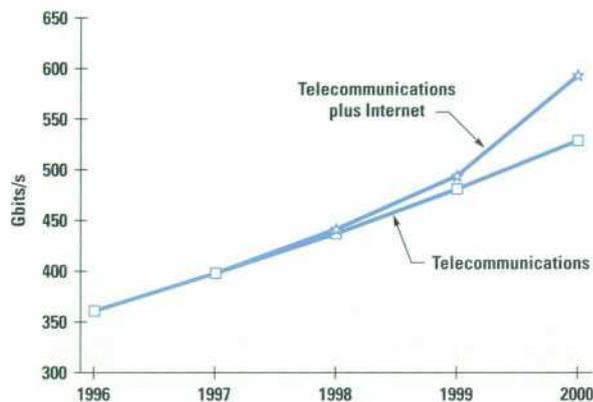


Figure 6

Telecommunications traffic increase from Internet activity.



Technologies to Increase Bandwidth

Single-mode optical fiber has enormous untapped bandwidth. Each nanometer of the spectrum near 1550 nm, where loss is at a minimum, represents 125 GHz of bandwidth.

Worldwide demand for increased network capacity has pushed core network service providers to find the best alternative to increase bandwidth beyond the 2.5-Gbit/s rate on a single optical fiber. The increased demand for bandwidth has forced a reevaluation of network designs and the development of network topologies based on wavelength-division multiplexing (WDM). The challenge for today's network designer is to exploit more of the bandwidth of the fiber.

The way information is physically encoded onto the transport layer involves a number of trade-offs between utilization of existing infrastructures, investment in new technologies, and scalability of the network. Traditionally, a

single optical carrier per fiber has been used in fiber-optic networks. Originally, optical networks were based on direct detection receivers and regeneration at periodic intervals along the core. Later research pointed towards the use of coherent receivers to increase the distances between regenerators.¹⁶ This was quickly put on the back burner with the emergence of the erbium-doped fiber amplifier (EDFA).^{4,7} The EDFA provides low-noise, polarization independent optical gain to overcome propagation losses in the fiber (~ 0.3 dB/km). Initial systems employed purely single-carrier information transport based on time-division multiplexing (TDM) until it became apparent that increased capacity could be economically implemented with a WDM overlay on the existing TDM network structure.

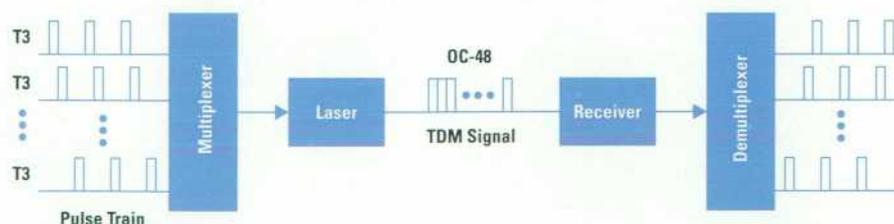
The goal of network designers is to use the right combinations of optical carriers and data rates to maximize the network performance in terms of reliability, cost, and future growth. In the next sections, we will discuss TDM and WDM techniques used to obtain higher data throughput and issues associated with deployment of these techniques.

Time-Division Multiplexing

TDM has been the traditional method for combining information channels. This approach is illustrated in Figure 7. Increased data rates are made possible by interleaving more and more pulses while shrinking the pulse width at the same time. This is similar to packing more cars on a single-lane highway with a fixed speed limit. At a certain point we must shrink the size of the car. Similarly, as the pulse packing increases, the electronics must operate at higher and higher frequencies to accommodate the shorter pulse widths.

Figure 7

Time-division multiplexing (TDM).



Starting from the most fundamental data rate, the voice call, the TDM system combines other calls through interleaving or multiplexing. The 64-kbit/s call, referred to as a DS0, is multiplexed with 23 other calls for a total of 1.544 Mbits/s, which is designated as a T1 (Figure 8). The next level up is the T3, which represents a data rate of 43.008 Mbits/s. Add about 7 Mbits/s for overhead and we have the first optical carrier (OC) designation, OC-1.

With the evolution of standards in various parts of the world, different nomenclatures and rates have been designated. At higher data rates, these standards have merged into the SONET/SDH designations. A comparison of the most common SONET hierarchies is shown in Table II.

Table II
SONET Signal Rates

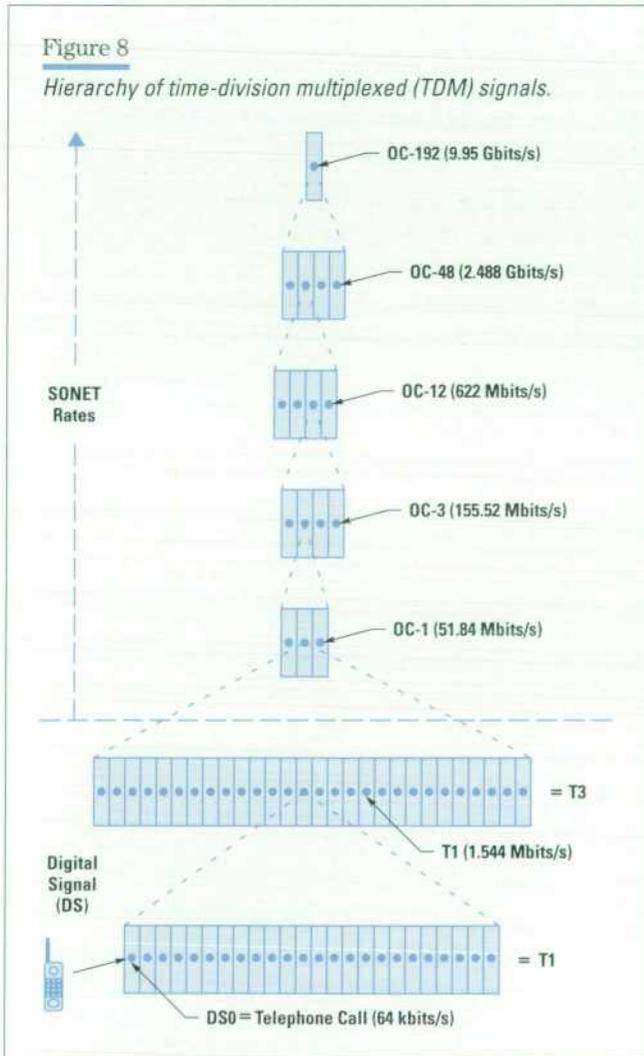
Electrical Signal	Optical Signal	Data Rate (Mbits/s)
STS-1	OC-1	51.84
STS-3	OC-3	155.52
STS-12	OC-12	622.08
STS-48	OC-48	2488
STS-192	OC-192	9953

The lower-rate OCs are multiplexed together to form the OC-48 transmission at a rate of 2.488 Gbits/s. OC-48 rates are practical with today's technologies and OC-192 is in early deployment. Obviously, a great deal of multiplexing is required to achieve these data rates. As the data rate increases, so does the cost of the electronic terminal equipment that converts the optical signals to electrical signals and demultiplexes the data to lower rates for electronic switching and routing. Other challenges also present themselves at high data rates, particularly the dispersion in the optical fiber. Fiber dispersion places limits on the permissible channel bandwidth to keep pulse distortion to an acceptable level.

Fiber Dispersion. The problem of dispersion is tied to the bandwidth that the signal occupies. Chromatic dispersion refers to the frequency dependence of the velocity of light in the optical fiber. The signal bandwidth scales inversely with twice the pulse width for NRZ (nonreturn-to-zero) transmission. The frequency chirp of a directly modulated laser significantly increases the bandwidth occupied by the signal. Chirping refers to the unwanted optical carrier frequency excursions created during intensity modulation.^{17,18} The magnitude of the chirp increases with the pulse rate. The chirped optical pulses propagating along the fiber are spread out by the combined effects of linear and nonlinear dispersion in the single-mode optical fiber. Eventually, it becomes difficult to distinguish between a logical 1 and a logical 0 as the dispersion causes the pulses to overlap in time.

Given that the dispersion of the installed fiber is fixed, we must apply dispersion compensation or reduce the bandwidth occupied by the channel. The bandwidth is reduced by minimizing transmitter chirp and employing different

Figure 8
Hierarchy of time-division multiplexed (TDM) signals.



encoding techniques. Low-chirp electroabsorption and Mach-Zehnder modulators have been developed. Electroabsorption modulators are integrated with DFB (distributed feedback) lasers to eliminate the extra fiber coupling required with external modulators. These devices provide low negative chirp to help compensate for pulse spreading caused by the Kerr effect.¹⁹

Approximately 95% of the world's installed fiber base has its dispersion zero near 1310 nm.⁸ At the low-loss wavelength of 1550 nm, this fiber has a dispersion of approximately 17 ps/nm · km. A simple relationship estimates the maximum span length allowed by linear chromatic dispersion spreading effects on the transmitted signal:

$$L_{\max} = \frac{0.8c}{DB^2\lambda^2}$$

where $c = 2.997925 \times 10^8$ m/s is the velocity of light, D is the dispersion parameter, B is the bandwidth, and λ is the wavelength of the light. This relationship is plotted in **Figure 9** for a wavelength of 1550 nm for standard non-dispersion-shifted fiber as well as for dispersion-shifted fiber with dispersion of ~ 2 ps/nm · km.

Obviously the use of zero-dispersion fiber would increase the maximum TDM rate. However, there are two compelling reasons not to use it. One is that installing new fiber under the ground is quite expensive, and the other is that zero-dispersion fiber is incompatible with present WDM

technology because of four-wave mixing effects.²⁰ An alternative and promising method is to compensate for the fiber dispersion.

Dispersion Compensation. A number of technologies are being investigated to reduce the pulse spreading in TDM links. Dispersion compensating fibers with special core designs are fabricated to yield the required wavelength dependent time delay to compensate for standard telecommunications fiber.²¹ Dispersion compensating fiber lengths on the order of 10 km are required to compensate for the installed-fiber dispersion characteristic for typical spans. Some issues encountered with the dispersion compensating fiber technique include insertion loss, polarization-mode dispersion, and Kerr-effect nonlinearity. Dispersion compensating fibers are often characterized by a figure of merit that describes the ratio of the dispersion to the loss. The loss of the dispersion compensating fiber is approximately 0.1 dB per compensated-span kilometer. Research is being directed towards improving the dispersion compensating fiber performance through careful fiber waveguide design.

Another approach uses the fiber Bragg grating. The phase response versus frequency for the fiber Bragg grating can compensate for fiber dispersion. Both reflection-type and long-period transmission gratings are under consideration. As an example, in **Figure 10** a fiber Bragg grating is used in conjunction with an optical circulator to compensate for fiber dispersion. Its operation is as follows. The circulator passes light in the direction of the arrow from port 1 to port 2. The fiber Bragg grating is a reflective type with a wavelength delay that compensates for

Figure 9

Line limitation resulting from uncompensated chromatic dispersion as a function of bit rate.

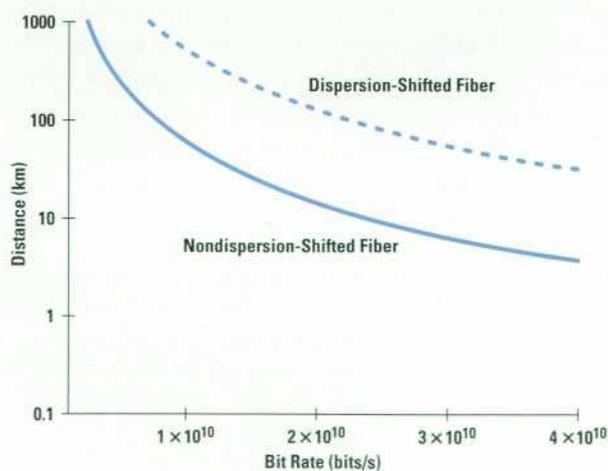
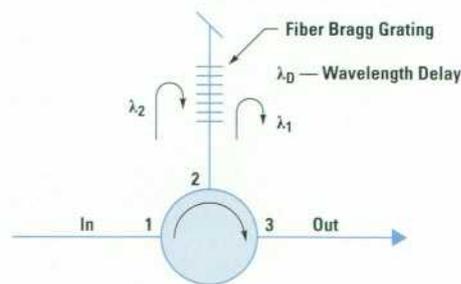


Figure 10

Circulator-based dispersion compensator.



the wavelength dependent delay of the fiber span. The compensated light then goes to the output, port 3. In this way, dispersion compensation is achieved with as little as 2 dB of optical path loss. Research is ongoing to extend the bandwidth over which the dispersion compensation can be employed. Multiple wavelengths can be independently compensated with separate fiber Bragg gratings. In a recent experiment, 6000 ps/nm, or equivalently, compensation for 5000 km of nondispersion-shifted fiber, was demonstrated by a group at the University of Southampton.²²

Dispersion compensation has allowed a significant increase in TDM rates using installed fiber. However, the cost of the terminal equipment including transmitters and receivers increases significantly at the higher bit rates. A way to achieve the increased aggregate bit rate using lower-bandwidth optoelectronics is to add a WDM overlay.

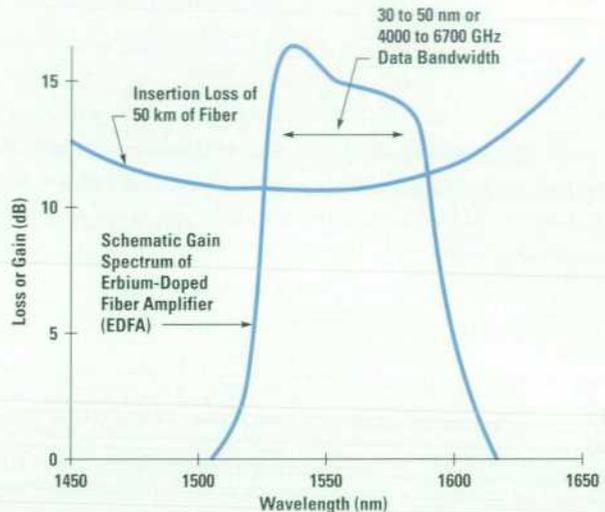
Wavelength-Division Multiplexing

The push towards WDM is being driven by cost considerations. The ability to compensate for chromatic dispersion combined with the expense of high-speed receivers makes WDM an appealing candidate for realizing the required aggregate data rates. While TDM is analogous to packing more cars on a single-lane highway, WDM adds more lanes to the highway.

Wavelength-division multiplexing permits true access to the tremendous bandwidth available from single-mode optical fiber. The available bandwidth is illustrated in **Figure 11**. The fiber itself supports approximately 25,000 GHz of bandwidth. This bandwidth is typically broken up into two bands, centered at 1300 nm and 1550 nm. The 1550-nm band overlaps fortuitously with the gain spectrum of the erbium-doped fiber amplifier (EDFA) as shown in **Figure 11**. In the region around 1550 nm, approximately 5000 GHz of gain bandwidth is available from the erbium-doped fiber amplifier (EDFA). The EDFA is simply a length of optical fiber whose center is doped with erbium ions. Pumping the fiber with a laser operating at 980 nm or 1480 nm causes the ions to absorb the pump energy. Later, the ions give up the energy to the incoming 1550-nm wavelength signal, resulting in amplification. Optical gains of 30 dB and noise figures of ~5 dB are routinely achieved from EDFAs, making them nearly ideal devices for overcoming the propagation loss in optical fibers. For WDM applications, the broad gain bandwidth

Figure 11

Bandwidth availability of installed single-mode optical fiber. The fiber itself supports approximately 25,000 GHz of bandwidth. In the region around 1550 nm, approximately 5000 GHz of gain bandwidth is available from the erbium-doped fiber amplifier (EDFA).

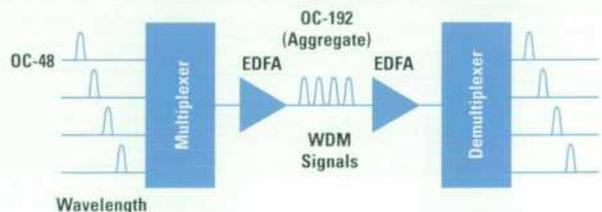


of the EDFA makes it capable of simultaneous amplification of many wavelength channels. This is a cost-effective use of the EDFA.

WDM terminals consist of multiple, independent TDM transmitters at different wavelengths in the 1550-nm band and an equal number of independent TDM receivers, as shown in **Figure 12**. The outputs of the transmitters are optically multiplexed together onto a common output interface to the optical transmission network. The optical input to the WDM terminal is optically demultiplexed and the separated signals are fed to TDM receivers.

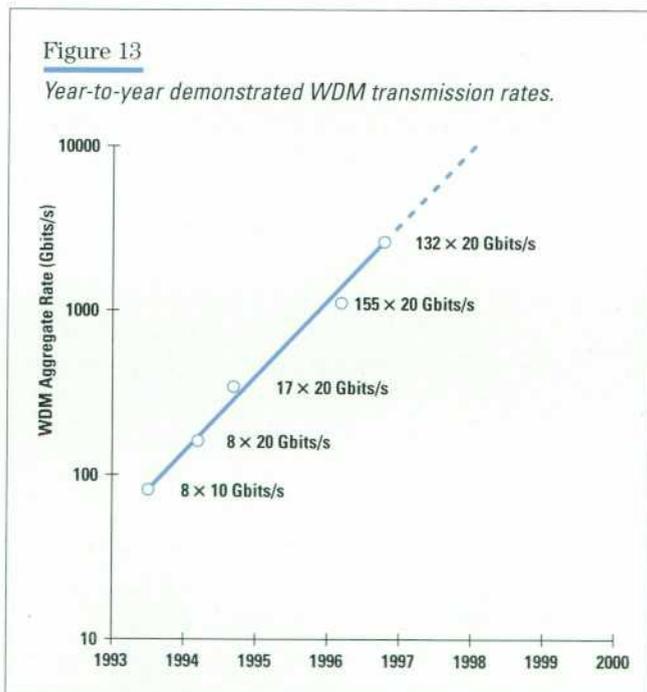
Figure 12

Wavelength-division multiplexing (WDM).



WDM can be considered as a way of summing together many TDM channels on a single fiber. The alternative solution would be to lay more fiber, which is costly. Increasing the aggregate transmission rate using WDM is technically viable and cost-competitive. Increasing TDM rates from say, 10 Gbits/s to 40 Gbits/s is very challenging by today's technical standards. On the other hand, using four WDM channels at 10 Gbits/s to achieve an aggregate rate of 40 Gbits/s is feasible.

WDM is making the "tera era" a reality. Laboratories have already demonstrated transmission capacity beyond 1 Tbits/s on a single fiber.^{23,24} In a recent laboratory experiment, 132 lasers were combined onto a single fiber to transmit information at a 2.6-terabit-per-second rate.²⁵ The tremendous progress in transmission capacity of a single fiber is shown in **Figure 13**, which shows laboratory WDM achievements over the last couple of years. Going beyond 6 Tbits/s on a single fiber will probably require adding another wavelength band such as 1300 nm, or a significant widening of the 1550-nm band.



WDM Network Elements. To make WDM systems possible, a number of new components are required to perform optical amplification, wavelength multiplexing, and routing. For example, wavelengths can be separated from the data stream using optical add/drop multiplexers (OADMs). OADMs and optical crossconnect switches are also envisioned for future wavelength-routed networks.

A key element in WDM networks is the EDFA, which provides nearly wavelength-transparent optical gain.^{4,7} As a result, the EDFA gain spectrum has determined the wavelength band of WDM systems. The EDFAs are placed along the links at varying intervals depending on the data rates, system quality, and other factors. A separation of 80 km is very reasonable. It is important that the EDFA gain response be fairly constant across the wavelength span of interest. If it is not, channel-by-channel power compensation is required in long-haul networks.

The 1300-nm telecommunications window is also available, since the fiber loss is still quite acceptable there. At this wavelength, however, optical amplifiers have not been able to achieve the spectacular results obtained by EDFAs at 1550 nm. Some candidates include gain-clamped polarization independent semiconductor amplifiers, praseodymium or neodymium fiber amplifiers, or Raman-effect fiber amplifiers.²⁶

While it is beyond the scope of this paper to discuss each element of the WDM network in detail, the summary in **Table III** provides a partial list of the network elements available to system designers and the technologies involved.^{27,28}

The use of specific technologies in realizing certain network elements depends on the requirements for network transparency. Ideally, an all-optical network could provide data-rate-transparent operation, allowing easy upgrade to higher data rates. In such a network, wavelength conversion would have to be performed with all-optical methods. Optoelectronic regeneration has limits on bandwidth set by the detection and regeneration circuitry and would not be strictly transparent. Discussion continues on the need for transparency in view of the difficulties of coherent cross talk and vendor interoperability associated with transparent network designs.

Table III*WDM Network Elements*

Device	Abbreviation	Function	Technologies
Erbium-doped fiber amplifier	EDFA	Provides flat gain spectrum to WDM channels	Silica-based or fluoride-based fibers, laser pumps
Multiplexer/demultiplexer	MUX/DMUX	Combines/separates multiple wavelength channels onto/from a single fiber	Waveguide arrays, zig-zag filters, interference filters, diffraction gratings, fiber gratings, fuse couplers
Wavelength add/drop multiplexer	WADM	Adds or drops one or more wavelength channels without terminating the entire layer	Fused couplers, interference filters, circulators and fiber Bragg gratings, Mach-Zehnder interferometers
Wavelength interchange crossconnect	WIXC	Crossconnects signals with allowance for wavelength interchange	Optoelectronic regeneration, cross-gain modulation, optical nonlinearity, mechanical
Wavelength-selective crossconnect	WSXC	Crossconnects individual wavelengths without wavelength interchange	Fused couplers, interference filters, circulators and fiber Bragg gratings, Mach-Zehnder interferometers
Optical crossconnect	OXC	Optical signal switching	Electromechanical, electrooptic

WDM Challenges. Other challenges facing WDM networks relate to optical frequency standards and network architectures. Setting wavelength and frequency standards is challenging in a competitive environment. The stability requirement on laser center wavelengths is stringent with narrow channel separations. It is further complicated by physical phenomena such as four-wave mixing, which discourages the use of uniformly spaced channel frequencies. Currently, the ITU-T (International Telecommunications Union, under United Nations charter) has allocated wavelength channels on a frequency grid with 100-GHz spacings referenced to 193.1 THz. This is helpful, but noting the approximately 5000 GHz of available bandwidth, the possibilities are still limited that vendors will choose the same points on the grid for their 16-channel systems. The ITU-T optical frequencies and their respective wavelengths from 1530 nm to 1560 nm are shown **Table IV**.

Four-wave mixing, in which the nonlinear behavior of the optical fiber causes different channels to mix, causes problems by scattering signal power to other wavelengths.²⁰

Low-dispersion fiber, while helpful for TDM, increases the efficiency of the undesirable four-wave mixing. Therefore, WDM links must be designed to have local dispersion to reduce four-wave mixing but achieve low global dispersion. Dispersion maps showing the dispersion with distance along the link are commonly used as part of the link design.²⁹

The complexity of network management for WDM systems can vary tremendously depending on the architecture. For point-to-point WDM systems, the management is not unlike a network that simply added more fibers along a span instead of using WDM. In transparent networks, on the other hand, signals pass without bandwidth limitations, which implies no optoelectronic regeneration. Thus, transparent networks require wavelength interchange and switching capability allowing wavelength reuse. The management of transparent networks is considerably more complex and is still in active investigation.

Table IV*ITU-T Standard Optical Frequencies*

Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)
192.1	1560.61	194.1	1544.53
192.2	1559.79	194.2	1543.73
192.3	1558.98	194.3	1542.94
192.4	1558.17	194.4	1542.14
192.5	1557.36	194.5	1541.35
192.6	1556.55	194.6	1540.56
192.7	1555.75	194.7	1539.77
192.8	1554.94	194.8	1538.98
192.9	1554.13	194.9	1538.19
193.0	1553.33	195.0	1537.40
193.1	1552.52	195.1	1536.61
193.2	1551.72	195.2	1535.82
193.3	1550.92	195.3	1535.04
193.4	1550.12	195.4	1534.25
193.5	1549.32	195.5	1533.47
193.6	1548.51	195.6	1532.68
193.7	1547.72	195.7	1531.90
193.8	1546.92	195.8	1531.12
193.9	1546.12	195.9	1530.33
194.0	1545.32	196.0	1529.55

Conclusion

Communications traffic in the world's fiber-optic backbone network is growing more than 10% per year. The growth rate shows promise of accelerating further. Ever-increasing communications bandwidth demands prompted by voice, fax, video, and Internet activity are being met by an array of technological innovations. Higher TDM transmission rates combined with WDM overlays have provided an economical alternative to putting more fiber in the ground. Development of standards for WDM networks will allow the possibility of interoperability between equipment from different manufacturers.

Acknowledgments

The authors would like to acknowledge the following people who contributed useful discussions: Tim Bagwell, Jerry Chappell, Waguih Ishak, Roger Jungerman, Chris Miller, Steve Newton, and Karl Shubert.

References

1. N.J. Frigo, "A Survey of Fiber Optics in Local Access Architectures," in I.P. Kaminow and T.L. Koch, editors, *Optical Fiber Telecommunications IIIA*, Academic Press, 1997, pp.461-522.
2. L.G. Roberts, "ATM Overview," *Proceedings of COMPCON '94*, IEEE Computer Society Press, 1994, pp. 94-104.
3. *ITU Statistical Yearbook 1994*, ITU, 1995.
4. R.J. Mears, L. Reekie, I.M. Jauncey, and D.N. Payne, "High-gain rare-earth-doped fiber amplifier at 1.54 μm ," *OFC/IOOC '87 Technical Digest*, 1987, paper W12, p. 167.
5. T. Li, "The impact of optical amplifiers on long-distance light-wave telecommunications," *Proceedings of the IEEE*, Vol. 8, no. 11, November 1993, pp. 1568-1579.
6. C.R. Giles and E. Desurvire, "Modeling erbium-doped fiber amplifiers," *Journal of Lightwave Technology*, Vol.9, no.2, February 1991.
7. J. Stimple, "Testing Erbium-Doped Fiber Amplifiers," *Hewlett-Packard Journal*, this issue, p. 120.
8. R. Mack, KMI Corp., Providence, RI, private communication, October 1996.
9. P. Trischitta, M. Colas, M. Green, G. Wuzniack, and J.Arena, "The TAT-12/13 Cable Network," *IEEE Communications Magazine*, Vol. 34, no. 2, 1996, pp. 24-28.
10. W.C. Barnett, H. Takahira, J.C. Baroni, and Y. Ogi, "The TPC-5 Cable Network," *IEEE Communications Magazine*, Vol. 34, no. 2, 1996, pp. 36-40.
11. L. Dadouris, M. Singhi, and S. Long, "The Impact of the Internet and Broadband Service Offerings on Submarine System Cable Capacity into the 21st Century," *SubOptic '97 Conference Proceedings*, 1997, pp. 48-54.
12. I.-M. Beaufils and J.-L. Chabert, "Interconnection of Future Submarine Cables in the Global Communications Web," *SubOptic '97 Conference Proceedings*, 1997, pp. 91-97.
13. See <http://www.merit.edu/nsfnet/statistics/history.bytes> and <http://www.mfst.com/MAE>
14. C. Partridge, *Gigabit Networking*, Addison-Wesley, 1993, pp. 225-251.

15. B. Ziegler, "Slow Crawl on the Internet," *Wall Street Journal*, August 23, 1996, p. B1.
16. R.A. Linke and A.H. Gnauck, "High-capacity coherent light-wave systems," *Journal of Lightwave Technology*, Vol. 6, no. 11, November 1988, pp.1750-1769.
17. T.L. Koch and J.E. Bowers, "Nature of wavelength chirping in directly modulated semiconductor lasers," *Electronics Letters*, Vol. 20, December 1984, pp. 1038-1039.
18. R.A. Linke, "Modulation induced transient chirping in single frequency lasers," *IEEE Journal of Quantum Electronics*, Vol. QE-21, 1985, pp. 593-597.
19. J.A.J. Fells, M.A. Gibbon, I.H. White, G.H.B. Thompson, R.V. Penty, C.J. Armistead, E.M. Kimber, D.J. Moule, and E.J. Thrush, "Transmission beyond the dispersion limit using negative chirp electroabsorption modulator," *Electronics Letters*, Vol. 30, July 1994, pp. 1168-1169.
20. A.R. Chraplyvy, "Limitations on lightwave communications imposed by optical-fiber nonlinearities," *Journal of Lightwave Technology*, Vol. 8, no. 10, October 1990.
21. Y. Liu, J. Antos, and M.A. Newhouse, "Large effective area dispersion-shifted fibers with dual-ring index profiles," *OFC '96 Technical Digest*, 1996, paper WK15.
22. R.I. Laming, W.H. Loh, X. Gu, M.N. Zervas, M.J. Cole, and A.D. Ellis, "Dispersion compensation with chirped fiber Bragg grating to 400 km at 10 Gbits/s in nondispersion-shifted fiber," *OFC '96 Technical Digest*, 1996, paper ThA5.
23. H. Onaka, H. Miyata, G. Ishikawa, K. Otsuka, H. Ooi, Y. Kai, S. Kinoshita, M. Seino, H. Nishimoto, and T. Chikama, "1.1 Tbits/s WDM transmission over 150 km 1.3 μ m zero-dispersion single-mode fiber," *OFC '96*, 1996, PD19.
24. A.H. Gnauck, A.R. Chraplyvy, R.W. Tkach, J.L. Zyskind, J.W. Sulhoff, A.J. Lucero, Y. Sun, R.M. Jopson, F. Forghieri, R.M. Derosier, C. Wolf, and A.R. McCormick, "1 Terabit/s transmission experiment," *OFC '96*, 1996, PD20.
25. Y. Yano, T. Ono, K. Fukuchi, T. Ito, H. Yamazaki, M. Yamaguchi, and K. Emura, "2.6 terabit/s WDM transmission experiment using optical duobinary coding," *22nd European Conference on Optical Communications*, 1996, paper ThB.3.1.
26. Y. Aoki, "Properties of fiber Raman amplifiers and their applicability to digital optical communication systems," *Journal of Lightwave Technology*, Vol. 6, no. 7, July 1988.
27. R.E. Wagner, R.C. Alferness, A.A.M. Saleh, and M.S. Goodman, "MONET: multiwavelength optical networking," *Journal of Lightwave Technology*, Vol. 14, no. 6, 1996, pp. 1349-1355.
28. S. Johansson, "Transport network involving a reconfigurable WDM network layer—A European demonstration," *Journal of Lightwave Technology*, Vol. 14, no. 6, 1996, pp. 1341-1348.
29. A.R. Chraplyvy, A.H. Gnauck, R.W. Tkach, and R.M. Derosier, "8x10-Gbit/s transmission through 280 km of dispersion-managed fiber," *IEEE Photonics Technology Letters*, Vol. 5, no. 10, October 1993.

Data Transmission Schemes for Higher-Speed IEEE 802 LANs Using Twisted-Pair Copper Cabling

Steven G. Methley

Alistair N. Coles

Eric Deliot

Transmission at 424.8 Mbits/s using Category 5 cable can meet both industrial and the more stringent domestic emissions regulations. The design is robust in operation and the complexity is not much greater than that used for the 100-Mbit/s rate.

In October 1995, two new 100-Mbit/s local area network standards were published by the Institute of Electrical and Electronics Engineers: IEEE 802.12, based on a demand priority (DP) access method, and IEEE 802.3u, based on a collision sense multiple access/collision detection (CSMA/CD) access method. Subsequently, there has been much interest in increasing the operating speed of these standards beyond 100 Mbits/s.¹ This imposes design challenges for the two media involved: optical fiber and copper cabling. Optical-fiber-based approaches are discussed elsewhere in this issue. Here we will examine schemes that are designed to use existing copper cable installations, specifically data-grade cable, Category 5. This cable is already installed in locations that have followed building wiring standards.

A critical objective for the copper solution was to cost less than the fiber solution and this meant low complexity was required. We show that transmission at 424.8 Mbits/s using Category 5 cable can meet both industrial and the more stringent domestic emissions regulations. Furthermore, the design is robust in operation and of a complexity not greatly in excess of that used for the 100-Mbit/s rate.² The data rate of 424.8 Mbits/s is equivalent to the Fibre Channel rate of 531 Mbits/s before 8B10B coding (mapping 8 bits to 10 bits) and was chosen in anticipation of other IEEE 802 physical layers (PHYs) also following the route of compatibility with Fibre Channel speeds to leverage existing components.

Compared with fiber, a copper system has a number of problems peculiar to it. A long metallic conductor is prone to act as a radio antenna, and this could lead to interference with other equipment (emissions) and unwanted pickup from other equipment (susceptibility). In addition, the copper cables are isolated by transformer coupling to avoid ground loops and other undesirable dc effects. These properties together with the transfer characteristic of the cable determine what can be successfully transmitted over the cable in a real-world environment.

Architectural Requirements: Speed, Bidirectionality

In a shared-medium access method such as DP or CSMA/CD, full-duplex data transmission is not possible, since only one station (or none) has access to the shared channel at any point in time. However, half-duplex data transmission is suitable and an important refinement is possible. Network control traffic, including, for example, requests for access to the shared media, can be allowed to travel upstream when data flow is downstream, or vice versa. This helps the efficiency of the network, making it into a hybrid duplex scheme in which data and control can flow simultaneously in opposite directions, but neither is full-duplex.³

Having four pairs in one Category 5 cable means that there are alternative duplexing schemes to the traditional single-channel frequency-division multiplexing (FDM) or hybrid-plus-echo-canceller approach. The bandwidth of one pair can be dedicated to a reverse control-signaling channel with the remaining three pairs for the forward data channel. This 3 + 1 scheme creates an asymmetric duplex scheme in terms of the bandwidth available in each direction. Having the still relatively high bandwidth of a single pair for control signaling is useful not because control traffic is high-bandwidth but because prompt detection of control codes is advantageous in terms of network performance. Asymmetric duplexing using selected pairs is notably less complex than FDM or hybrid plus echo canceller because of the lower component count. As shown in **Figure 1**, only two of the four pairs need be half-duplex; the other two can be simplex. Near-end crosstalk (NEXT) is no longer the dominant noise source in this Category 5 system as it was in the Category 3 100-Mbit/s systems.⁵ Externally induced noise is dominant. Having a solution without FDM, echo cancellation, or NEXT cancellation is the pivotal step in forming a low-complexity system design.

Speeds greater than 100 Mbits/s are of interest for extending the existing standards, and in particular, speeds matching the Fibre Channel rates offer the possibility of leveraging existing components including drivers and clock recovery circuits. The first two Fibre Channel rates offering a marked speed increase over 100 Mbits/s are 531 Mbits/s and 1062.5 Mbits/s. However, these include the overhead of an 8B10B code designed assuming a single serial transmission medium. Since the copper solution divides the data among three pairs, the 8B10B code could be replaced with something more appropriate for this application. The raw data rates for the two Fibre Channel rates then become 424.8 Mbits/s and 849.6 Mbits/s. Extending these rates to the 3 + 1 asymmetric duplex scheme gives per-pair rates of 141.6 Mbits/s and 283.3 Mbits/s.

Signaling: Multilevel Signaling, Coding, Control

Earlier work⁶ had shown that transmitting basic NRZ data at 155 Mbits/s was unlikely to satisfy domestic emissions regulations (e.g., FCC B) and might even prove problematical in meeting the less stringent industrial regulations (e.g., FCC A). Regulations begin at 30 MHz, so reducing transmitter spectral energy above this frequency is an obvious approach to reducing emissions. Thus, a per-pair

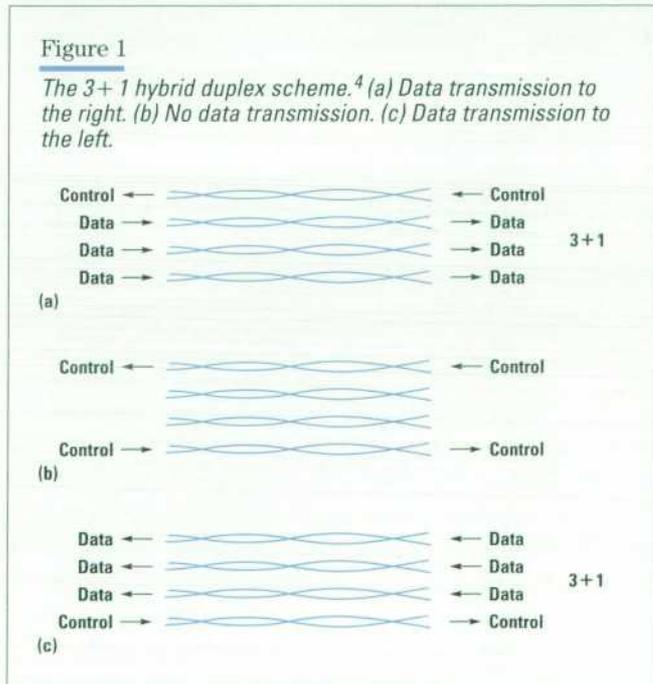
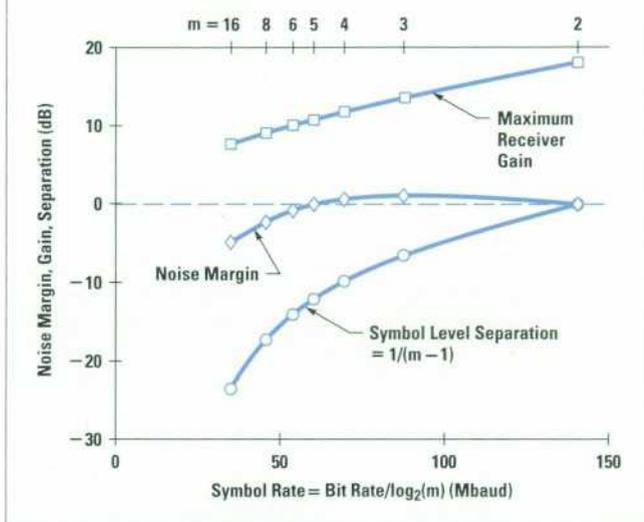


Figure 2

PAM noise margin as a function of PAM index (m) for a data rate of 141.6 Mbits/s.⁸



rate of 283.3 Mbits/s immediately seems far less suitable for a copper implementation than 141.6 Mbits/s. Several bandwidth compressing modulation schemes were studied, including quadrature amplitude modulation (QAM), partial response (PR) classes I and IV, and pulse amplitude modulation (PAM).^{7,8} QAM is a two-dimensional scheme requiring complex in-phase and quadrature filters, and while PR has good bandwidth compression, this comes at the expense of clock recovery. PAM requires dc balancing but is a relatively straightforward scheme to implement and when m levels are used reduces bandwidth requirements by $\log_2(m)$. If excess bandwidth α is also reduced below 100% then an extra factor $(1 + \alpha)/2$ is gained to give the overall relationship:

$$\text{Transmitter bandwidth} = \frac{\text{bit rate}}{\log_2(m)} (0.5 + \alpha/2).$$

Using the measured external noise level (discussed later), noise margin was calculated for PAM systems from 2 to 16 levels with 80% excess bandwidth, 100-m Category 5 worst-case attenuation, and a per-pair bit rate of 141.6 Mbits/s. **Figure 2** shows the noise margin plotted against baud rate, or equivalently, PAM index at the given data rate. For these calculations, equal peak transmitter voltage was assumed regardless of number of levels. Also shown are the maximum receiver gain and the symbol separation that were used to calculate the margin. A comparison of

the transmitter spectra for 2-level, 4-level, and 8-level PAM are shown in **Figure 3**, where a fair comparison has been introduced by adjusting the transmitter output voltage to yield the same noise margin at the receiver in each case. The start of emissions regulations is marked by the dotted line at 30 MHz. Clearly, 8-level PAM has lower energy above this frequency. To determine if the transmitters will actually pass the regulations, EMC measurements are required, which must also include the cable.

The absence of a dc response is a problem for a baseband system such as PAM, but can be solved by the addition of a balancing block code having the property of reducing the running digital sum (RDS, the sum of transmitted symbol levels), which reduces low-frequency components. A block code that has more codewords than data words—that is, redundancy—allows control signaling to be readily incorporated into the transmitted symbol stream. Block coding can also ensure a high transition density, which aids clock recovery schemes.

Table I shows five possible block codes.⁴ The best all-around codes are 8B3N (mapping 8 bits to 3 nonary or 9-level symbols) and 16B6O (mapping 16 bits to 6 octary or 8-level symbols) since, in addition to good bandwidth compression and redundancy, the inputs are multiples of eight bits, which can lead to implementation convenience. Additionally, in the 8B3N case, the RDS at a codeword boundary is constrained to be between -4 and 5 . Choosing a subset of the 8B3N code for data allows inband control signaling, and if the subset is suitably constrained, the

Figure 3

2-level, 4-level, and 8-level PAM transmitter spectra for equal noise margin.⁸ Only the first lobes of the spectra are shown.

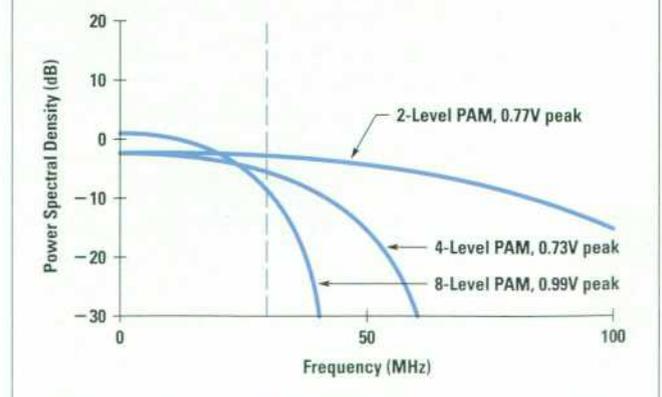


Table I*Properties of Possible Block Code Options⁴*

Code	Output Levels	Bandwidth Compression Factor	Redundancy Factor	Baud Rate (Mbaud) for Bit Rate of 141.6 Mbits/s	Comments
3B1O	8	3	1	47.2	Cannot balance
8B3O	8	2.67	2	53.1	Byte aligned
10B4O	8	2.5	4	56.6	
8B3N	8	2.67	2.85	53.1	Byte aligned
16B6O	8	2.67	4	53.1	Word aligned

* Redundancy factor = number of output permutations divided by number of input permutations.

spectral magnitude of the control codes does not greatly exceed that of the data (see **Figure 4**). Finally, by ensuring that the control codewords have a Hamming distance of two with respect to data codewords, then single error events will always be detected in the control codewords. When 8B3N is encoded, the per-pair rate of 141.6 Mbits/s equates to a symbol rate of 53.1 Mbaud.

Electromagnetic Compatibility (EMC)

Electromagnetic compatibility was one of the first areas to be investigated and measurements were performed before the 8B3N code had been chosen for the system. To evaluate the general effectiveness of PAM bandwidth compression with respect to emissions, uncoded 8-level PAM was used for the measurements at 50 Mbaud.^{10,11}

European (EN) and U.S. (FCC) emissions regulations limit the radiation level that the system can be allowed to generate (see **Figure 5**). The industrial Class A levels are 10 dB less stringent than the domestic Class B levels shown. 100VG-AnyLAN⁵ used all four pairs of the cable to good effect by sending only 30 Mbaud on each pair, hence staying below the emissions limits. However, at higher rates this is not enough; even using all four pairs for data, the per-pair rate is 125 Mbits/s for 500 Mbits/s total, and using the 3 + 1 scheme, 141.6 Mbits/s is required. This would require a bandwidth in excess of 30 MHz using 100VG-AnyLAN signaling. Bearing in mind that some balancing overhead would be required, 155 Mbits/s NRZ and 50 Mbaud 8-level PAM were convenient rates to be investigated in terms of emissions.

EMC measurements were made to ascertain the levels of radiation and susceptibility of the Category 5 cable so that the results could be used to examine the trade-off between bandwidth improvement and increased noise susceptibility when moving to a multilevel system. A 100-m length of Category 5 cable including short patch cords and a punch-down block (connector block) was tested with 150-Mbit/s and 50-Mbit/s binary data and a prototype 50-Mbaud, 8-level PAM source. Transmitter output level was 1V peak-to-peak into 50 ohms and 1.4V peak-to-peak when matched to the 100-ohm Category 5 transmission line using a balun. In each case the free cable was wound noninductively on a 1-m-by-1-m wooden frame. Emissions were measured at 3 meters with each transmitter and susceptibility was measured when the cable was subjected to a standard 3V/m field. The field strength was leveled at each of 11 spot frequencies using a pair of optically coupled field

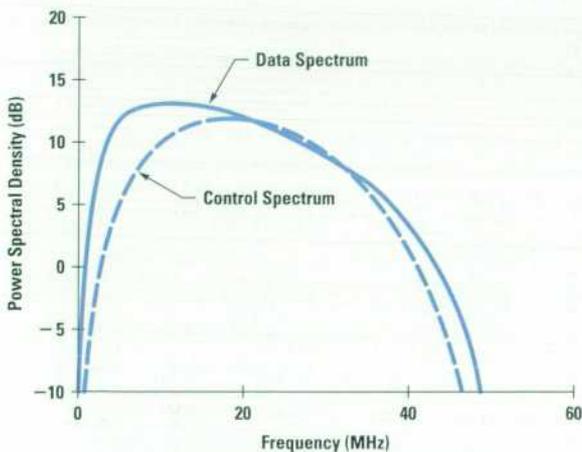
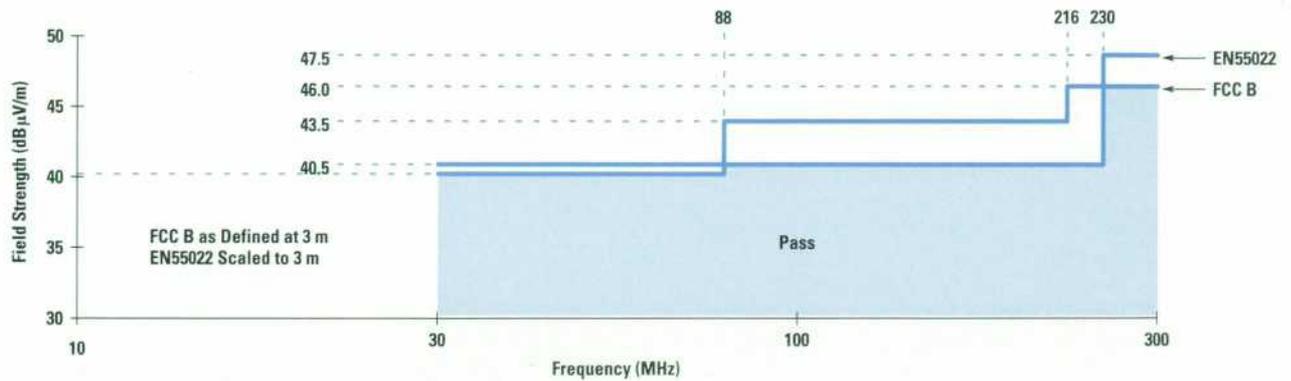
Figure 4*Spectra of 8B3N data and control signaling.⁹*

Figure 5

FCC and EN emissions regulations.⁶



probes before each measurement. The results are shown in **Figure 6**.

Note that externally induced noise is the major noise source in the Category 5 half-duplex system under consideration, principally due to the low level of cross talk. **Figure 7** shows that the 155-Mbit/s case breaks the FCC B emissions mask (barely), whereas **Figure 8a** shows that the 50-Mbit/s case does not. The emissions mask has been adjusted to account for the EMC measurement site attenuation and is thus no longer flat with frequency.

The prototype 50-Mbaud trace (**Figure 8b**) shows some clock breakthrough, but otherwise meets the mask. The

breakthrough is a result of prototype construction rather than a feature of the code. Development of the prototype would enable the spectrum to approach the 50-Mbit/s trace from the test instrument (**Figure 8a**), which is shown for comparison purposes. The 155-Mbit/s EMC data was within 1 dB of emissions measurements made a few years earlier for presentation to the ATM Forum.⁶

The Channel: Cable and Transformers

The measured frequency and phase responses of a sample 100-m Category 5 channel and transformers are shown in **Figures 9a and 9b**. Note the increasing attenuation with frequency, the absence of a dc response, and the nearly linear phase slope. The phase slope equates to a time delay of 0.486 μ s which, as expected, is equal to the transit time of 100 meters of cable assuming a propagation velocity of 0.7c. Oversampled FIR (finite impulse response) models of the measured and EIA/TIA 568 worst-case channels were generated for simulation purposes, and the measured

Figure 6

Measured induced noise levels in a 3V/m field.¹²

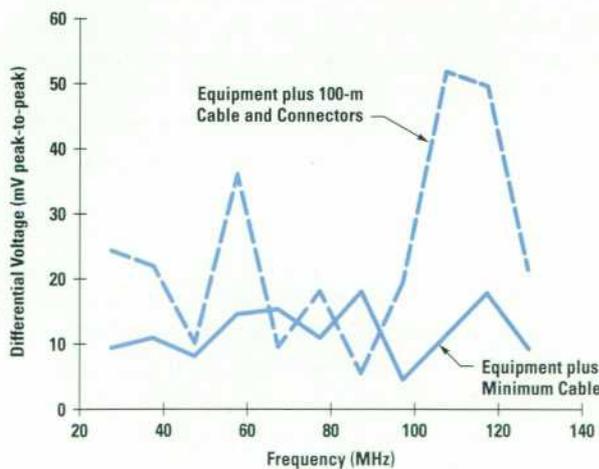


Figure 7

Measured 155-Mbit/s emissions with FCC B limit overlay.⁸

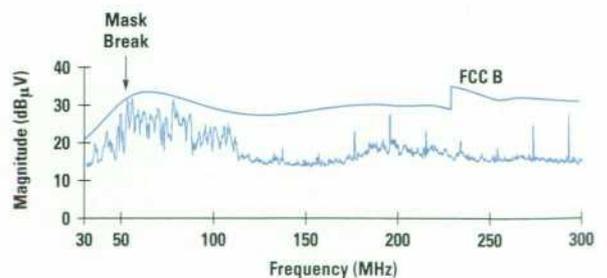
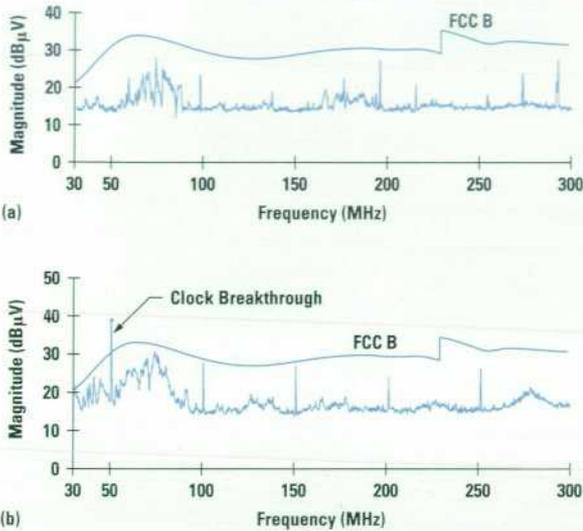


Figure 8

Measured emissions with FCC B limit overlay.⁸ (a) 50-Mbit/s NRZ. (b) 50-Mbaud 8-level PAM.



case is shown in **Figure 9c**. The rather long delay to the punctual tap of this impulse response is 192 taps \times 1/400 MHz = 0.48 μ s, equal to the cable delay as before. The decay after the punctual tap is the result of the low-pass nature of the cable and the very small final tail, which is offset below zero, is a result of the low-frequency corner introduced by the transformers.

Simulation

Figure 10 shows the simulated data transmission path. The functional blocks were written using the MATLAB package. Transmitter block coding was 16B6O or 8B3N, finite precision was used for the signal conversion blocks, and the receiver was configured to use a reference clock. Both T-spaced (baud-spaced or synchronous) and T/2-spaced FIR filter equalizers were used, using the gradient LMS (least mean squares) algorithm in reference-directed mode for convergence and including finite-precision effects.

A T/2 fractionally spaced equalizer was preferred for its insensitivity to timing phase (due to the absence of aliased band-edge components). If a lower sampling rate had been preferred, a T-spaced equalizer could have been used at the expense of slightly more clock recovery complexity. The simulation was used to study the performance of the

block codes for varying degrees of transformer low-frequency loss.

The 3+1 scheme was used with a bit rate of 141.6 Mbits/s on each pair, giving a bandwidth over the three data pairs of 424.8 Mbits/s, which corresponds to the 531-Mbit/s Fibre Channel rate without 8B10B encoding as described above. Mapping 16 bits to 6 octary (8-level) symbols or mapping 8 bits to 3 nonary (9-level) symbols results in the same bandwidth compression factor (2.65) and yields a symbol rate of $T = 53.1$ Mbaud in each case. Transmit filtering was root raised cosine with $\alpha = 0.8$, followed by a fifth-order analog Butterworth filter with a cutoff at 40 MHz. The channel was simulated by the cascade of transformers modeled as first-order high-pass filters with a cutoff at 100 kHz and cable modeled as an FIR filter

Figure 9

Measured (a) frequency and (b) phase response of 100-m typical Category 5 cable to 200 MHz. (c) 305-tap FIR filter model.

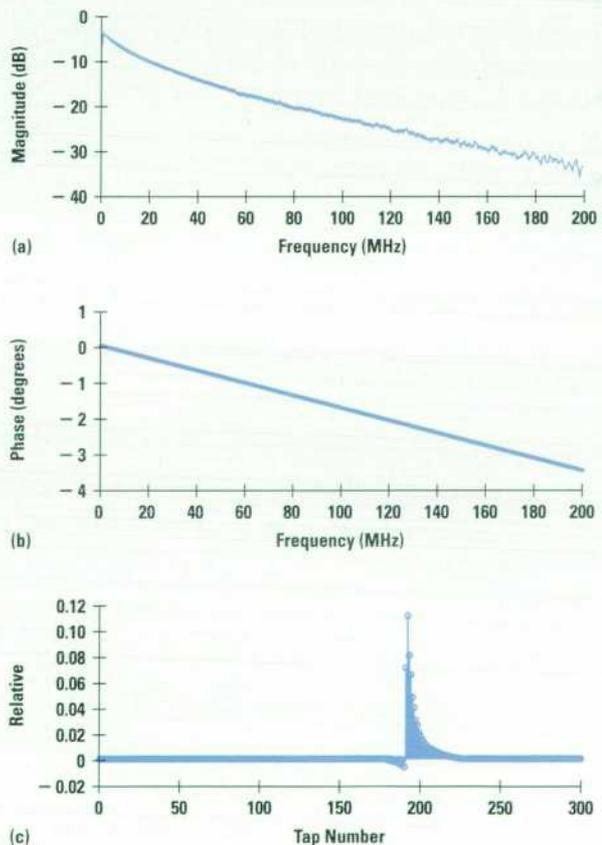
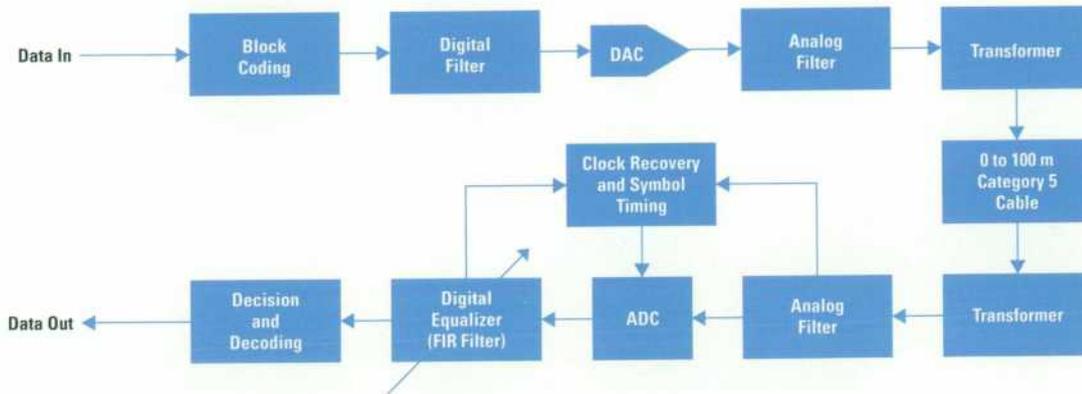


Figure 10

Simulated transmission path.¹³



with a transfer function approximating the worst-case propagation loss as set out in EIA/TIA 568.

At the receiver a fifth-order Butterworth filter at 40 MHz provided anti-aliasing and out-of-band noise rejection. The analog-to-digital converter (ADC) had variable resolution and the effect of an AGC circuit was modeled by ensuring that the applied signal occupied the full scale of the converter. The equalizer structure was T/2 and used

the well-known gradient LMS adaptation algorithm. Internal accuracy was maintained at 12 bits although the coefficient resolution and the number of taps were variable.

Figure 11 shows the effect of transformer coupling in terms of baseline wander. The system output level diagram is shown with and without a block code. Clearly the baseline wander in the first case is intolerably high.

Figure 11

Output level diagram (a) without and (b) with 16B60 block code.¹

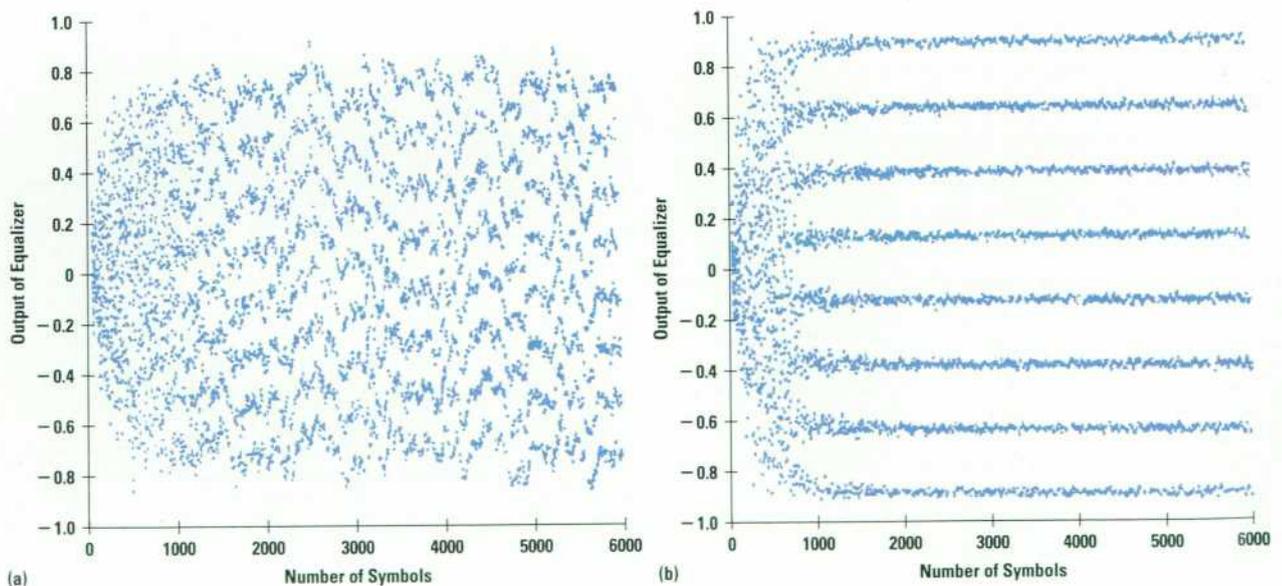
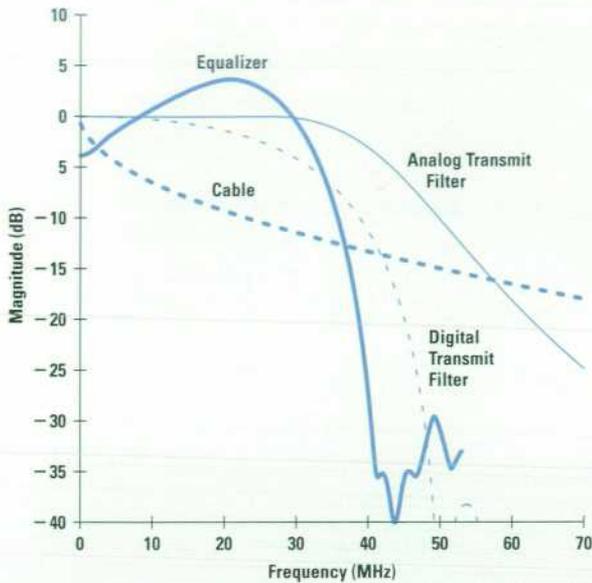


Figure 12

Frequency responses of cable, digital transmit filter, analog transmit filter, and equalizer.

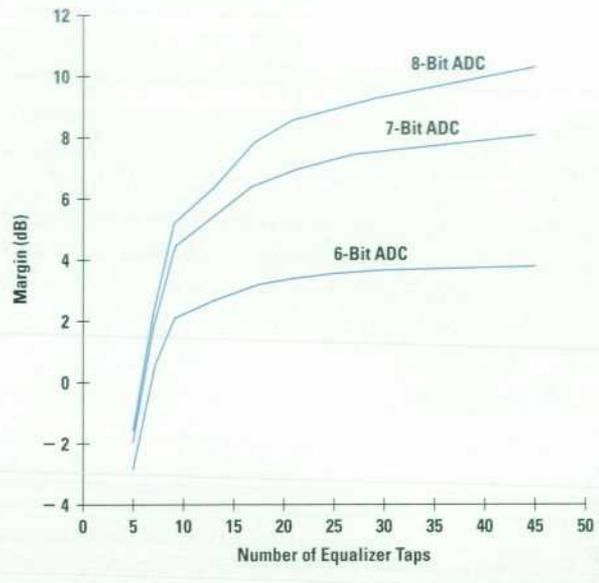


To compare the results obtained with different numbers of taps and resolutions, the metric of SNR margin over $BER = 10^{-10}$ was used. A positive margin of x dB indicates that the system has exceeded the signal-to-noise ratio (SNR) required to obtain a bit error rate (BER) of 10^{-10} by x dB. Below the roll-off region the equalizer adapts to whiten the incoming signal spectrum, effectively inverting the cable response as seen in **Figure 12**.

Figure 13 clearly shows that positive margins begin to be achieved for systems using at least a 6-bit ADC and at least 7 T/2 taps. More important, practical margins are possible for quite low-complexity systems. As an example, a 7-bit ADC used with a 25-tap T/2 fractionally spaced equalizer offers a 7.3-dB margin beyond a BER of 10^{-10} . Preliminary designs targeting HP's CMOS14 process used an 8-bit ADC and a 27-tap filter for each channel to yield a total power consumption of approximately 4 watts. Because the ADC and FIR elements are critical blocks in the design, reducing to a 7-bit data path as in the simulation above would help area and power performance markedly. For example, reducing the resolution of a flash ADC by one bit halves the area and power consumption.

Figure 13

Margin as a function of T/2 equalizer length with 10-bit coefficients (12 internal) for 6-bit, 7-bit, and 8-bit ADCs.



Hardware Demonstration

A basic hardware proof-of-concept demonstrator was constructed using off-the-shelf components following the block diagram shown in **Figure 14**. A balanced 8-level transmitter using a basic 5B2O code was written in VHDL

Figure 14

Hardware demonstrator block diagram.

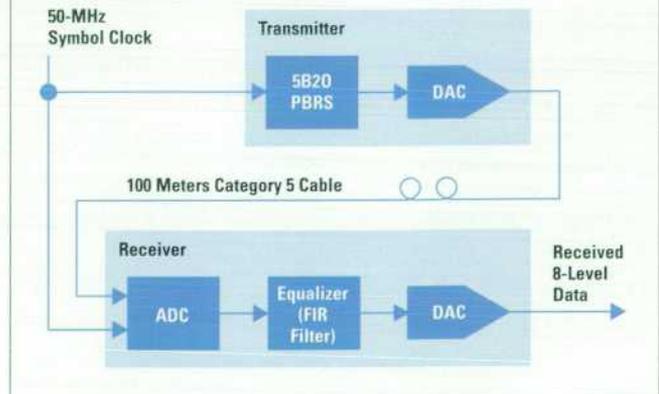
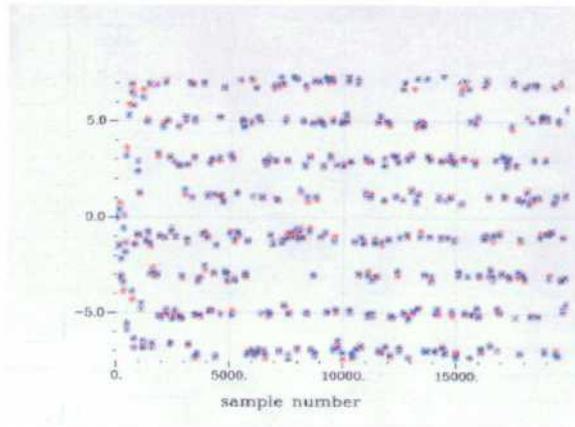


Figure 15

Eye diagram using the simulated equalizer. The vertical scale is voltage in arbitrary units. The graph shows rapid convergence of the output signal to the quantized levels of 8-level PAM.

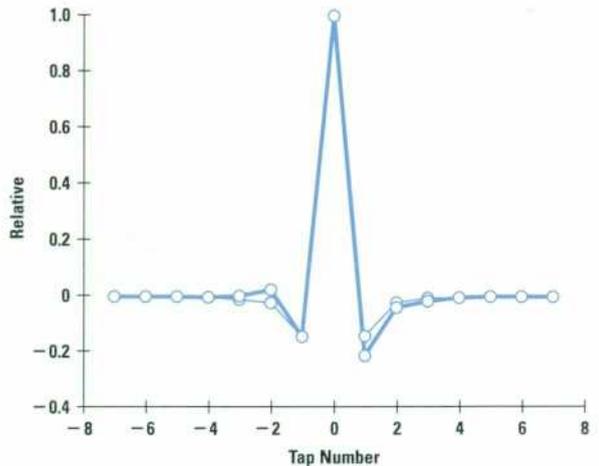


and synthesized into an Altera 7K CPLD (complex programmable logic device) and used to drive a digital-to-analog converter (DAC) clocked at the symbol rate of 50 Mbaud. After traversing the channel of 100 meters of Category 5 cable and transformers, the resulting signal was clocked into an analog-to-digital converter by a delayed version of the transmitter symbol clock. Seven-tap and 15-tap FIR (finite impulse response) filter equalizers were produced in the Altera Hardware Description Language by modifying parameterized macros to accommodate limited-precision effects and odd-order symmetric filters. To make the demonstrator much faster than serial FIR implementations, lookup tables containing precomputed partial products were used in a pipelined manner, assuming variable input data but fixed coefficients. The partial product method is also called distributed arithmetic (www.xilinx.com) and vector multiplication (www.altera.com). The FIR device was synthesized into an Altera 8K part having 12K usable gates.

The demonstrator had a T-spaced filter whose coefficients were calculated offline using the COSSAP analysis package with measured cable data. An eye diagram using the converged simulated equalizer described above is shown in **Figure 15** and the taps are shown in **Figure 16**. When running, the demonstrator had no feedback or trimming of coefficient values and relied wholly upon the accuracy

Figure 16

Taps for the filter used in the demonstrator. Heavy line: asymmetric taps. Light line: forced symmetry in taps (used in demonstrator).

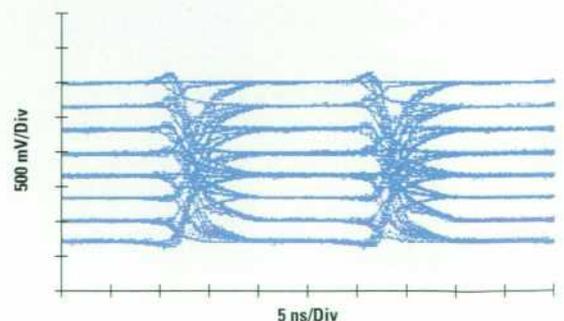


of the simulation. This method was adequate for a demonstration, but clearly a feedback algorithm such as LMS would be employed in a practical implementation to cope with changes in cable frequency response, delay, temperature, and other parameters. The transmitter clock source was delayed and used to clock the receiver, phasing being critical because of the synchronous equalizer. Moving to fractional spacing would largely remove this dependence on clock phase.¹⁴

While the open-loop nature of the system precluded long-term BER testing, the resulting eye diagrams (**Figure 17**) confirmed the operation of the hardware.

Figure 17

Eye diagrams of the received 8-level data.



Conclusion

At a rate of 424.8 Mbits/s it has been shown that a low-complexity solution exists for a higher-speed shared-medium IEEE 802-style physical layer using Category 5 cable. Neither echo nor NEXT cancellers are required and the overall complexity is only slightly greater than current 100-Mbit/s systems. The system is designed to work on the installed base of Category 5 cable and is capable of meeting both industrial and domestic EMC regulations. The data rate facilitates interworking with other physical layers using Fibre Channel bit rates.

Acknowledgments

The authors are pleased to thank Simon Crouch and Miranda Mowbray for their contribution to the block code analysis and Jim Barnes, Derek Knee, and Rajeev Badyal for their work on chip design and performance prediction.

References

1. A. Coles, E. Deliot, and S. Methley, "Proposed Higher-Speed 802.12 Physical Layer," *Presented to IEEE 802.12*, March 1996.
2. A. Coles, D. Cunningham, and S. Methley, "100-Mbit/s Data Transmission on UTP and STP Cabling for Demand Priority Networks," *IEEE Journal of Selected Areas in Communications*, Vol. 13, no. 9, December 1995, pp. 1684-91.
3. A. Coles, "Architectures for a High-Speed 802.12 UTP Physical Layer," *Presented to IEEE 802.12*, November 1995.
4. A. Coles, "Block Coding for Higher-Speed UTP PHYs," *Presented to IEEE 802.12*, March 1996.
5. A. Coles, D. Cunningham, J. Curcio, D. Dove, and S. Methley, "Physical Signaling in 100VG-AnyLAN," *Hewlett-Packard Journal*, Vol. 46, no. 4, August 1995, pp. 18-26.
6. S. Methley, "155-Mbit/s NRZ over Category 5: EMC Issues, Test Methods, Results, and Conclusions," *ATM Forum PHY Group Presentation*, February 1994.
7. A. Coles, "Modulation Schemes for a High-Speed UTP PHY," *Presented to IEEE 802.12*, November 1995.
8. A. Coles, "Multilevel Modulation Schemes: Further Analysis," *Presented to IEEE 802.12*, January 1996.
9. A. Coles and M. Mowbray, "Control Signaling for Higher-Speed UTP PHYs," *Presented to IEEE 802.12*, March 1996.
10. S. Methley, "Recent Emissions Measurements for 2-TP and Higher-Speed VG," *Presented to IEEE 802.12*, November 1995.
11. S. Methley, "An Introduction to Proposed Changes to EN55022 Radio Disturbance Characteristics of IT Equipment," *Presented to IEEE 802.12*, January 1996.
12. S. Methley, "Recent Susceptibility Measurements for 2-TP and Higher-Speed VG," *Presented to IEEE 802.12*, January 1996.
13. E. Deliot, "Simulation of Digital Equalization of Block Coded PAM Signals on UTP Category 5," *Presented to IEEE 802.12*, March 1996.
14. E. Deliot, "Digital Equalizer Structures," *Presented to IEEE 802.12*, March 1996.



Steven G. Methley

Steve Methley received his PhD degree from Imperial College, University of London in 1986. He is a senior member of the technical staff of HP Laboratories Bristol, England and is working on broadband Internet technology. Before coming to HP he worked on optical transmission at British Telecom Research Laboratories. His leisure interests include swimming, cycling, kit car construction, and films.



Alistair N. Coles

Alistair Coles is a technical contributor in the link technology group of HP Laboratories Bristol, England. He holds a BSc degree in electrical engineering (1988) from the University of Bath. A native of Plymouth, UK, he is married and enjoys hiking and guitar.



Eric Deliot

Eric Deliot is a senior member of the technical staff of HP Laboratories Bristol, England. He has worked on digital signal processing and simulations and more recently on IEEE 1394 technology. He received his PhD degree in electrical engineering from the University of Manchester and joined HP in 1994.

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 (see article, page 19). 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.

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

Figure 1

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

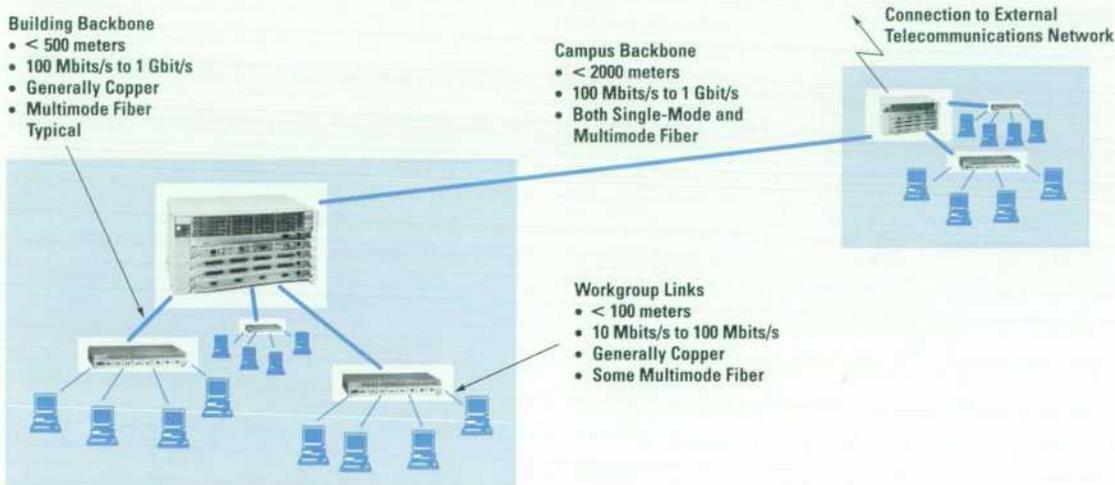
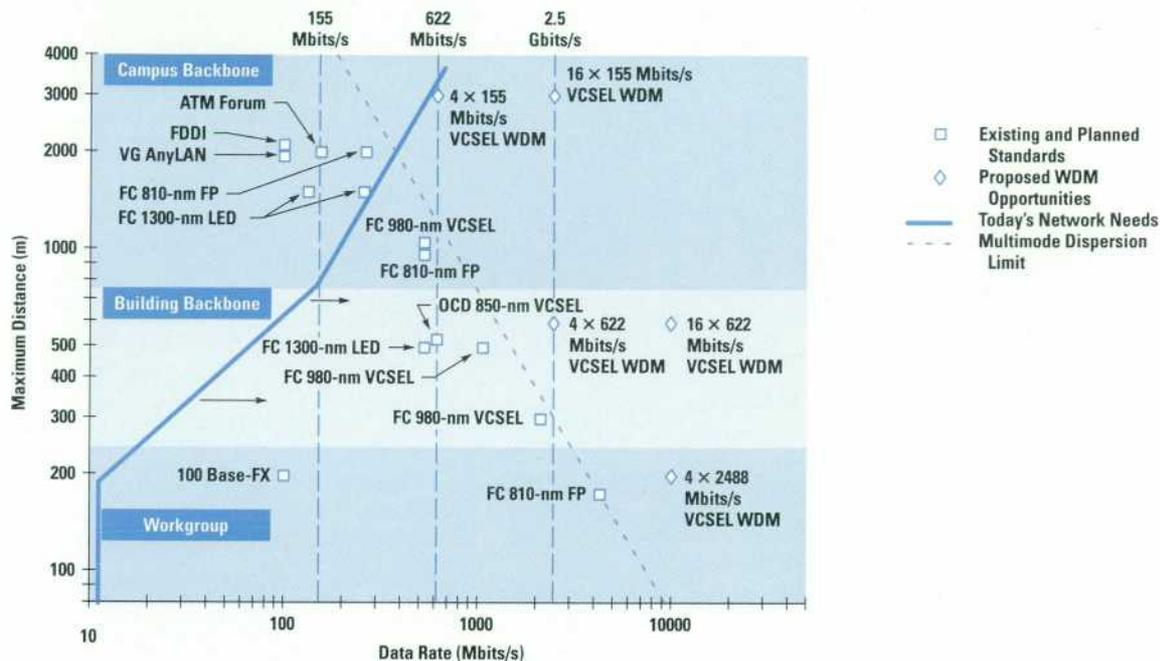


Figure 2

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- μ m multimode fiber.



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- μ m 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

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- μ m 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.

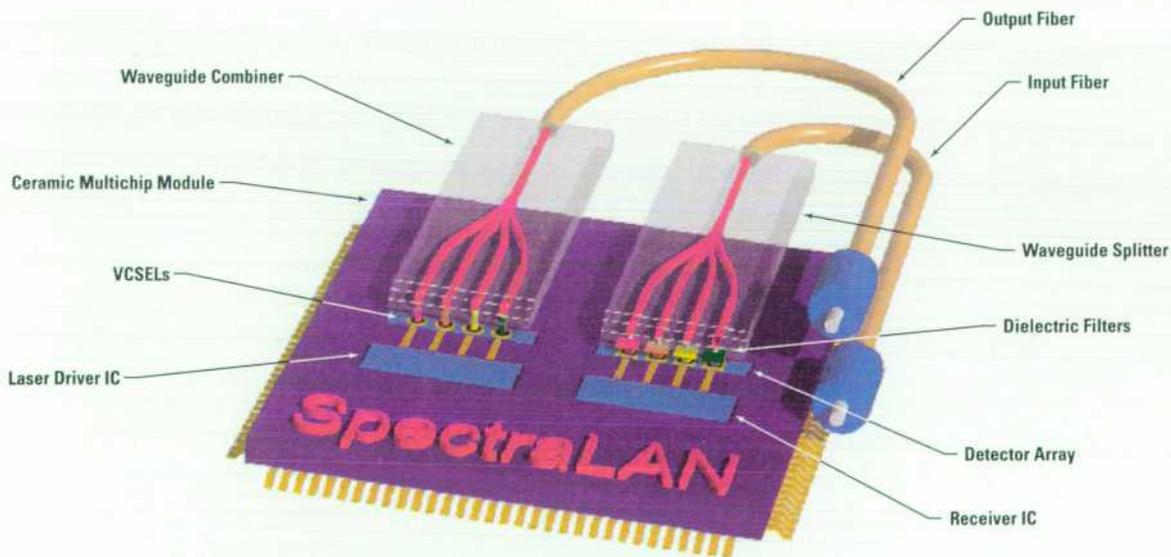
SpectraLAN Project

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 on page 53.

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

Figure 3

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



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

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.

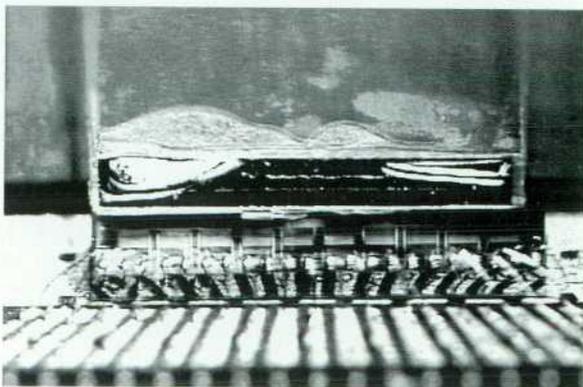
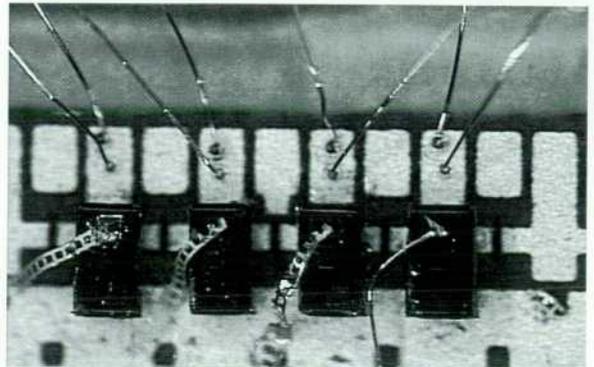


Figure 5

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



separate wafers, each of a different wavelength. **Figure 5** shows a set of four discrete VCSELs packaged in a multi-chip 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.

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 multi-wavelength 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

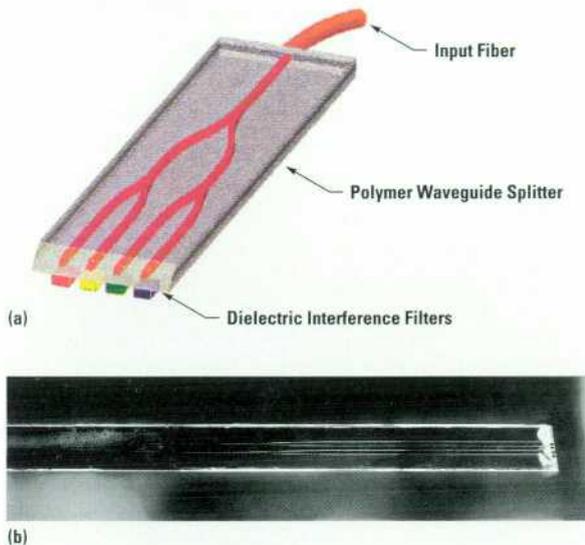
waveguide is then coupled to the fiber either through permanent attachment (pigtailed) or through a detachable connector. All of the devices built to date have been pigtailed. We have observed total insertion losses, from VCSEL to fiber, of 2.1 dB in a Polyguide™ combiner. A 4-to-1 combiner using single-mode 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 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.

Figure 6

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



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

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.

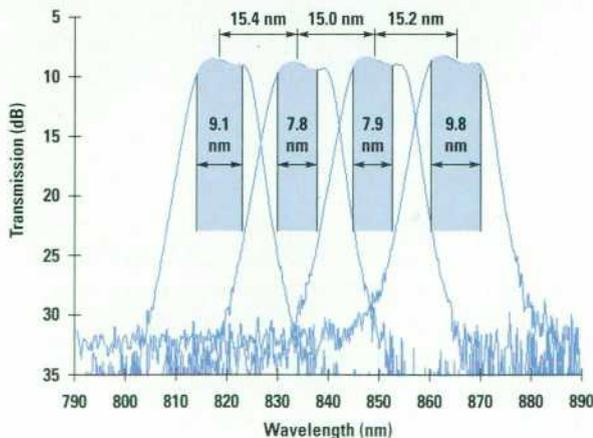
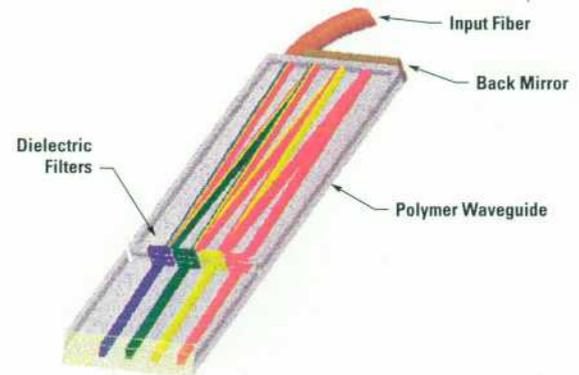


Figure 8

Conceptual drawing of zigzag demultiplexer.

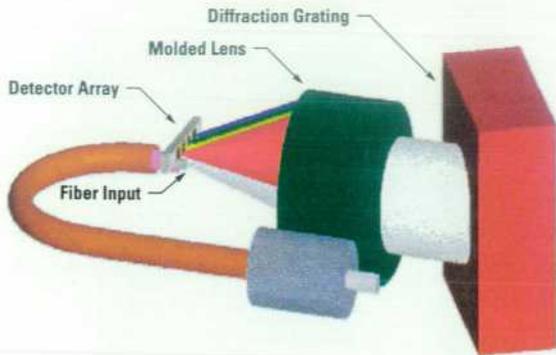


filter technologies found in the splitter demultiplexer. The geometry, however, is quite different. In the zigzag demultiplexer, the input light impinges upon the first dielectric filter at an angle. One wavelength is transmitted through the filter, while the remaining wavelengths are reflected. The 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

Figure 9

Conceptual drawing of grating-based demultiplexer.



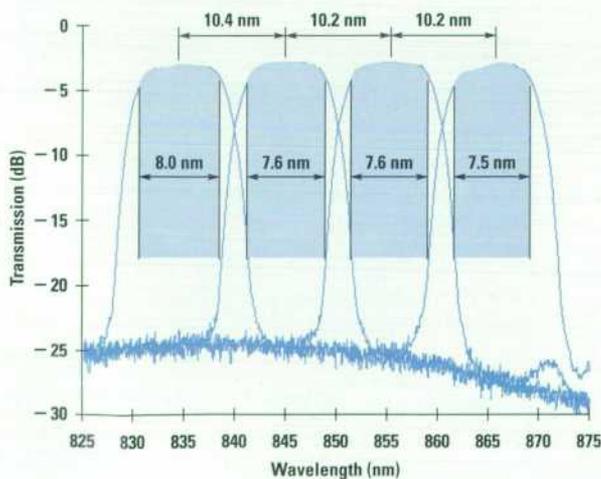
economical for the LAN application. Recent advances in precision plastic molding may hold promise for accomplishing this goal.

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

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.



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- μ m 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 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- μ m multimode fiber is 500 m. **Figure 13b** shows the

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**.

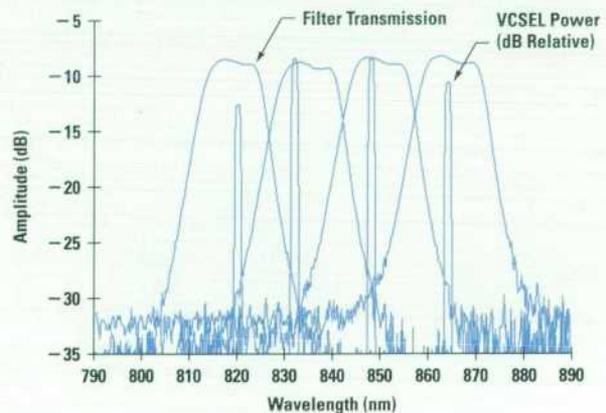


Figure 12

Experimental arrangement for four-channel WDM system measurements.

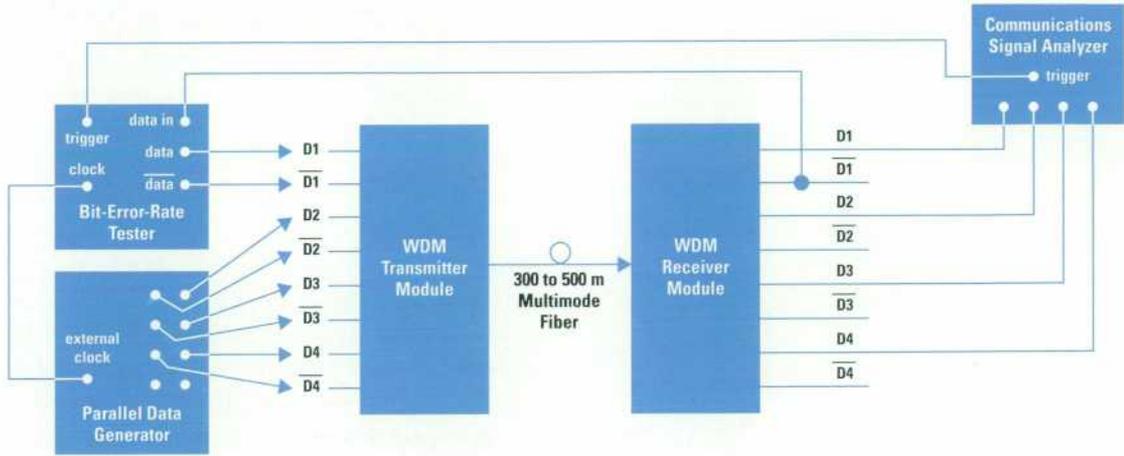


Figure 13

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- μ m-core graded-index multimode fiber.

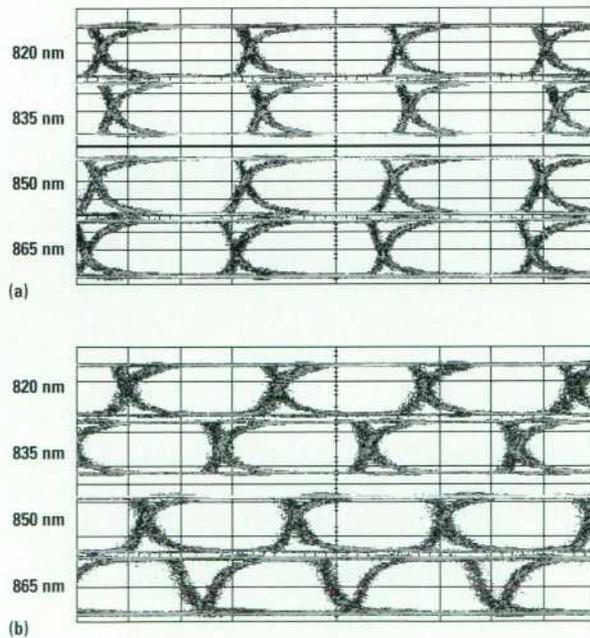
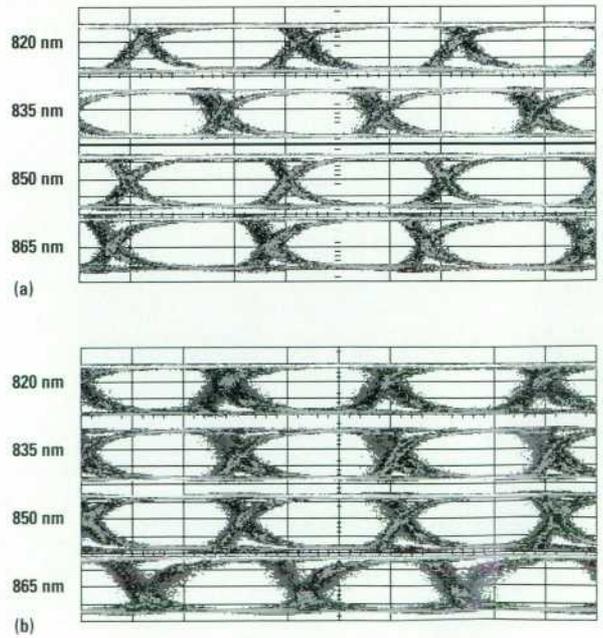


Figure 14

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- μ m-core graded-index multimode fiber.



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 **Figure 14b**. 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.

Table 1a

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

Input Channel Wavelength (nm)		Output Channel Wavelength (nm)			
		820	835	850	865
820	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 1b

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

Input Channel Wavelength (nm)		Output Channel Wavelength (nm)			
		820	835	850	865
820	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

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- μ m 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 manager. 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.

Gigabyte-per-Second Optical Interconnection Modules for Data Communications

Kenneth H. Hahn

Kirk S. Giboney

Robert E. Wilson

Joseph Straznicky

A ten-channel parallel optical link module operating at 1 Gbit/s per channel has been developed in the POLO (Parallel Optical Link Organization) program. Key components include vertical-cavity surface emitting laser and detector arrays, bipolar transceiver ICs, a high-speed ball-grid array package, polymer waveguides, and multichannel ribbon fiber connectors. Applications of the POLO module include computer clusters, switching systems, and multimedia.

The performance of advanced computer and communications systems is increasingly limited by the constraints of electrical interconnects. Future demand for interconnect bandwidth in computing and switching systems will be met by optical interconnects. Evolving communications standards such as ATM, gigabit Ethernet, and Fibre Channel require serial data rates approaching and often exceeding 1 Gbit/s. The next generation of high-performance processors will have clock speeds in excess of 300 to 400 MHz and aggregate processor bus bandwidths of more than 2 to 3 Gbytes/s. The increasing performance of such systems has driven the development of Gbyte/s interconnection standards such as SCI and Super HIPPI.

Such demands, when combined with stringent low-cost and high-performance specifications, cannot be met by any existing commercially available interconnect technology. Given the constraints of standard optical and electrical interconnections, parallel optical interconnection solutions operating at Gbyte/s data rates offer a number of advantages. The input and output data is inherently in parallel format, which reduces latency of multiplexing and demultiplexing functions and simplifies system integration. Parallel optical links amortize packaging costs over multiple channels while minimizing link latency and module footprint. By comparison, serial links will be expensive and

bulky in multiple-channel applications. Copper wire has a very limited bandwidth-length product and is unsuitable for Gbyte/s data communications in the local area (10 to 300 m). In the telephone central office environment, for example, electrical interconnects between high-capacity switching systems are creating serious bottlenecks in terms of the sheer bulk of the cable required, the limited backplane real estate available for connections, and the resultant EMI created by large electrical cable bundles.¹

The key performance parameters in such interconnections are bandwidth-length and interconnection density. Flexible film cables with impedance-controlled transmission lines offer high bandwidth and density. However, effective operating lengths are limited by attenuation and noise to less than 1 to 2 m. Twisted-pair and microcoaxial cables can accommodate transmission lengths of approximately 10 to 20 meters. However, they are bulky and relatively expensive. Connectors also limit performance. Today, an 18-twisted-pair cable for Gbyte/s Scalable Coherent Interface (SCI) measures 4 inches by 1.25 inches. In

comparison, a 10-to-12-fiber optical connector will be less than 0.4 inch by 0.3 inch, representing an order of magnitude reduction in cross-sectional form factor. Optical fibers in ribbon form have much higher density, lower attenuation and skew, competitive cable cost, and future scalability to multi-Gbit/s line rates.

Parallel optical data links are expected to be widely used as interconnections for computer clusters, switching systems, and multimedia (**Figure 1**). Cost reduction and demonstration of reliable and robust operation will be critical to the success of parallel optical links in markets presently dominated by copper-wire and serial optical data links.

POLO Program

To demonstrate the technical feasibility of low-cost, high-performance parallel optical data links, a three-year collaborative program was launched by HP, AMP, Du Pont, SDL, and the University of Southern California (USC) in August 1994.² Led by HP, the POLO (Parallel Optical Link

Figure 1

Applications of parallel optical data links.

- Parallel Optical Data Links**
- > 1 Gbit/s per Channel
 - > 1 Gbyte/s Aggregate Bandwidth
 - < 300 meters
 - Ribbon Fiber Cable and Connectors
 - N Channels
 - 20 Fibers with 1-Inch-Wide Package

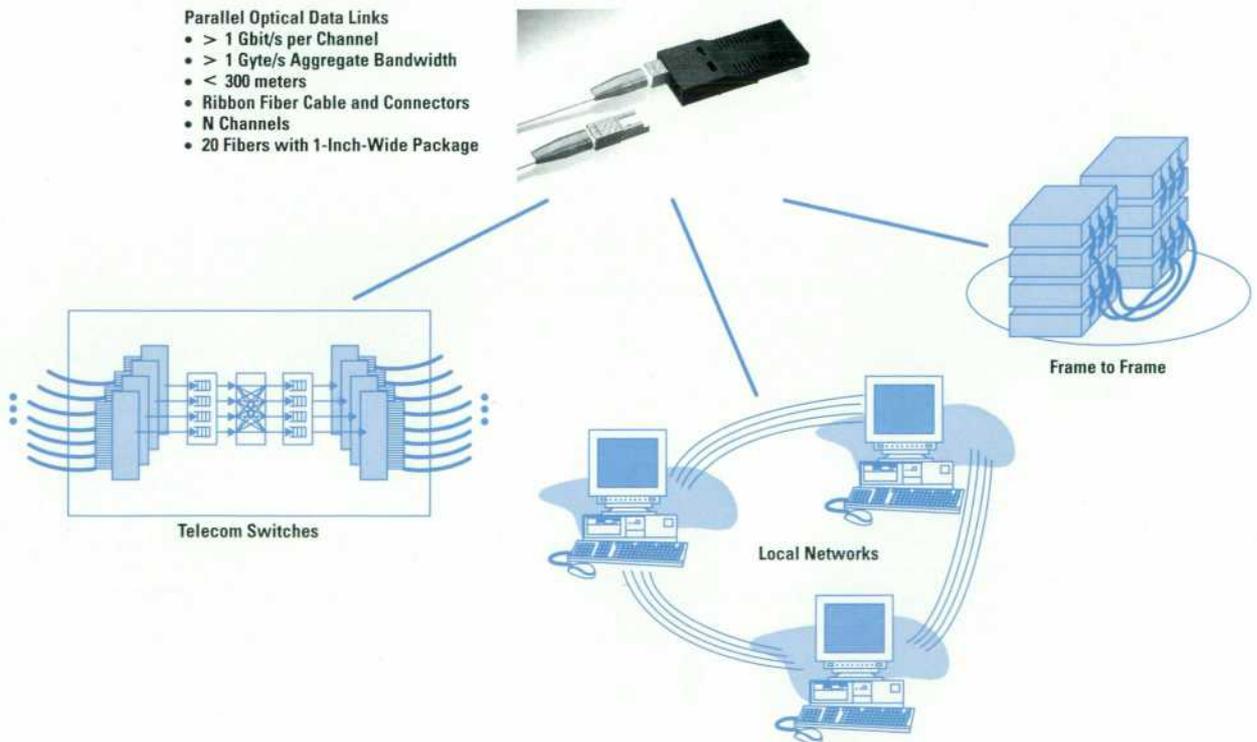
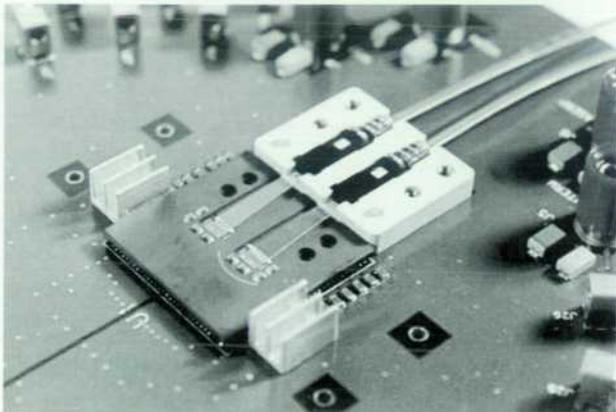
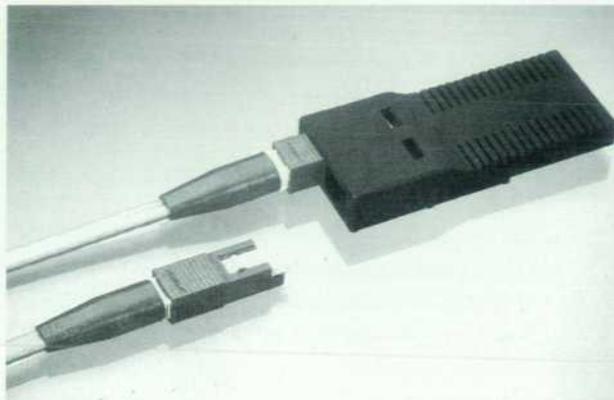


Figure 2

(a) POLO-1 module on evaluation board. (b) POLO-2 module with push/pull connectors for ribbon fiber interface.



(a)



(b)

Organization) program has demonstrated 20-Gbit/s throughput in a one-inch-wide footprint, developed a high-density optical connector for ribbon fiber, and demonstrated the operation of POLO modules in a workstation communication testbed at USC. Two generations of modules have been developed, as shown in **Figure 2**. The first generation (POLO-1) module operated at 622 Mbits/s per channel with 980-nm vertical-cavity surface emitting lasers (VCSELs) and featured a prototype connector assembly and a 1.75-inch-wide leadframe package. The second generation (POLO-2) module operates at 1 Gbit/s per channel and incorporates a higher-density (1-inch-wide) ball-grid array (BGA) package, a multichannel ribbon fiber connector from AMP, and 850-nm VCSELs and MSM (metal-semiconductor-metal) detector arrays.

Module Design and Performance Summary

The performance of the POLO-2 module is summarized in **Table I**. The maximum length is limited by the modal bandwidth of standard multimode fiber. While interchannel skew in ribbon fiber can limit length for synchronous operation of parallel channels, recent results indicate that synchronous parallel transmission for more than 1 km is possible at 622 Mbits/s per channel by cutting each fiber in sequence from the same fiber pull, limiting interchannel skew to less than 1 ps/m.³

Table I

POLO-2 Module Performance Summary

Number of channels	10 transmit and 10 receive (9 data plus 1 clock or 10 data)
Data rate per channel	0 to 1 Gbit/s
Length	< 300 m
Electrical interface	Differential ECL, latched or unlatched by clock channel
Multichannel module package	Ceramic ball-grid array
Module width	2.5 cm
Wavelength	850 nm
Connector	Lightray MPX (based on MT ferrule)
Optical interface	62.5/125- μ m graded-index ribbon fiber to polymer waveguide
Power dissipation	< 2W or < 100 mW/channel
Receiver sensitivity	-17 dBm (-20 dBm at detector), single channel only

Figure 3 shows a rendering of the POLO-2 module. The key components integrated into the package have been extensively described in the literature, including vertical-cavity surface emitting lasers (VCSELs)⁴ and Polyguide™ polymer optical waveguides.⁵

The design of the optical-electrical interface is shown in **Figure 4**. The VCSEL and detector arrays are packaged with the transceiver ICs in a 324-pin ceramic BGA package. Polyguide waveguides couple light between the VCSEL and detector arrays and the ribbon fiber using 45° out-of-plane mirrors and fiber-to-waveguide connectors. The ceramic package features impedance-controlled traces and integrated resistors for termination of input ECL signals. The use of a 45° optical interface allows the VCSELs and detectors to be packaged in close proximity to the transceiver ICs, allowing control of electrical parasitics and GHz-bandwidth operation. Because the waveguides are multimode, simultaneous alignment of ten channels to the VCSEL and detector arrays is possible with loose alignment tolerances.

VCSELs and MSM detectors

VCSELs are ideal sources for optical data links. The devices are processed and characterized at the wafer level, and one-dimensional or two-dimensional arrays are easily fabricated. Light is emitted perpendicular to the substrate with a circular beam that enables efficient, direct, fiber or waveguide coupling. For parallel links, VCSEL arrays can be fabricated to match the pitch of the optical waveguide array. Large-area top emitting 850-nm VCSELs are used

Figure 3
POLO-2 module design.

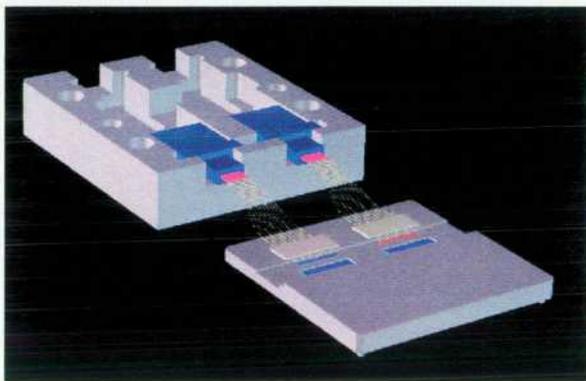
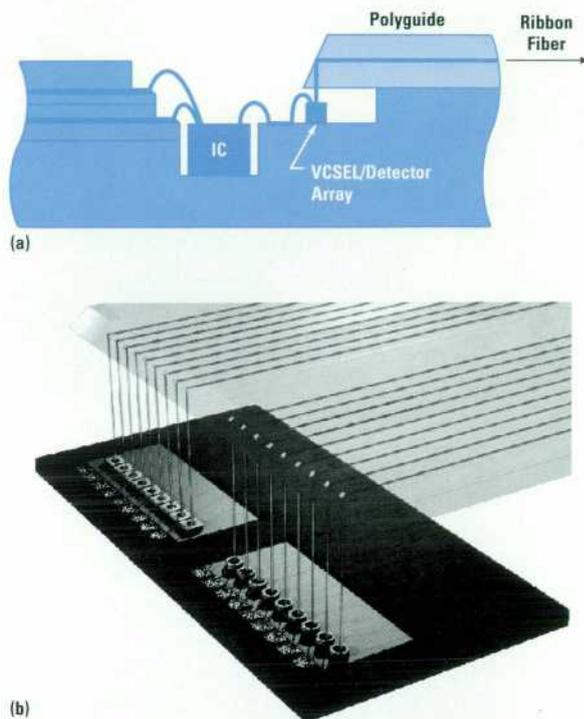


Figure 4

(a) Optical-electrical interface design. (b) Coupling of VCSEL and detector arrays with optical waveguides.

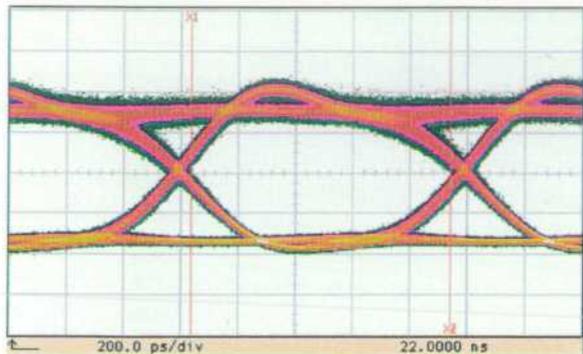


in the POLO-2 module. The threshold currents of these 18- μm -diameter VCSELs are about 3 to 4 mA. The lasers are typically prebiased near threshold to guarantee a high extinction ratio for all channels, and modulated to peak output power of ~ 2 mW. The low relative intensity noise (RIN) and reflection sensitivity of the VCSELs allows Gbit/s data rates in multimode fiber links with low BER. RIN is typically less than -130 dB/Hz under typical operating conditions. We have previously shown that large-area VCSELs emit in multiple transverse modes, leading to reduced coherence.⁶ This reduces the susceptibility of the multimode fiber link to modal noise, making these sources ideal for such applications. **Figure 5** shows an eye diagram of an 850-nm VCSEL biased below threshold and driven with a pseudorandom binary sequence (PRBS) at 1 Gbit/s. The eye is wide open and the measured BER is $< 10^{-13}$.

An attractive feature of VCSELs is their ability to scale to higher data rates. Modulation of greater than 3 Gbits/s per

Figure 5

Eye diagram of an 850-nm VCSEL biased below threshold and driven with a pseudorandom binary sequence at 1 Gbit/s.



channel has been successfully demonstrated. **Figure 6** shows the frequency response of a 980-nm VCSEL at two bias currents, showing a small-signal -3 -dB electrical frequency response of 6.6 GHz at the larger bias.

A 1-by-10 MSM detector array is shown in **Figure 7**. These devices are fabricated with a straightforward two-mask process. The interleaving metal fingers on a bulk GaAs layer have very low capacitance, allowing large area detectors with a measured -3 -dB frequency of 1.5 GHz. The detector area of $200 \times 200 \mu\text{m}^2$ provides greater than ± 50 - μm alignment tolerance to the optical waveguides. The measured responsivity at 850 nm is $> 0.4\text{A/W}$ and fall times are < 200 ps.

Figure 6

Frequency response of 980-nm VCSEL at two bias currents.

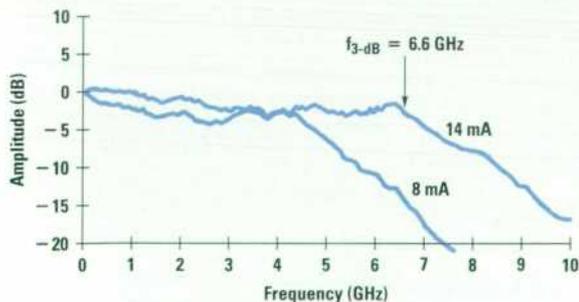


Figure 7

1-by-10 MSM detector array.



Transmitter and Receiver ICs

Transmitter and receiver ICs fabricated with Hewlett-Packard's HP-25 silicon bipolar process are used in the POLO module. The transmitter IC contains ten laser drivers that use common reference voltages to set the VCSEL prebias and modulation currents. The transmitter input and receiver output interfaces are differential ECL.

Since the receiver determines the link architecture, several versions of the receiver IC have been designed to provide maximum user flexibility, including arrays of latched digital receivers, unlatched digital receivers, and analog transimpedance amplifiers for linear testing. The latched receiver has nine data channels and one clock channel. The output data is synchronized by the clock channel at the receiver output, removing any accumulated skew and jitter. The unlatched receiver allows the module to operate as ten independent serial links.

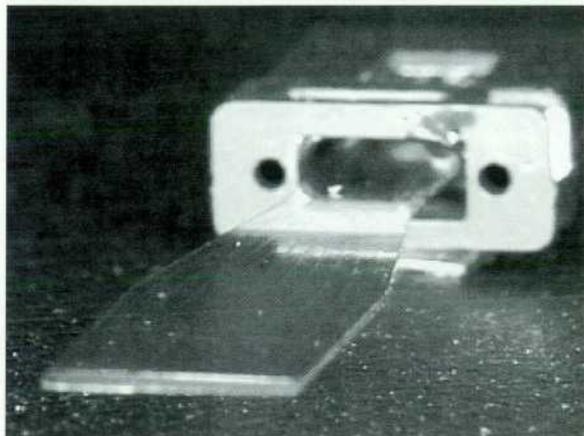
Both ac-coupled and dc-coupled versions of the receivers have been fabricated and tested. Because each channel determines its own threshold, the ac-coupled system has much higher channel-to-channel dynamic range. However, data needs to be encoded because of a low-frequency cut-off. The dc-coupled version can handle any data pattern, but channel-to-channel uniformity in received power (within several dB) is required because a single threshold is used across all channels.

Polyguide™-Ribbon Fiber Optical Interface

The use of polymer waveguides allows the waveguide design to be easily tailored to system requirements, including waveguide dimensions, pitch, and numerical aperture. For example, the waveguide pitch is $360 \mu\text{m}$ at the p-i-n detector interface and $500 \mu\text{m}$ at the VCSEL interface, but a smooth taper allows a waveguide pitch of

Figure 8

Polymer waveguide assembled with an MT-style ferrule for the multichannel optical interface.



250 μm at the ribbon fiber interface. The width and numerical aperture of the polymer waveguide are optimized to increase coupling efficiencies and optical alignment tolerances at each interface.

The Polyguide™ waveguides are assembled with an MT-style ferrule and aligned to the VCSEL and MSM detector arrays on the ceramic package. **Figure 8** shows a 10-channel polymer waveguide integrated with an MT-style ferrule.

Guide pins in the MT ferrule allow for accurate optical alignment of this assembly with ribbon fiber.

Assembly with BGA package

The POLO-2 module is the first fiber-optic module based on a ball-grid array (BGA) electrical interface. A significant advantage is the high pin density of the BGA. For example, the use of a BGA enabled a $3\times$ reduction in package size compared to the leadframe package of POLO-1. Other advantages of BGA technology include compatibility with standard surface mount processes, high thermal conductivity, and low electrical parasitics.

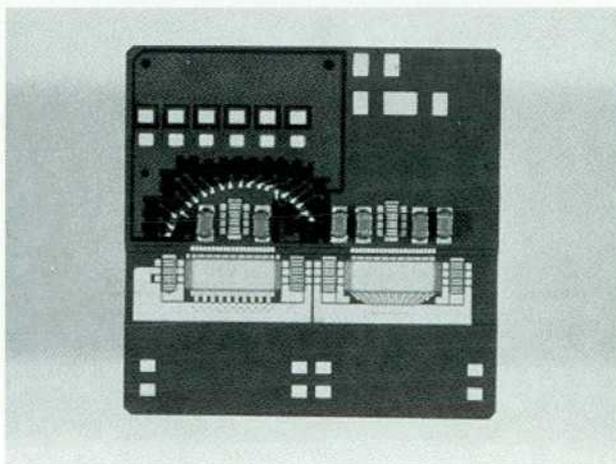
The VCSELs and p-i-n detectors, laser driver and receiver ICs, and bypass capacitor arrays are mounted on the BGA substrate and wire-bonded (**Figure 9**). Polyguide™ waveguides are aligned to the optoelectronics and attached to the package, forming the optical interface to the ribbon fiber. The 18-by-18 BGA is on standard 0.050-inch pitch, resulting in a total module width of 1 inch. Integrated 50Ω resistors in the ceramic package allow termination of the input ECL signals.

Push-Pull Connector for Ribbon Fiber Interface

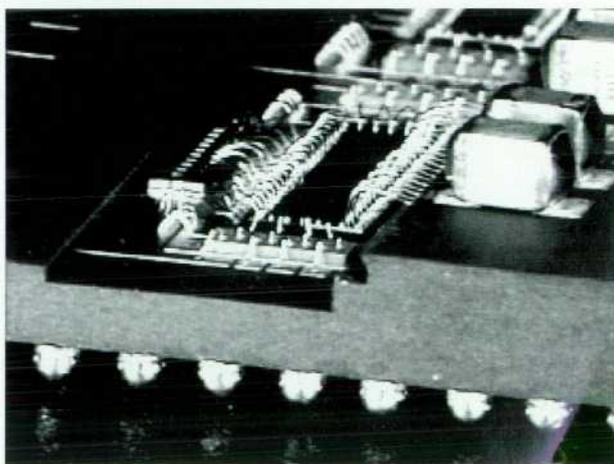
POLO-2 features a multichannel ribbon fiber connector developed by AMP. Only 9 mm wide and capable of handling 12 fibers, this connector is based on the precision molded MT array ferrule housed inside a push-pull SC-style

Figure 9

(a) 324-pin BGA package for optoelectronic integration. (b) Wire-bonded IC and detector arrays.



(a)



(b)

housing. The ribbon fiber cable uses 62.5/125- μm fiber and meets the requirements of GR-001435 *Generic Requirements for Multifiber Optical Connectors for Type IR Media (Ribbonized Fiber Enclosed in Reinforced Jacket)*. The design and construction of the push-pull connector is also in accordance with the optical, environmental, and mechanical testing requirements of the same Bellcore generic requirement specifications. Thus, the uniformity of the insertion loss across 10 channels will be kept below 0.6 dB throughout the service life, which includes 200 durability mating cycles, and the optical insertion loss for the interface will be less than 2 dB at the end of the service life. **Figure 2b** shows the assembled POLO-2 module with ribbon fiber connectors and a plastic housing that provides a receptacle for the connectors.

System Results

The module is mounted on an evaluation board for characterization. To prevent the transfer of any mechanical loads from the ribbon fiber cable to the internal module components, the module housing mounts rigidly to the printed circuit board. Supply voltages of -5V and -3V are required for transmitter and receiver operation. An additional -2V supply is also required for ECL termination.

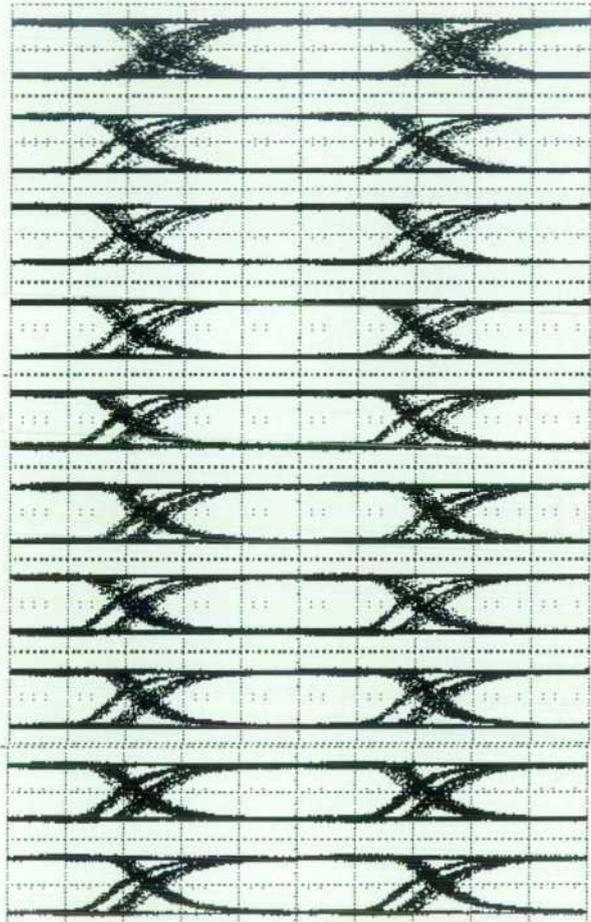
To test BER with worst-case cross talk conditions, all of the transmit and receive channels of one module are operated in loopback mode, with the transmitter and receiver of one module connected by a single ribbon fiber. A multi-channel data generator (HP 80000) is used to modulate the ten transmitter channels with independent PRBS streams. **Figure 10** shows the eye patterns of all 10 channels in simultaneous operation at 1 Gbit/s at the receiver output.

The BER for each channel is $< 10^{-11}$, and an extended measurement of one channel results in $\text{BER} < 10^{-14}$ with 300 m of low-skew ribbon fiber. While some pattern dependent jitter is observed, the eyes are clearly open at 1 Gbit/s. The rise and fall times are < 350 ps, and channel-to-channel skew (excluding ribbon fiber skew) is < 100 ps. The eye opening (timing margin for $\text{BER} < 10^{-9}$) is typically $> 70\%$ for most channels. **Figure 10b** shows ten simultaneous output eye patterns of the module on a single oscilloscope trace. The aggregate timing margin for all channels is better than 50%.

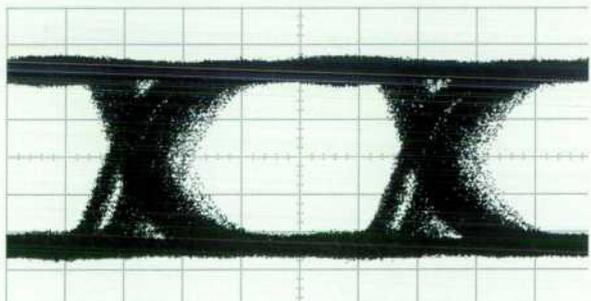
We have also obtained operation of the POLO-2 module at 1 Gbit/s per channel with an ac-coupled receiver. Initial

Figure 10

Output eye patterns of unlatched, dc-coupled module at 1 Gbit/s per channel. (a) 10 channels shown separately. (b) 10 channels aggregated.



(a)



(b)

measurements show significantly improved channel-to-channel dynamic range. Latched receivers have been operated previously with the POLO-1 module at 622 Mbits/s. Performance comparable with unlatched systems has been demonstrated. A 622-MHz clock channel synchronized nine data channels, eliminating accumulated skew and jitter at the receiver output.

Network Interface Testbed

A prototype POLO module with evaluation board has been successfully integrated into a Gbit/s experimental workstation network at the University of Southern California (USC). The network uses experimental high-speed network interface boards called Jetstream, which were developed at Hewlett-Packard Laboratories Bristol.⁷ One of these boards is inserted into each of two Hewlett-Packard 9000 Series 700 workstations, which form the two nodes of the network. Eight channels (4 transmit, 4 receive) of the POLO module, each running at 1 Gbit/s, were exercised between the two workstations, which were connected by 500 m of low-skew fiber ribbon. The POLO module successfully transmitted and received multi-Gbit/s data packets error-free in this network.

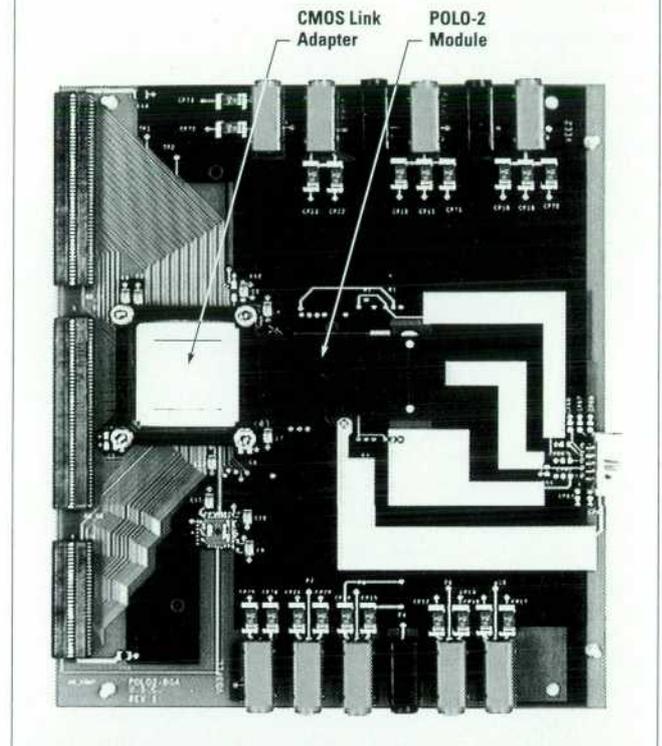
In addition, a link adapter board for a host interface has been fabricated at USC (**Figure 11**). This board contains a CMOS link adapter chip, which will directly interface to the POLO module and to external synchronous FIFO buffers. This will allow the use of the hardware interface with generic bus architectures such as PCI or other open bus standards. USC has recently demonstrated 1-GHz clocking and 1-Gbyte/s throughput in the link adapter chip.

Future Developments

The bandwidth demands of processing and communications systems will continue to multiply in the foreseeable future. While ribbon fiber has the highest bandwidth, density, and length capability, its implementation at very short distances (less than 1 to 2 meters) is limited by termination costs (i.e., parallel optical modules and connectors). To demonstrate the cost/performance superiority of parallel optics for short-distance applications such as shelf-to-shelf and board-to-board interconnections, greater system integration and functionality need to be demonstrated to the end user. The functionality includes interface compatibility with common communications standards and direct integration with network and processor buses. The integration and packaging include small footprint, low

Figure 11

Link adapter board for interfacing a POLO-2 module with a computer bus.



power consumption, and cheap optical subassemblies packaged as standard electronic components.

Conclusion

Parallel optical links that offer the highest bandwidth-length and bandwidth-density performance available have been designed and demonstrated. 1 Gbyte/s duplex operation over several hundred meters of ribbon fiber has been obtained with a 1-inch-wide optical interconnection module with ten transmit and ten receive channels.

Acknowledgments

The authors would like to thank Eric Wong and Tsu-Yau Chuang for assembly and testing, and Mike Tan and Chris Kocot for VCSEL and photodetector development. Thanks also to Jean Norman and Susan Davies for packaging, Tim Schoen for printed circuit board design and layout, Larry Hilliard and Lewis Dove at the Colorado Springs Technology Center for module assembly, and Alan Peters for board assembly. We would also like to thank Jim Dudley,

Chun Lei, and Nobu Nakamoto at the Optical Communications Division for their work on VCSELs and photodetectors, Rick Schneider for p-substrate and oxide-confined VCSEL development, Brian Lemoff for Polyguide™ support, and Dennis Brzezinski for helpful discussions on system applications. We are grateful to David Dolfi, Steve Newton, Waguih Ishak, Ron Kaneshiro, and Steve Joiner for their support of this work.

References

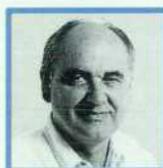
1. G.J. Grimes, S.R. Peck, and B.H. Lee, "User perspectives on intrasystem optical interconnection in SONET/SDH transmission terminals," *1992 IEEE Global Telecommunications Conference*, IEEE, New York, 1992, pp. 201-207.
2. K.H. Hahn, K.S. Giboney, R.E. Wilson, J. Straznicki, E.G. Wong, M.R. Tan, R.T. Kaneshiro, D.W. Dolfi, E.H. Mueller, A.E. Plotts, D.D. Murray, J.E. Marchegiano, B.L. Booth, B.J. Sano, B. Madhavan, B. Raghavan, and A.F.J. Levi, "Gigabyte/s Data Communications with the POLO Parallel Optical Link," *Proceedings of the 46th Electronic Components and Technology Conference*, 1996, pp. 301-307.
3. S. Siala, A.P. Kanjamala, R.N. Nottenburg, and A.F.J. Levi, "Low-skew multimode ribbon fibers for parallel optical communication," *Electronics Letters*, Vol. 30, October 1994, pp.1784-1786.
4. M.R. Tan, K.H. Hahn, Y.M. Houng, and S.Y. Wang, "SELS for short-distance optical links using multimode fibers," *Conference on Lasers and Electro-Optics 1995*, Optical Society of America, Washington D.C., 1995, pp. 54-55.
5. B.L. Booth, "Polymers for integrated optical waveguides," *Polymers for Lightwave and Integrated Optics*, C.P. Wong, editor, Academic Press, New York, 1993. Also B.L. Booth, "Optical interconnection polymers," *Polymers for Lightwave and Integrated Optics: Technology and Applications*, L.A. Hornak, editor, Marcel Dekker, New York, 1993.
6. K.H. Hahn, M.R. Tan, Y.M. Houng, and S.Y. Wang, "Large-area multitransverse-mode VCSELs for modal noise reduction in multimode fiber systems," *Electronics Letters*, Vol. 29, August 1993, pp. 1482-1483.
7. A. Edwards, et al, "User-space protocols deliver high performance to applications on a low-cost Gb/s LAN," *ACM SIGCOMM*, 1994.



Kenneth H. Hahn

A principal project scientist with HP Laboratories, Ken Hahn is currently on

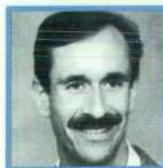
leave, attending Harvard Business School. He has been a project leader and a principal investigator of the POLO Consortium and has worked on high-speed modulators, polymer optical waveguides, and gigabit/s optical links. His PhD degree in physics is from Stanford University (1991). He is married and enjoys travel, golf, and tennis.



Joseph Straznicki

A project scientist at HP Laboratories since 1980, Joe Straznicki is responsible

for the development of high-performance optical receivers, transmitters, and subsystems in various technologies. He received his Diplom Ingenieur in cybernetics in 1969 from the Technical University of Brno, Czech Republic. Before joining HP he was with the University of California at Los Angeles and Hughes Aircraft Company.



Kirk S. Giboney

Kirk Giboney is a hardware design engineer with HP Laboratories,

responsible for parallel optical link module design and testing. He received his PhD degree in electrical engineering from the University of California at Santa Barbara in 1995, and has forty publications to his credit. Before coming to HP he was with McDonnell Douglas Astronautics Company and Hughes Research Laboratories. His interests include cycling, flying, hiking, archery, and the arts.



Robert E. Wilson

Rob Wilson received his BSME degree in 1978 from Stanford University.

A development engineer with the HP Communications Semiconductor Solutions Division, he is a packaging specialist, responsible for optical subassemblies for small-form-factor modules. He is married, has three children, and is interested in photography and elementary school science teaching.

Developing Leading-Edge Fiber-Optic Network Link Standards

David G. Cunningham

Delon C. Hanson

Mark C. Nowell

C. Steven Joiner

Advances in fiber-optic network technology within Hewlett-Packard are achieved by close cooperation between Hewlett-Packard Laboratories (HPL) and Hewlett-Packard's Communications Semiconductor Solutions Division (CSSD). This paper explores the interaction between HPL and CSSD for the advancement of high-speed LAN standards, particularly in the ATM Forum and IEEE 802.3z (Gbit/s Ethernet). Details of major technical contributions to 622-Mbit/s ATM and Gbit/s Ethernet specifications are presented.

Standardization has created a global fiber-optic LAN market in which Hewlett-Packard competes. However, successful open LAN standards are developed by consensus. Consensus is fundamental to the standardization process, since it ensures that the technological advances embodied in a final standard will be implemented by a number of vendors. In addition, standardization gives customers confidence that LAN products will not originate from a single source with corresponding higher prices. Given the need for consensus, it is important that Hewlett-Packard continuously participate in LAN standards development so that the company remains aware of current industry and future standards requirements. This understanding is a very important input into the strategic planning process for the HP Communications Semiconductor Solutions Division (CSSD) and Hewlett-Packard Laboratories (HPL).

Consensus-based standards make it impossible for any company to dominate the global optical-fiber LAN market. However, a long-term collaboration between HPL and CSSD has enabled Hewlett-Packard to be a market leader in high-speed optical-fiber LANs. Responsibility for developing future business and standards strategy is jointly owned by CSSD's strategic Pathfinders and HPL. CSSD and HPL engage the standardization process as early as possible. Involvement in embryonic standards provides valuable insight into the

capabilities and needs of Hewlett-Packard's competitors, partners, and customers. Based on this insight we adapt our standards and research strategy so that there is maximum likelihood of it being accepted by consensus-driven standards bodies. Actively influencing emergent standards ensures that a good return is achieved from Hewlett-Packard's investment in research and development. Obviously, this is a long-term commitment involving the continuous seeding of the Hewlett-Packard Laboratories research agenda many years in advance of emerging standards or customer needs.

Recently, two major high-data-rate fiber-optic LAN standards have emerged: 622-Mbit/s (OC-12) Asynchronous Transfer Mode (ATM) and Gbit/s Ethernet. Building wiring standards require optical fiber transmission over 500 m of 62.5/125- μ m (core/cladding diameter) multimode fiber (62MMF) for backbone links. This building wiring standards requirement determines the choice of transceiver technology, as illustrated in **Figure 1**. The ATM Forum, which requires a 622-Mbit/s line rate, considered long-wavelength LEDs and short-wavelength laser diodes. VCSELs (vertical-cavity surface emitting lasers) operating near wavelengths of 980 nm were also discussed in the context of Gbit/s ATM links. The Gbit/s Ethernet standards committee considered short-wavelength laser diodes and long-wavelength laser diodes. CSSD and HPL made major technical contributions to both standards, which led to the inclusion of 1300-nm LEDs in the OC-12

ATM Forum specification and 1300-nm laser diodes in the Gbit/s Ethernet specification.

This paper will provide insight into the long-term collaboration between CSSD and Hewlett-Packard Laboratories in the area of fiber-optic LAN standards. The collaboration will be explored through the chronology of Hewlett-Packard's involvement in the development of recent LAN standards, particularly in the ATM Forum and IEEE 802.3z (Gigabit/s Ethernet).

Wiring Link Length and Transceiver Technology

Building wiring is a large capital investment usually amortized over approximately 15 years. To protect this investment, International Standards Organization (ISO) building wiring system standard ISO/IEC 11801 specifies cabling architectures, link lengths, and type. Historically, fiber-optic backbones for LANs are developed to support the link length requirements defined in ISO/IEC 11801. These requirements evolved out of the recognized need in the mid-1980s to achieve a more unified approach to developing and installing LANs. As illustrated in **Figure 2**, a key aspect of this building wiring standard is the definition of link lengths: 100 m horizontal from hubs to the desktop (90 m of cable plus up to 10 m for patch cords), 500 m for building backbone, and 2 km for campus backbone. These link length requirements dictate the choices of fiber-optic cable and transceiver technology as a function of data rate. In particular, because of its suitability for use with low-cost light-emitting diode (LED) transceivers, the installed base of optical fiber is predominantly 62MMF in

Figure 1

Illustration of transceiver technology choices for OC-12 (622-Mbit/s) and Gbit/s Ethernet 62MMF links based on overfilled launch bandwidth. Building wiring standards require optical-fiber transmission over 500 m of 62.5/125- μ m multimode fiber (62MMF) for building backbone links. (LD = laser diode. VCSEL = vertical-cavity surface emitting laser. LED = light-emitting diode).

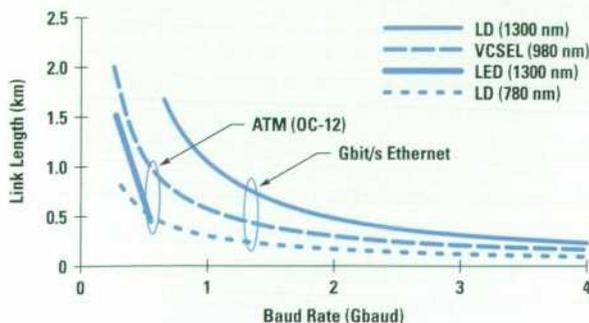
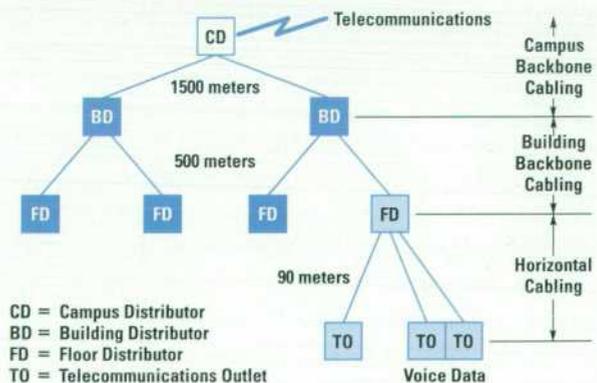


Figure 2

ISO/IEC 11801 customer premises cabling model.



both the U.S.A. and Europe. Transceiver vendors and fiber-optic LAN standards must develop transceiver technology that operates in harmony with the building wiring standards to protect the capital investment of both LAN users and LAN equipment suppliers.

The HP Communications Semiconductor Solutions Division is a leading supplier of optical-fiber optoelectronic components used to communicate over both premise backbones using primarily multimode fiber and public networks using single-mode fiber. Because of link performance and cost trade-offs, multimode fiber transceivers are developed using 650-nm, 850-nm, and 1300-nm technology. Visible 650-nm LEDs match the transmission window of large-core (980- μm diameter) plastic optical fiber, which has high attenuation but yields the lowest-cost transceivers and optical connectors as a result of relaxed mechanical tolerances. Infrared 850-nm and 1300-nm technology matches the transmission characteristics of glass multimode fiber having smaller core diameters, that is 62MMF and 50MMF (50/125- μm core/cladding diameter). These fibers have lower attenuation and higher bandwidth at wavelengths near 1300 nm compared to 850-nm operation but yield more expensive systems compared to plastic optical fiber. Single-mode fiber transceiver technology operating at 1300-nm and 1550-nm wavelengths supports the 10-to-50-km distance requirements of telecommunications single-mode fiber links and is still more expensive. Nevertheless, 1300-nm single-mode fiber links have extended transmission capabilities and are being deployed on the campus to extend beyond the distance and data rate limits of multimode fiber.

Fiber Optic LAN Standards Development

The initial fiber-optic backbone link standards developed in the mid-1980s support a 2-km campus backbone length using 62MMF. This requirement influenced the subsequent ISO/IEC 11801 campus backbone link length. The 10-Mbit/s, 2-km IEEE 802.3 Ethernet standard uses 850-nm LEDs while the 100-Mbit/s, 2-km ANSI X3T12 Fiber Distributed Data Interface (FDDI) standard requires 1300-nm LEDs because of the impact of fiber spectral dispersion at this higher data rate. Subsequently, based on the FDDI backbone link standard, a 2-km 62MMF link length specification using 1300-nm LEDs was developed for transmitting Asynchronous Transfer Mode (ATM) cells over Synchronous Optical Network (SONET) links at 155.5 Mbits/s, also referred to as optical carrier level 3 (OC-3).

This rate standard, initiated in the ATM Forum, was formalized in the T1E1.2 T1.646 broadband ISDN customer interface standard.

Long-Wavelength LED Specification

It was generally assumed that low-cost 1300-nm LEDs would be too slow for operation at 622 Mbits/s (OC-12). However, exploratory work at HPL Bristol and other manufacturers of 1300-nm LEDs indicated that the necessary 1-ns optical response time was achievable with low-cost designs. This resulted in a development program at CSSD yielding the necessary data to support an OC-12 specification for a 500-m 62MMF link length in both the ATM Forum and T1E1.2 T1.646 specifications. This is the highest data rate at which 1300-nm LEDs can reasonably be specified in multimode fiber link applications.

It was obvious to HPL researchers that a new low-cost, LED-like laser technology was required for multimode fiber Gbit/s LANs and computer interconnects. This realization was key to the initiation of vertical-cavity surface emitting laser (VCSEL) development within Hewlett-Packard Laboratories during the early 1990s.^{1,2} Since 1300-nm LEDs reach their limit at 622 Mbits/s, Hewlett-Packard developed a link length and data rate extension to Gbit/s ATM based on VCSELs operating at wavelengths near 980 nm.³ HPL demonstrated that 980-nm VCSELs could support building backbone link lengths at Gbit/s data rates with 62MMF. CSSD and HPL felt that this proposal was very suitable for Gbit/s LAN standards since it was in harmony with ISO/IEC 11801. By comparison, VCSELs operating at 850 nm were felt to be an inferior choice since they cannot support building backbone link lengths at Gbit/s data rates with 62MMF (see **Figure 1**) based on the standard overfilled launch (OFL) modal bandwidth for 62MMF.

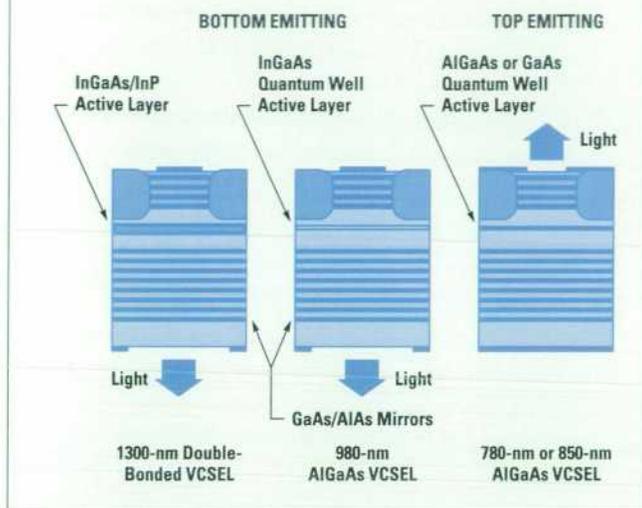
Vertical-Cavity Surface Emitting Lasers (VCSELs)

The ATM Forum OC-12 multimode fiber specification development provided an interesting first view of VCSELs entering the standards arena. **Figure 3** shows VCSEL cross sections for devices operating at 850 nm, 980 nm, and 1300 nm.

A noteworthy aspect of the OC-12 link development in the ATM Forum was the short-wavelength (780-nm and 850-nm) specification developed in competition with the 1300-nm LED specification by short-wavelength laser

Figure 3

Vertical-cavity surface emitting laser (VCSEL) cross-sections. VCSELs operating at 980 nm and 1300 nm have transparent substrates and can emit light through the substrate. These are termed bottom emitting VCSELs. They have the advantage that the active region can be placed very close to a heat sink placed on the top mirror stack.



diode transceiver vendors. The short-wavelength proposal was optimized for 50MMF rather than the dominant 62MMF. It was based on developments in the Fibre Channel (FC) standard for computer interconnect.

The short-wavelength proposal was written in such a way that it included both existing 780-nm Fabry-Perot edge emitting laser diodes, commonly referred to as compact disk (CD) laser diodes, and 850-nm VCSELs. After much contentious debate, the eventual result was that the ATM Forum adopted the two optically incompatible OC-12 multimode fiber specifications—long-wavelength LEDs and short-wavelength (780-nm and 850-nm) laser diodes—and left the marketplace to resolve the choice. In general, the winner will be the interface standard chosen by the major early adopters since other equipment suppliers will need to be compatible. CSSD sells many more long-wavelength LED-based transceivers than short-wavelength laser-based transceivers for OC-12 ATM.

Hewlett-Packard demonstrated to the ATM Forum that VCSELs operating at 980 nm are extremely reliable, are relatively simple to manufacture, are optically compatible with the specified OC-12 1300-nm receivers, and support building backbone link lengths at Gbit/s data rates with

62MMF. However, in the give and take of the standards arena, the Hewlett-Packard 980-nm VCSEL proposal for Gbit/s ATM links did not progress. One reason for this was opposition from multimode fiber suppliers to specify the higher modal bandwidth available at 980 nm with installed 62MMF. As a result of the OC-12 developments, HP focused its efforts on 850-nm VCSELs for use at Gbit/s data rates (described in the next sections) rather than 980-nm VCSELs and concentrated longer-term research on 1300-nm VCSELs, which have multiple uses with both multimode fiber and single-mode fiber.

OC-12 Technical Challenges

The ATM Forum OC-12 multimode fiber specification development brought into focus two major technical issues that must be addressed before Gbit/s optical link specifications will be acceptable to customers and LAN standards bodies. These are modal noise and robust transmission methods for 500-m backbone links using installed 62MMF. Both issues were widely recognized within the industry and ad hoc industry groups were formed to address them. The work of the ad hoc groups has now largely been transferred to Telecommunications Industry Association (TIA) fiber-optic test procedure committees. LAN standards committees have agreed to reference the new TIA fiber-optic test procedures once they mature. HPL and CSSD are very active in both the TIA committees and the LAN standards bodies.

Ethernet Frame-Based Gbit/s LANs

Because of its ease of use and relatively low cost, Ethernet has become the most pervasive LAN, with over 100 million nodes in service. The initial standard operating at 10 Mbits/s was finalized in 1985. This version of Ethernet connected all nodes via a central coaxial bus, which proved to be somewhat inflexible as users changed locations or were added to the network.

Hewlett-Packard proposed the 10Base-T star topology for Ethernet in 1987 and it became part of the IEEE 802.3 standard in 1990. In November of 1992, to meet user requirements for higher data rates and to support emergent multimedia applications, the higher-speed study group of the IEEE 802.3 committee was formed to develop 100-Mbit/s Ethernet. The efforts of the higher-speed study group culminated in July 1995 when the LAN MAN Standards Committee (LMSC) of the IEEE ratified two new

100-Mbit/s standards that use Ethernet frames: IEEE 802.3u⁴ and IEEE 802.12.⁵ These new 100-Mbit/s standards were designed to provide an upgrade path for the many tens of millions of 10Base-T and token ring users worldwide. Both of the new standards support installed customer building cabling as well as existing LAN management and application software. During the development of IEEE 802.3u and IEEE 802.12, Ethernet frame switching gained momentum as a method for increasing network capacity. Switch-based LANs are now being standardized by the LMSC.

The broad market acceptance of the new 100-Mbit/s shared media access LAN technologies is indicated by market estimates that during 1995, 25% of all network interface cards were 100-Mbit/s capable. International Data Corporation has estimated that two million Ethernet LAN switch ports were shipped during 1995.

The new 100-Mbit/s repeater and switch-based LANs require a higher-speed backbone. To address this need, both IEEE 802.12 and IEEE 802.3 initiated Gbit/s projects. It is generally agreed that the pervasiveness of Ethernet frame-based LANs will ensure that the LMSC Gbit/s standards will be the dominant Gbit/s LAN technology. Eventually, Gbit/s Ethernet frame-based LANs may reach the desktop.

Gbit/s IEEE 802.12 (Demand Priority) LANs

During 1995, IEEE 802.12 initiated development of Gbit/s demand priority LAN specifications initially as a higher-speed backbone for the 100-Mbit/s systems. HPL and CSSD demonstrated that higher-speed IEEE 802.12 could leverage some of the physical layers and control signaling developed for Fibre Channel.⁶ We also demonstrated that multimode fiber and short-wavelength VCSEL transceivers can be used to connect repeaters or switches separated by less than 300 m within a building.⁷ For longer campus backbone links, laser transceivers operating at a wavelength of 1300 nm are used. To ensure that only one Gbit/s physical layer solution set will be developed by the LMSC, HP has shifted a major portion of its physical layer activities from IEEE 802.12 to Gbit/s IEEE 802.3. This is sensible because Gbit/s IEEE 802.12 and IEEE 802.3 can use the same physical media interfaces.

IEEE 802.3 Gbit/s Ethernet

The IEEE 802.3 committee is responsible for development of the Ethernet LAN standards. In 1996, a crescendo of

activity began to extend this data rate by another factor of 10 to interconnect 100-Mbit/s hubs and switches. The fast pace of the Gbit/s Ethernet standards development effort drove the need for joint efforts between CSSD and HPL Bristol to formulate and quantify HP's optical technology strategy to meet the requirements of this standard. The development of physical layers at Gbit/s data rates brought to the forefront the two key technical challenges previously debated during the ATM Forum OC-12 specification development: modal noise and robust transmission methods for 500-m backbone links using installed 62MMF.

Since the OC-12 debates, HPL and CSSD have worked actively to understand and quantify these issues both internally and externally in various ad hoc industry groups and standards task forces. Much of the knowledge and experience gained from our internal research and the external forums was used to develop the successful dual technology proposal that HP presented to IEEE 802.3z in July of 1996.^{8,9}

Because of their higher modulation rates and narrower spectral widths (compared to LEDs) Fabry-Perot laser diodes and VCSELs are the candidates for Gbit/s Ethernet standards. All new leading-edge LAN standards evolve from the progress made in previous or related current standards. In the case of Gbit/s Ethernet, the launching vehicle was the Fibre Channel standard operating at 100 Mbytes/s or an 800-Mbit/s serial data rate. Because Fibre Channel uses an 8B10B line code and incorporates a 62.5-Mbaud control channel, its line rate is 1.0625 Gbaud.

The vast majority of installed LAN building and campus backbones in the U.S. and Europe use 62MMF. A recent survey conducted as part of the Gbit/s Ethernet standard development process concluded that 44% of the installed 62MMF links in the U.S. have greater than 500-m link length. However, Fibre Channel is an interconnect standard and it was perceived satisfactory to achieve a worst-case maximum link length of 500 m with 50MMF at 1.0625 Gbaud if this resulted in the lowest-cost interconnect solution. The Fibre Channel 50MMF link specification is not sufficient for LANs, which are dominated by 62MMF. 50MMF has a 500-MHz · km worst-case modal bandwidth specified with overfilled launch (see **Figure 4**) at 850 nm, while only 160 MHz · km worst-case modal bandwidth under the same conditions is achieved with 62MMF. Therefore, Fibre Channel operating on 62MMF achieves less than 340-m link lengths at 1.0625 Gbaud, as shown in

Figure 4

Schematic representation of overfilled launch and restricted mode launch.

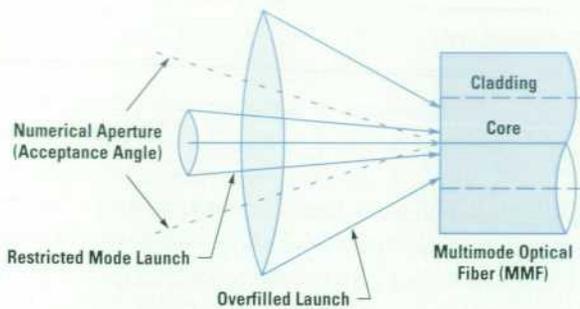
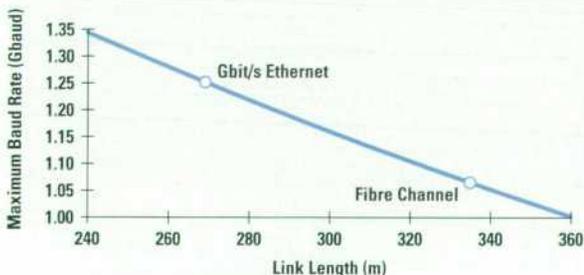


Figure 5. The inclusion of 780-nm lasers would reduce the Fibre Channel link length to less than 300 m. Gbit/s Ethernet uses the Fibre Channel 8B10B line code, which results in a 1.250 Gbaud line rate. As shown in **Figure 5**, the 18% higher data rate of Gbit/s Ethernet results in 19% shorter link length. This reduces the worst-case link length to ≈ 270 m with 62MMF, far less than both the installed 62MMF LAN link lengths and the lengths specified in ISO/IEC 11801. This disharmony with ISO/IEC 11801 and installed LAN requirements provided the motivation for the dual link technology strategy proposed by HP at the July 1996 meeting of the IEEE 802.3 Gbit/s Ethernet committee. This strategy was launched by joint efforts at HPL Bristol and CSSD. The technical problems raised by the HP strategy and our proposed solutions are the primary focus of the remainder of this paper.

Figure 5

Link lengths with 820-to-860-nm lasers and worst-case 62MMF.



Long-Wavelength Laser and Multimode Fiber Links for Gbit/s Ethernet

Because of the limited link length capability at 1.250 Gbaud of 850-nm laser diodes with 62MMF, a proposal was needed to service the installed base. The logical choice was to use existing 1300-nm laser diode single-mode fiber transceivers for single-mode fiber links up to 3 km and to extend the link length of multimode fiber. However, the use of long-wavelength laser diode single-mode fiber transmitters in conjunction with multimode fiber has traditionally been resisted by fiber-optic link designers because of the potentially catastrophic problem of modal noise. To support HP's proposal, it was necessary to develop theoretical and experimental evidence that a long-wavelength solution was robust to modal noise.

Modal Noise

Since the mid 1980s, there has been concern about modal noise caused by mode-selective loss when using relatively coherent laser diodes with multimode fiber links (see reference 10 and the references therein). Early research conducted at HPL Bristol indicated that 1300-nm laser diodes operating with 62MMF have the same modal noise performance as the 850-nm laser diodes operating with 50MMF within the ATM Forum OC-12 specification. Additionally, with 1300-nm laser diodes, the 62MMF link length supported (with the existing long-wavelength 500-MHz·km overfilled launch modal bandwidth specification) is ≈ 850 m, as shown in **Figure 6**. This link length exceeds the 500-m building backbone length defined in ISO/IEC 11801 but falls short of the 2-km 62MMF campus backbone

Figure 6

Link lengths with 1300-nm lasers and 62MMF with 500 MHz·km and 1200 MHz·km bandwidth-distance product.

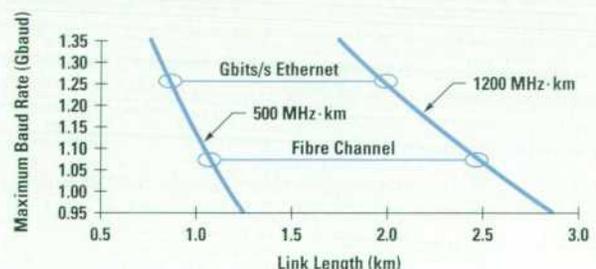
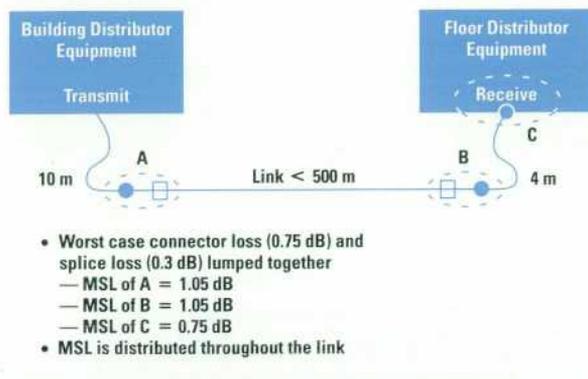


Figure 7

Worst-case ISO/IEC 11801 mode-selective loss (MSL) model. Fiber-optic connectors are represented by solid circles and optical splices by open squares. ISO/IEC 11801 allows two worst-case connectors and splices per link. For a worst-case mode-selective loss model all losses are assumed to be mode-selective. Patch cords at the transmit and receive ends of the link are assumed to be 10 m and 4 m long, respectively, to maximize modal noise. Total link mode-selective loss is 2.85 dB.



length specified for many lower-data-rate LANs. **Figure 6** shows that a modal bandwidth of ≈ 1200 MHz · km is needed to support the 2-km 62MMF link length at 1.250 Gbaud.

Allocation for Mode-Selective Loss (Modal Noise): Theory. Many standards (ATM Forum, Fibre Channel, Serial HIPPI) contain power penalty allocations to allow for mode-selective loss. A modal noise theory¹⁰ has been developed and used to predict the worst mode-selective loss allocation for all these standards.

In addition, an ad hoc industry group (Hewlett-Packard, Honeywell, IBM, VIXEL), sometimes called the *modal noise test methodology group*, developed an initial mode-selective loss power penalty measurement test procedure. A PC-based simulation tool developed by HPL Bristol implements the theory of reference 10 and the tool has been accepted by the group for calculating worst-case power penalties. Recently, a draft test procedure was transferred to the TIA FO 6.5 committee for standards development.

Distributed Mode-Selective Loss. An important conclusion from the original theory¹⁰ and the work of the modal noise test methodology group was that mode-selective loss is distributed throughout a fiber-optic link.

Theory predicts that mode-selective loss close to the transmitter will usually generate the most modal noise. However, even for the worst-case link model of ISO/IEC 11801 (**Figure 7**), only 2.85 dB of loss can be placed near the transmitter for short link lengths. To maximize modal noise, a 10-m patch cord is assumed at the transmit end and a 4-m patch cord at the receive end of the link. The initial 10-m patch cord ensures that both low-frequency and high-frequency modal noise is present.¹⁰ For the calculation, the worst-case loss of the connector and splice at each end of the link are lumped together. The resulting 1.05 dB of loss is assumed to be totally mode-selective loss and the minimum separation between the two 1.05 dB mode-selective loss points is 4 m. An additional 0.75 dB of mode-selective loss is assumed to be present at the connection to the optical receiver. The total amount of mode-selective loss is 2.85 dB.

Figure 8 shows the calculated power penalties as a function of link length for ISO/IEC 11801 links for both short-wavelength and long-wavelength lasers. The calculations assumed three laser modes having relative intensities of 0.1, 1, and 0.1, a linewidth of 5 GHz (laser modulated) for each mode, and a mode partitioning factor (k) of 1. It is clear from the theoretical model that the worst-case power penalties for the short-wavelength 50MMF and the long-wavelength 62MMF are equal. This is primarily because both multimode fiber transmission systems support the same number of modes at these respective wavelengths.

Figure 8

Calculated power penalties for 850-nm 50MMF and for 1300-nm 62MMF, laser diodes, and the worst-case ISO/IEC 11801 model.

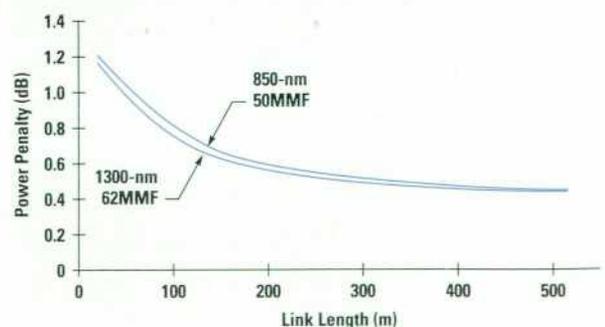


Figure 9

Calculated worst-case power penalty for 50MMF, 850-nm laser links as a function of rms source width using the modal noise test methodology group model.

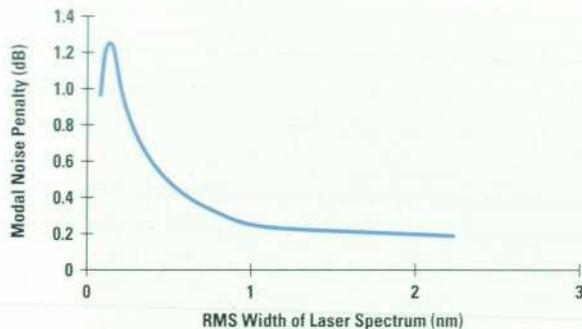
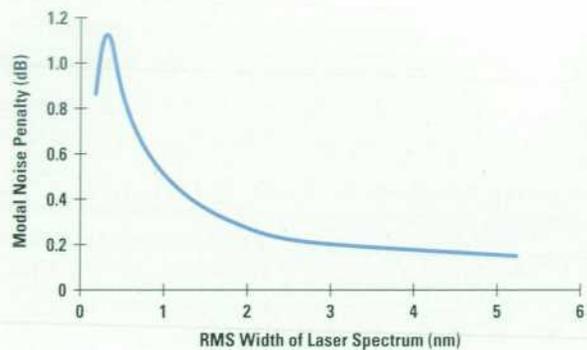


Figure 10

Calculated worst-case power penalty for 62MMF, 1300-nm laser links as a function of rms source width using the modal noise test methodology group model.



Modal Noise Test Methodology Group Worst-Case Link

Model. To ensure that a reasonable worst-case link is analyzed and experimentally tested, the ad hoc modal noise test methodology group assumed that three 1-dB points of mode-selective loss separated by 4 m are placed 12 m from the transmitter output connector. The distance of 12 m ensures that the link has enough bandwidth for both high-frequency and low-frequency modal noise to be present and close to their maximum levels.¹⁰

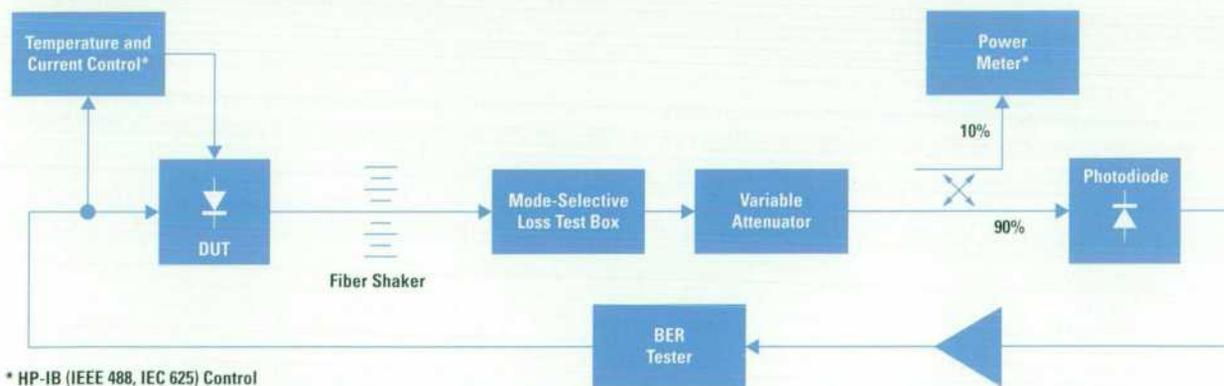
Predicted worst-case power penalties, according to the modal noise test methodology group model, for short-wavelength and long-wavelength lasers as a function of the laser rms spectral width are plotted in **Figure 9** and **Figure 10**. Maximum power penalties of approximately

1 dB are predicted as specified by ATM Forum, Fibre Channel, and Serial HIPPI standards for operation at either wavelength.

Modal Noise: Experimental Results. Modal noise testing has been concentrated on a selection of Hewlett-Packard low-complexity, 1300-nm coaxial lasers, which are expected to produce worst-case modal noise performance. All tests were computer-controlled, as depicted in **Figure 11**. The modal noise test box contains the three points of mode-selective loss and a reference path as required by the draft modal noise test procedure. During the testing the fiber was mechanically agitated and the temperature of the laser under test was continuously ramped. Although

Figure 11

Computer-controlled modal noise power penalty measurement setup used for modal noise testing.



near worst-case long-wavelength lasers have been tested, the maximum power penalty observed to date is 0.3 dB. This is consistent with the predicted maximum power penalty of 0.6 dB for the coaxial lasers. A set of modal noise test results is shown in **Figure 12**.

The modal noise data summarized in this section led to consensus within the IEEE 802.3z committee that the HP dual technology strategy was the best option for Gbit/s Ethernet.

Extended Link Lengths: Restricted Mode Launch

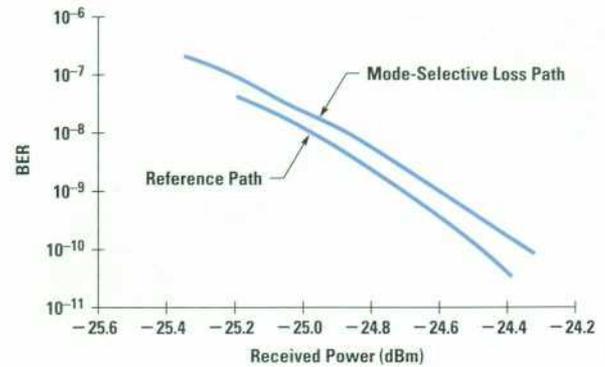
In contrast to short-wavelength transceivers, long-wavelength transceivers easily meet the 500-m link length needs for building backbones even with worst-case 62MMF bandwidths. However, achieving link lengths of 2 km as required for campus backbones requires new technical developments even at wavelengths of 1300 nm. For this reason, there is interest in using restricted mode launch to increase the bandwidth-distance product of multimode fiber. Such techniques are applicable to both short-wavelength and long-wavelength operation and to both 62MMF and 50MMF systems. The TIA has started a task group, TIA FO 2.2, to investigate restricted mode launch into multimode fiber. If the work of TIA FO 2.2 is successful, the IEEE 802.3 and IEEE 802.12 Gbit/s LANs are expected to incorporate restricted mode launch into their specifications.

Promising results have been presented to TIA FO 2.2 using various restricted mode launches (see **Figure 4**) to increase the bandwidth of multimode fiber. Unfortunately, the restricted mode launch effect with installed fibers is not completely understood and many questions need to be answered (installation effects, mechanical stability, connector effects, etc.) before LAN specifications using restricted mode launch can be developed.

Preliminary data reported to TIA FO 2.2 has indicated that the modal bandwidth can also be reduced below overfilled launch values with some forms of restricted mode launch. Even when pure single-mode launch is used, connector effects have been observed to significantly reduce modal bandwidth. These results indicate that some forms of restricted mode launch may not be a good solution for currently installed fiber.

Figure 12

Measured modal noise power penalty of a low-complexity, coaxial, 1300-nm laser using the ad hoc modal noise group test methodology. The power penalty is $\ll 1$ dB.

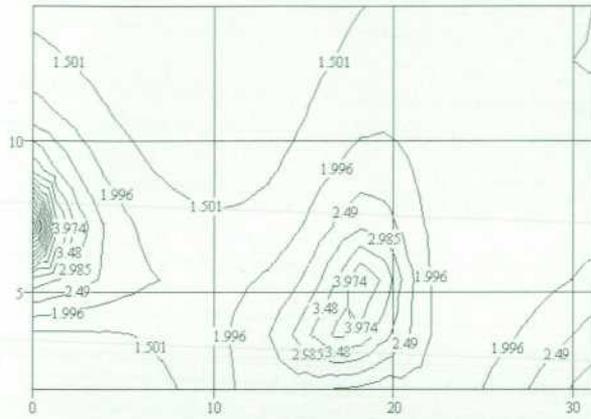


Restricted Mode Launch Theory. Modern graded-index multimode fibers have approximately square-law refractive index profiles. The modal field distribution can therefore be modeled, at least to a first approximation, as the field distribution of the modes of a square-law medium. In addition, it is well established that the WKB method¹¹ can be used to calculate the delay time of the modes. For power-law refractive index profiles, simple analytical expressions for the delay exist.¹¹

Using the analytical expressions for the field distributions and the delay times, it is possible to calculate the rms impulse width of a multimode fiber. In addition, other relevant parameters such as the coupled power can be estimated. **Figure 13** shows a plot of the bandwidth gain (restricted mode launch bandwidth divided by overfilled launch bandwidth) for single-mode excitation of an approximately square-law 62MMF as a function of the offset of the single-mode beam from the center of the fiber core and as a function of the single-mode beam 1/e waist. There is a large gain in bandwidth for center launch when the single-mode excitation has a waist equal to the waist of the fundamental mode of the 62MMF. Surprisingly, there is also a gain peak for offsets near one-half the fiber radius. The 62MMF bandwidth again increases for large offsets. While the intermediate and center launch would

Figure 13

Bandwidth gain contours for single-mode excitation compared to overfilled launch fiber bandwidth, at a wavelength of 1300 nm, for approximately square-law 62MMF. The x axis is the offset from the center of the fiber core in μm and the y axis is the beam waist in μm of the single-mode exciting beam.



generally exhibit high coupling efficiencies, the larger offsets would suffer from substantial power loss.

Experimental Results: Single-Mode Fiber Center Launch. We have investigated restricted mode launches of low-complexity, 1300-nm coaxial lasers into 62MMF. The fiber used for the experiments had an overfilled launch bandwidth-distance product of $638 \text{ MHz} \cdot \text{km}$. The output of a

Figure 14

Measured eye diagram at 1.25 Gbaud with 1300-nm, back-to-back transceivers

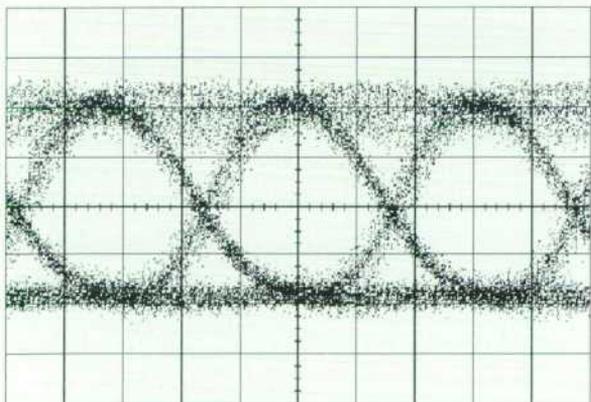
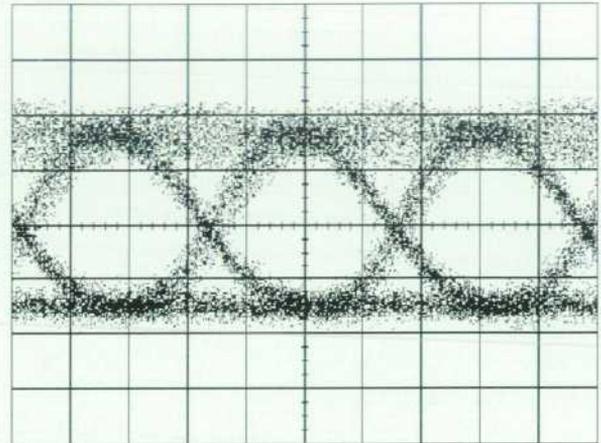


Figure 15

Measured eye diagram at 1.25 Gbaud with 1300 nm transceivers separated by 2 km of 62MMF.

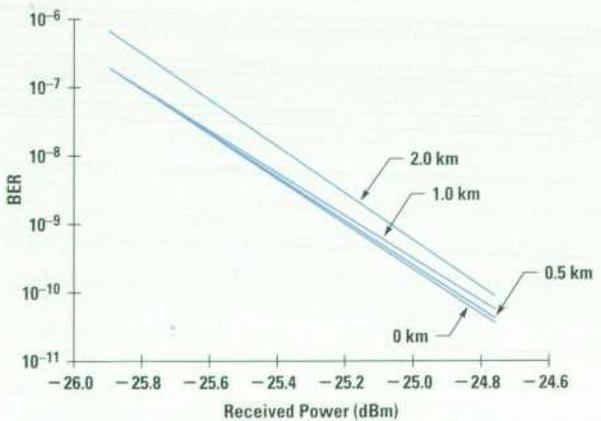


1300-nm coaxial laser module with an SC connector was connected directly to various lengths of 62MMF, center launched. Each fiber length was made up by concatenation of 500-m reels of cable. **Figure 14** and **Figure 15** show the measured eye diagrams of zero and 2-km length links. Clearly, the eye is open to distances of 2 km. **Figure 16** plots the measured power penalties for various link lengths up to 2 km. The power penalty at 10^{-10} BER.

Figure 16

Figure 16

BER curves for 0-km, 0.5-km, 1-km, and 2-km link lengths.



However, with some fibers, particularly fibers with center-line refractive index defects, small connector offsets of the order of 5 micrometers result in very low bandwidth compared to overfilled launch. The use of restricted mode launch results in links much shorter than those obtained with overfilled launch for these fibers.

Experimental: Offset Restricted Mode Launch. A surprising prediction of the theoretical model (see **Figure 13**) is the occurrence of three bandwidth gain peaks. Experimental measurements of the bandwidth gain versus offset for a single-mode fiber-to-62MMF launch are plotted in **Figure 17**. The wavelength of the Fabry-Perot laser diode used was 1300 nm. The measured refractive index profile of the 62MMF had no dips or peaks and was approximately square-law. Three gain peaks were observed near the theoretically predicted offset values.

Although offset single-mode launches are restricted, for 62MMF and offsets in the range of 15 μm to 25 μm , the output intensity distribution is only slightly restricted compared to overfilled launch. This favorable excitation of the multimode fiber results in bandwidths equal to or greater than the overfilled launch bandwidth. We have observed that at least overfilled launch bandwidth can be achieved with offset launch on all fibers tested to date.

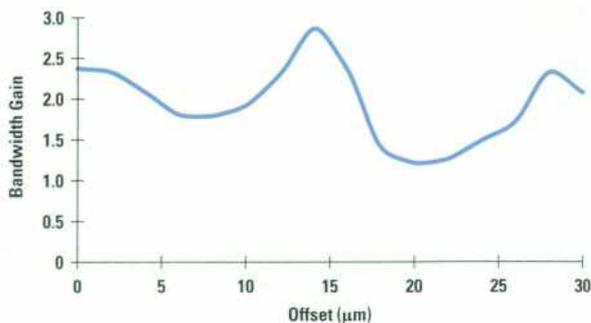
We have also investigated the modal noise penalties, connector effects, and mechanical stability of offset launches. No problems have been encountered in these areas. However, it is not clear that offset launches could be easily integrated into HP's manufacturing processes. HPL continues to investigate both center launch and offset launch to determine which type of launch is best for various LAN and interconnect systems.

Gbit/s IEEE 802.3 Link Specification Status

The July 1996 IEEE 802.3z Gbit/s Ethernet plenary meeting agreed to adopt two optical interface specifications. One interface is defined for 850-nm lasers to achieve link lengths of ≈ 250 m with 62MMF and ≈ 500 m with 50MMF. This gives up some extended link length performance to achieve the lowest possible cost. The second interface is defined to meet the installed 62MMF building backbone link lengths with 1300-nm transceivers, which achieve 3 km with single-mode fiber and ≈ 500 m with both 62MMF and 50MMF. Because the bandwidth achieved with center launched transceivers may be less than the overfilled launch bandwidth, IEEE 802.3z has defined its

Figure 17

Measured bandwidth gain versus offset for single-mode fiber-to-62MMF launch at a wavelength of 1300 nm. Three gain peaks are observed as predicted in



link lengths using a worst-case modal bandwidth. If conditioned launches that guarantee overfilled launch bandwidth are used, the IEEE 802.3z draft standard allows longer link lengths.

Conclusion

CSSD's and HPL's involvement in embryonic standards provides valuable insight that allows us to adapt our standards and research strategy to increase the probability of it being accepted. Our involvement in the ATM Forum OC-12 specification led us to refocus our VCSEL development to operation at wavelengths near 850 nm rather than 980 nm. This was because multimode fiber suppliers were opposed to specifying the higher multimode fiber modal bandwidth available at 980 nm and because all other laser transceiver suppliers were 850-nm focused. Our longer-term research was directed towards long-wavelength VCSELs, which have multiple uses with both multimode fiber and single-mode fiber.

CSSD and HPL have successfully championed long-wavelength LEDs for 622-Mbit/s ATM links. We also have shown, experimentally and theoretically, that long-wavelength lasers have similar modal noise power penalties to short-wavelength lasers when used with multimode fiber. For Gbit/s Ethernet, this enabled us to introduce a dual technology solution: short wavelength vertical-cavity surface emitting lasers for extended horizontal links and long-wavelength lasers for backbone links. The intriguing discovery that restricted mode launch can extend the data rate-distance performance of multimode fiber LAN links

without increased modal noise penalties could revolutionize laser/multimode fiber system design. For this reason CSSD and HPL continue to investigate restricted mode launch.

References

1. K.H. Hahn, M.R. Tan, Y.M. Houn, and S.Y. Wang, "Large-Area Multitransverse-Mode VCSELs for Modal Noise Reduction in Multimode Fiber Systems," *Electronics Letters*, Vol. 29, no. 16 1993, pp. 1482-1484.
2. M.R.T. Tan, K.H. Hahn, Y.M.D. Houn, and S.Y. Wang, "Surface Emitting Lasers for Multimode Data Link Applications," *Hewlett-Packard Journal*, Vol. 46, no. 1, February 1995, pp. 67-71.
3. D.C. Hanson, *Roadmap for Extending Distance and Data Rate Limits of 622-Mbit/s 1300-nm LED Multimode Fiber Links*, ATM Forum Subworking Group: Physical Layer, Document 94-1021, 1994.
4. IEEE Standard 802.3u-1995, *CSMA/CD Access Method, Type 100Base-T*.
5. IEEE Standard 802.12, *Demand Priority Access Method, Physical Layer and Repeater Specification for 100-Mbit/s Operation*.
6. D.G. Cunningham, *Fibre Channel-Based PMD for Gbit/s IEEE 802.12—Initial Observations*, Contribution to IEEE 802.12, Montreal, November 1995.
7. D.G. Cunningham, *Fiber and Transceiver Technology for Gbit/s IEEE 802.12*, Contribution to Higher-Speed IEEE 802 Tutorial, La Jolla, March 12, 1996.
8. D.C. Hanson, *IEEE 802.3 1.25-Gbaud MMF Link Specification Development Issues*, Contribution to IEEE 802.3z, Enschede, The Netherlands, July 8-12, 1996.
9. D.G. Cunningham, M.C. Nowell, and D.C. Hanson, *1.25-Gbaud, 550-m Links on Installed 62MMF for IEEE 802.3: Leveraging Existing Long-Wavelength Specifications*, Contribution to IEEE 802.3z, Enschede, The Netherlands, July 8-12, 1996.
10. R.J.S. Bates, D.M. Kuchta, and K.P. Jackson, "Improved Multimode Fiber Link BER Calculations due to Modal Noise and Non Self-Pulsating Laser Diodes," *Optical and Quantum Electronics*, Vol. 27, 1995, pp. 203-224.
11. D. Marcuse, *Light Transmission Optics, Second Edition*, Van Nostrand Reinhold, New York, 1982.



David G. Cunningham

David Cunningham is a departmental scientist and project manager

with HP Laboratories Bristol, England (HPLB). He has been managing and leading HPLB efforts for the physical layer of the Gigabit Ethernet standard. He received his PhD degree in laser physics from The Queens University, Belfast in 1985. Before joining HP in 1987, he worked on integrated optics at British Telecom Research Laboratories. He is married and has a son



Delon C. Hanson

Del Hanson is a principal engineer at HP's Communications Semiconductor

Solutions Division. A 30-year HP employee, Del is currently focused on fiber-optic network standards development and strategic product planning. He has published some 20 articles on microwave subsystems as well as fiber-optic components and networks. He is a native of Baldwin, Wisconsin, has been married for 38 years and has three children. Fishing, golfing, and gardening are among his outside interests.



Mark C. Nowell

Mark Nowell is a member of the technical staff of HP Laboratories Bristol, England. He is responsible for optical

physical layer work and did the experimental measurements necessary to support the Gigabit Ethernet standards effort. He received his PhD degree from Cambridge University in 1994, specializing in optoelectronics. He is married, plays squash, and enjoys playing with his two children.



C. Steven Joiner

A principal engineer at HP's Communication Semiconductor Solutions

Division (CSSD), Steve Joiner is a member of a group of engineers called Pathfinders by CSSD, who are responsible for the development of product definitions. With HP since 1978, he received his PhD degree in Physics from Rice University in 1979. He is married, has three sons, and is active in youth organizations and education.

1300-nm Strained Quantum Well Lasers For Fiber-Optic Communications

William S. Ring

Simon J. Wrathall

Adrian J. Taylor

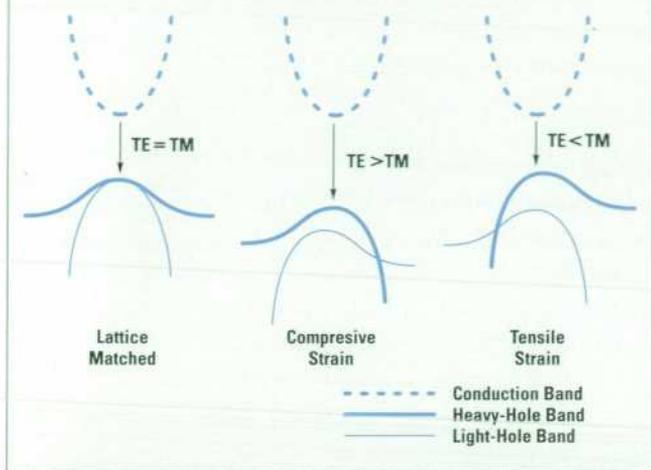
This paper describes new uncooled strained quantum well lasers for SONET/SDH systems. New Fabry-Perot lasers for short-haul and intermediate link applications are extremely reliable, have high ex-facet power, and have record low threshold currents, making lower packaging costs possible. Uncooled distributed feedback lasers for the long-haul market at 622 Mbits/s and 2.488 Gbits/s are discussed. These operate from -40°C to $+85^{\circ}\text{C}$ with extremely good threshold and power characteristics.

In recent years there has been increased emphasis on cost and performance issues for telecommunications lasers. The drive to push fiber to the curb (FTTC) and fiber to the home (FTTH) can only be fully realized when inexpensive laser transmitter modules become available.

The drive for lower cost has forced the move to uncooled semiconductor lasers, which require high output powers and low threshold currents over the temperature range of -40°C to $+85^{\circ}\text{C}$, combined with inexpensive packaging technology. The type of laser used depends primarily on the link length and operating speed. The Synchronous Digital Hierarchy (SDH/SONET/ATM) standards pertaining to the use of either multimode or single-mode lasers specify the basic requirements for spectral purity and operating wavelength range appropriate to the fiber characteristics. They also dictate whether a multimode, Fabry-Perot laser or a single-mode, distributed feedback (DFB) laser is used in the system.¹ For most short-haul and intermediate link length applications, the STM-1 (155 Mbits/s) and STM-4 (622 Mbits/s) standards dictate that a Fabry-Perot laser should be suitable. For long-haul applications a distributed feedback laser is recommended for wavelength stability and dispersion. As systems move to higher speeds, such as STM-16 (2.5 Gbits/s) the emphasis moves to distributed feedback lasers, even for shorter link lengths.

Figure 1

Diagram showing how strain splits the two degenerate valence bands in a III-V semiconductor and how the sign of the strain governs whether the TE (transverse electric) or TM (transverse magnetic) gain is dominant.



For the STM-1 and STM-4 short-haul and intermediate links, 1300-nm Fabry-Perot lasers are the main devices employed. For operation without a thermoelectric cooler there are difficulties in designing a laser that is temperature insensitive and has high output power. It has been known for some time that long-wavelength telecommunication lasers suffer from an inherent temperature sensitivity problem that is not observed in short-wavelength, GaAs-based devices. This is believed to result from the small bandgap and the associated nonradiative and absorptive processes of Auger recombination and inter-valence band absorption, which then become significant at longer wavelengths.

Fabry-Perot 1300-nm Strained Quantum Well Lasers

It was predicted theoretically in 1986 that introducing strained layers in the active region of long-wavelength lasers should improve the temperature performance.^{2,3} The physics behind this is that the strain introduced in the crystal lattice splits the degeneracy of the two valence bands in the semiconductor, enhancing either TE (transverse electric) polarization optical gain or TM (transverse magnetic) polarization optical gain (see **Figure 1**). The strain can be accommodated elastically if the total thickness of the strained layer is less than the critical thickness of the material system. Critical thickness is the point at which the elastic energy in the layer is equal to the energy

required to introduce dislocations into the crystal lattice. The value of the critical thickness was originally calculated by Mathews and Blakeslee⁴ and is equivalent to a strain-thickness product of approximately 10 nm% for InGaAsP/InP (e.g., a 10-nm thickness with 1% strain). For practical device purposes, only thin layers, that is, strained quantum well devices, can realistically be grown.

The first realization of strained quantum well devices occurred for 1550-nm strained quantum well lasers in 1989 and 1990 with the demonstration of increased ex-facet power (power out of the laser chip into a broad-area detector) and reduced threshold current. A high characteristic temperature (T_0) of 90K was reported.⁵ Initial effort mainly concentrated on 1550-nm devices for long-haul links to improve the power budget. Meanwhile, work on 1300-nm quantum well devices has highlighted another factor that leads to increased temperature sensitivity: loss of carriers out of the quantum well.^{6,7}

Design Issues

The most common material system used to grow 1300-nm and 1550-nm long-wavelength lasers is the InGaAsP/InP system. This material system suffers from the fact that the ratio of the conduction band to the direct bandgap is 0.4, compared to an estimated 0.6 to 0.85 for GaAs-based systems. For 1300-nm quantum well lasers this leads to an energy difference between the electron confined state and the barrier material on the order of 1 to 2 $k_B T$ or 25 to 50 meV at room temperature. This can lead to a significant excitation of carriers out of the quantum well or a large carrier density in the barrier material at elevated temperatures. A large carrier density in the barrier material degrades the performance of the laser, causing increased temperature sensitivity and reduced ex-facet power.

An alternative material system is the InGaAlAs/InP system, in which the ratio of the conduction band to the direct bandgap is 0.7. This gives increased electron confinement. The electron confined state energy can then be designed to be 4 to 5 $k_B T$ or 100 to 150 meV at room temperature,⁸ significantly larger than the InGaAsP system. Unfortunately, because of the aluminum content and associated reliability issues, only *ridge laser structures* are presently fabricated using this material system. To achieve low threshold currents, *buried heterostructure lasers*, which have better electrical and optical confinement, are preferred.

We have investigated 1300-nm strained InGaAsP/InP active regions to try to optimize the device performance for low threshold current and high ex-facet power from -40°C to $+85^{\circ}\text{C}$. The following section outlines the main results from our experiments, which led to the development of a highly successful 1300-nm quantum well laser.

Optimization of Confined State and Barrier Energy

The energy difference between the electron confined state in the quantum well and the barrier layer can be calculated using a crystal energy band modeling method⁹ (model solid method) combined with a k.p effective mass one-dimensional square well potential model. The k.p method is a band structure modeling technique by which the effective mass of the electrons and holes is calculated from the energy gaps of the semiconductor.

The model solid method calculates the band edge energies relative to the vacuum level using a density functional mathematical method. These can then be used to calculate the conduction band and valence band energy offsets, that is, the depth of the quantum wells. These, combined with the effective mass and its energy dependence from k.p theory, are used to find the positions of the energy states in the quantum well.

To optimize the energy difference, the well, barrier composition, and well width are varied. To illustrate the enhancement that can be achieved using a strained quantum well compared to a lattice matched quantum well structure, we have plotted in **Figure 2** the energy difference as a function of well width. For the equivalent emission wavelength of 1300 nm there is an improvement in energy difference for the strained layer structure. This, coupled with (1) the smaller carrier density in the quantum well due to increased differential gain and (2) the polarization discrimination of the strain layers, improves the device characteristics.

Fabrication

To understand the design limits for the Fabry-Perot laser device, we grew several wafers with the same active-region strain, but different numbers of quantum wells and different guide layer (separate confinement heterostructure) thickness. The devices were grown using atmospheric-pressure metal organic chemical vapor deposition (MOCVD) in a horizontally configured Thomas Swan reactor. The strain in the quantum wells was fixed at 0.8% and the well width was 6 nm. All wafers were fabricated into

our standard buried heterostructure, which uses a p-n-p-n current blocking structure to confine the current to the active region.

Internal Loss

The external ex-facet efficiency is related to the internal optical absorption loss in the cavity by the following expression, originally given by Biard.¹⁰

$$\frac{1}{\eta_d} = \frac{1}{\eta_i} + \frac{\alpha_i L}{\eta_i \ln(1/R)}, \quad (1)$$

where L is the cavity length, R is the facet reflectivity, α_i is the internal absorption loss, η_i is the internal quantum efficiency, and η_d is the external differential quantum efficiency. Using equation 1 and plotting the inverse of the differential efficiency against the cavity length we can obtain a value for the internal loss of the laser cavity.

In a quantum well laser, when there is no intervalence-band absorption, the internal loss per quantum well should be linear with the number of wells, since the internal losses do not depend on carrier density. We plotted the calculated internal loss for devices with different numbers of wells and found a linear dependence (see **Figure 3**). The lowest loss obtained was 4.5 cm^{-1} for four quantum wells. This equates to an ex-facet slope efficiency of 0.42 mW/mA for a 350- μm -long chip, with $\eta_i = 0.98$. Higher output powers can be obtained if a high-reflectivity coating is then applied to one of the facets, because the low internal absorption loss allows more

Figure 2

Calculated energy difference between the electron confined state in the well and the barrier energy for strained and lattice matched material. For the equivalent emission wavelength from the well, the strained multiple quantum well laser has a higher barrier height.

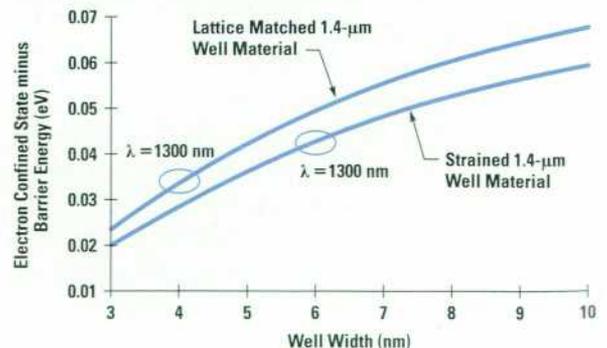
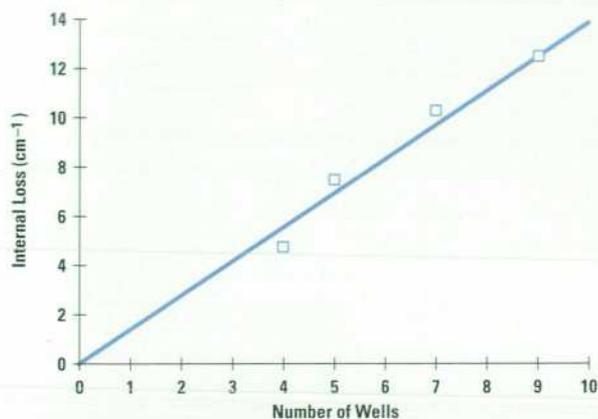


Figure 3

Measured variation in the internal absorption loss for 0.8% strained multiple quantum well lasers as a function of the number of wells at 25°C. This demonstrates the negligible effect of intervalence band absorption in 1300-nm strained material.



photons to travel in one round trip in the cavity and escape out of the uncoated facet. Typically the uncoated cleaved facet efficiency can be 0.6 mW/mA for a cleaved and high-reflectivity coated device.

Gain-Current Curves at 25°C

A useful tool for understanding the performance of the laser is a gain-current curve. This can be obtained experimentally by combining the measured internal loss with the variation of the threshold current density per unit length. Normally the gain of a semiconductor laser can be approximated by the following logarithmic expression, which takes into account the saturation of the gain spectrum at high current densities:

$$G = G_0 \left[\ln \left(\frac{\eta_i J_{thr}}{J_{tr}} \right) + 1 \right], \quad (2)$$

where G is the modal gain and is related to the material gain g by the optical confinement factor Γ (the overlap of the optical field with the active material), that is, $G = \Gamma g$, η_i is the internal quantum efficiency, J_{thr} is the threshold current density in A/cm^2 , J_{tr} is the transparency current density in A/cm^2 , and G_0 is the modal gain coefficient. If we rearrange equation 2 and plot $\ln(J_{thr})$ versus the inverse of the length, the slope of the graph gives the modal gain coefficient G_0 and the intercept is related to the

transparency current density J_{tr} . The modal gain varies with the number of quantum wells in the structure (see **Figure 4**). The smaller the number of wells the faster the gain rolls off with increasing current density. For 9-well and 7-well material the gain-current curves are very similar at 25°C, but as we will see later from the temperature dependence of the threshold current, a difference becomes apparent, especially when trying to optimize for good 85°C characteristics.

Threshold Current between 25°C and 85°C

The relationship between threshold current and temperature for semiconductor lasers is commonly characterized by an exponential function, which is used to associate a characteristic temperature T_0 with the device structure.¹¹ The expression is:

$$J_{th}(T_2) = J_{th}(T_1) \exp(\Delta T/T_0), \quad (3)$$

where J_{th} is the threshold current or threshold current density and $\Delta T = T_2 - T_1$ is the temperature difference. This expression is essentially valid only above 250K, when the threshold current begins to increase rapidly with temperature.

The target threshold current for the product was a maximum of 30 mA at 85°C. Different active region structures were characterized over temperature and a strong correlation of the threshold current with the number of quantum wells at elevated temperatures was observed (see **Figure 5**). For 5-well structures the threshold current

Figure 4

Measured gain-current curves for 9-, 7-, 5-, and 4-well active regions at room temperature. The smaller the number of wells the flatter the gain curve at higher current density.

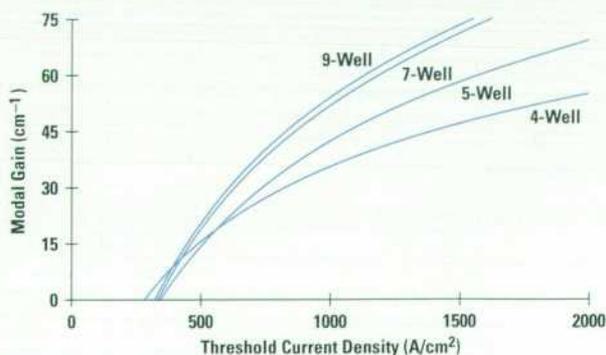
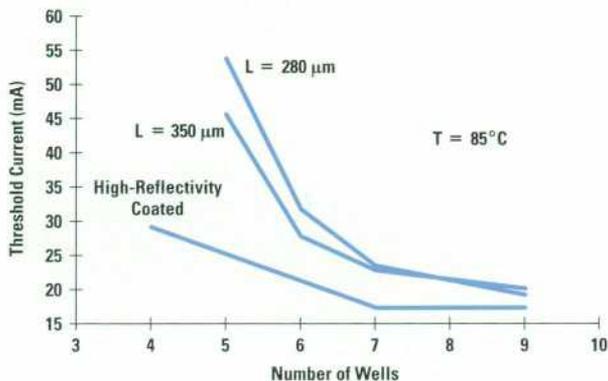


Figure 5

Plotted is the trend in threshold current for 280- μm , 350- μm , and cleaved high-reflection coated devices for different numbers of quantum wells in the active region at 85°C. The smaller the number of wells the worse the temperature performance due to increased escape from the quantum wells.



increased from 5 mA to 44 mA for 350- μm -long devices ($T_0 \approx 28\text{K}$). Meanwhile, 9-well devices showed increases from 5 mA to only 18 mA ($T_0 \approx 48\text{K}$) over the same temperature range. To understand the role played by the optical guide layer thickness on either side of the quantum wells, we compared 5-well and 7-well material (see **Figure 6**).

Figure 6

Temperature performance of 5-well and 7-well lasers at 25°C and 85°C as a function of waveguide (separate confinement heterostructure, SCH) thickness. This demonstrates that the number of quantum wells dominates and hence carrier density per well has a more dominant effect than the waveguide width.

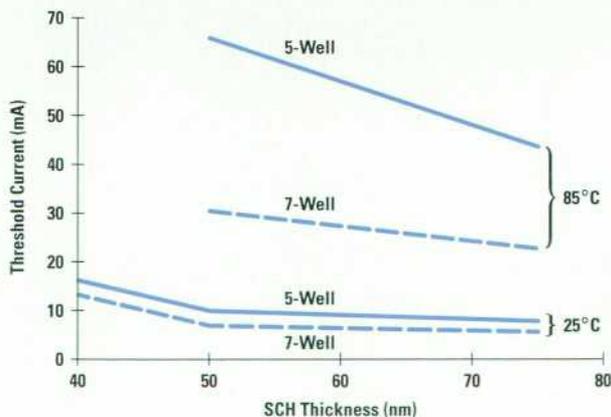
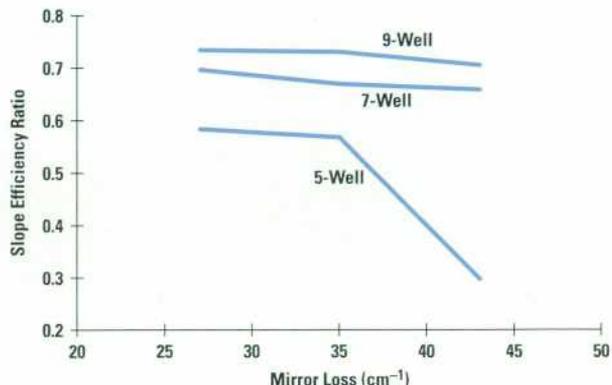


Figure 7

The ratio of the slope efficiency between 85°C and 25°C for 9-, 7-, and 5-well devices is plotted as a function of mirror loss. The main observation is the breakpoint observed for the 5-well device, indicating that the internal quantum efficiency has passed a critical point.



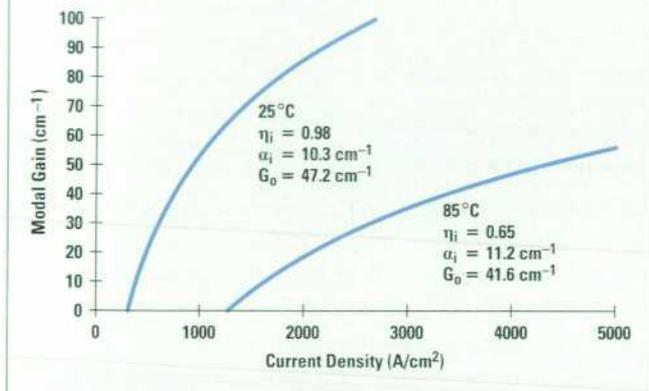
It was very noticeable that the threshold current increased more dramatically for 5-well material than for the 7-well material for the same guide layer. This showed us that the current density, or equivalently, the carrier density per quantum well, was the dominant mechanism affecting the temperature sensitivity. This indicated that the excitation out of the well material into the barrier layer was a primary loss mechanism at elevated temperatures for InGaAsP-based lasers at 1300-nm wavelength. To further understand the significance of this effect we compared the external differential efficiency between 25°C and 85°C.

Differential Efficiency between 25°C and 85°C

The change in the external differential efficiency is characterized using equation 1. The change in internal loss and internal quantum efficiency can be used as a tool to understand the dominant temperature dependent loss mechanism. When we measured the slope efficiency for the different number of wells and compared the ratio of the 85°C slope efficiency divided by the 25°C slope efficiency, a definite trend became apparent. As the device becomes shorter, the mirror loss of the device increases. When the slope efficiency ratio is plotted as a function of mirror loss (See **Figure 7**), it is evident that 9-well and 7-well devices show a linear dependence, which is expected because the internal loss or quantum efficiency

Figure 8

Modal gain-current curve for a 7-well laser measured at 25°C and 85°C. The internal loss α_i only increased from 10.3 cm^{-1} to 11.2 cm^{-1} while the internal injection efficiency η_i varied considerably, falling from 0.98 to 0.65 between the two temperatures.



did not vary. For the 5-well device a breakpoint is observed, which is consistent with either a rapid increase in internal loss or a rapid degradation in the internal quantum efficiency. In a laser above threshold it is assumed that the internal efficiency is equal to 1, but in the method of Biard,¹⁰ η_i is not equal to 1 because it contains contributions of leakage current, excitation out of the quantum well, and other effects. The evidence suggests that at a critical carrier density, excitation out of the well dominates the laser threshold current and the slope efficiency. Further evidence to support this was found when we constructed a gain-current curve for 7-well devices at both temperatures (see Figure 8). The internal absorption loss was found to increase only from 10.3 cm^{-1} to 11.2 cm^{-1} , but η_i decreased from 0.98 to 0.65, showing a marked reduction in electrical-to-optical conversion efficiency.

Reliability

The reliability requirement for telecommunications products is a field lifetime of approximately 25 years.¹⁰ To validate whether the active region design is suitable, an accelerated aging test is performed and predictions for the median and lower tenth percentile lifetimes are calculated. The standard testing regime is to hold the device at 85°C and monitor the change in threshold current with the device driven approximately at rated power.

As initially pointed out, the introduction of strain in the active region of the device can be accommodated elastically only if the total thickness is less than the critical thickness for the material. Dislocation-free material is ideal for a laser device, but typically a defect density of less than $10^4/\text{cm}^2$ is sufficient for operation. For a strained multiple quantum well stack it has been shown that a factor of 4 to 6 times the critical thickness can be accommodated before dislocation formation is so severe as to degrade the device performance. A rule of thumb is a maximum critical thickness of $4 \times 10 \text{ nm}\%$ for a strained multiple quantum well stack.

Dislocation formation in the active region can lead to hot spots and leakage current paths that degrade the device performance over its life. For 9-well devices the estimated strain-thickness product is $9 \times 0.8 \times 6 \text{ nm}$ or $43 \text{ nm}\%$, which is over the maximum limit. The 7-well structure is $33 \text{ nm}\%$ and under. Strain compensation can be introduced by having the opposite mismatch in the guide layers. This is not equivalent to growth of zero-net-strain structures, but can partially compensate the overall strain-thickness product. Both 9-well and 7-well wafers were life tested for 5000 hours and exhibited linear degradation rates less than 0.75% per 1000 hours (see Figure 9). Median lifetimes were predicted assuming a power-law dependence at

Figure 9

Accelerated aging test at 85°C for 9-well and 7-well devices. The line indicates a degradation rate of 0.75% per khr. Both types of device show very good reliability, with predicted lives of $2.6 \times 10^5 \text{ hr}$ at 85°C.

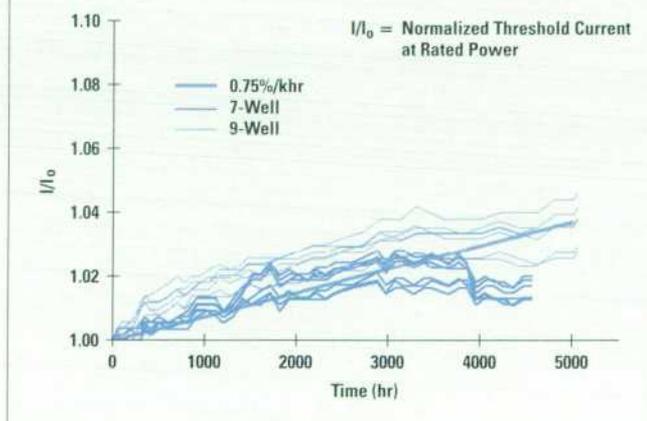
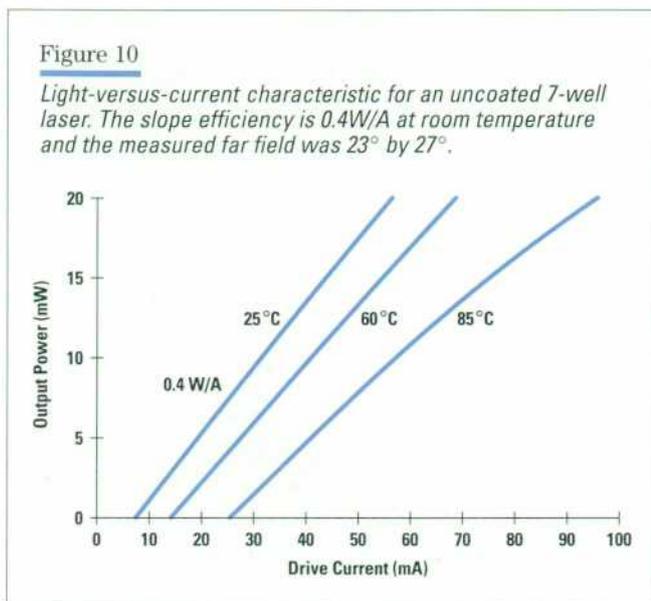


Figure 10

Light-versus-current characteristic for an uncoated 7-well laser. The slope efficiency is 0.4W/A at room temperature and the measured far field was 23° by 27°.



85°C. The median life for 7-well lasers was estimated at 2.6×10^5 hr (the failure criterion is median life $< 40,000$ hr at 85°C). The 25°C median life can be calculated, assuming an activation energy of 0.6 eV to convert the accelerated test to its room temperature value. The estimated median life is 1.9×10^7 hr or 2100 years, which is suitable for all telecommunications system applications.

Uncoated 5-well laser devices were also life tested at 85°C, but because of their extremely poor temperature performance, they exhibited large changes in threshold current within the first 600 hours. The typical increase in threshold was approximately 10%. This was unacceptable from both a threshold current and a reliability standpoint and the design was rejected on that basis.

Fabry-Perot Laser Device Conclusions

The main conclusions from the chip investigation above show that the dominant factor affecting the high temperature operation of the 1300-nm InGaAsP/InP buried heterostructure laser is the carrier density per quantum well. To achieve good reliability, high ex-facet power, and low threshold currents, an optimized design of a 7-well device was chosen. The number of quantum wells dominated over the guide layer thickness, allowing us to achieve a nearly circular far field of 27° by 23° by modifying the guide layer thickness at some cost in temperature performance (see Figure 10). This approach has been beneficial from a manufacturing and packaging viewpoint. The

existing coaxial laser package has recently been superseded by a new welded version that incorporates a ball lens, offering a significant cost reduction compared to the previous package. The lower internal loss and high ex-facet power achieved with a cleaved high-reflectivity coated chip makes it possible to achieve 2 mW out of the fiber without the use of expensive aspheric lenses. The combination of chip and package now becomes a low-cost solution that has manufacturing, marketing, and business advantages.

Uncooled 1300-nm Distributed Feedback Lasers

Uncooled distributed feedback lasers are emerging in the marketplace to replace cooled transmitters. The arena they address is the long-haul SDH/SONET specifications for 622 Mbits/s and 2.488 Gbits/s, where the dispersion and the window of operation cannot be achieved by a Fabry-Perot laser because of its broad spectral width and wavelength temperature coefficient. The principal difference between the Fabry-Perot laser and the distributed feedback laser is the addition of a grating layer in the device, which controls the fundamental operating wavelength. The difference in the characteristics of Fabry-Perot and distributed feedback lasers are highlighted in Table I.

Table I

Operating Characteristics of Fabry-Perot and Distributed Feedback Laser Diodes

Characteristic	Fabry-Perot	Distributed Feedback
Wavelength	Gain Peak	Bragg Mode
Wavelength Temperature Coefficient	0.3 to 0.5 nm/°C	< 0.1 nm/°C
Spectral Width (rms)	1.7 to 2.5 nm	< 1 nm at -20 dB
Side Mode Suppression	Multimode Operation	> 30 dB

Distributed feedback lasers have some inherent problems achieving single-mode operation over the required operating temperature range of -40°C to $+85^\circ\text{C}$. In the next sections, we will outline the main areas of concern and show how they can be tackled to produce a more manufacturable product.

Single-Mode Behavior

To gain insight into the basic principles of operation of distributed feedback lasers, we need to derive the coupled wave equations and then apply some approximations. This will highlight a significant problem found in conventional index-guided distributed feedback lasers, that is, single-mode yield (the percentage of devices that operate in one distributed feedback mode).

The introduction of a grating layer into the device structure produces a periodic variation in both the real and imaginary (gain) parts of the refractive index, which can be represented by the following equations:

$$n(z) = n_0 + \Delta n \cos(2\beta z + \phi) \quad (4a)$$

$$\alpha(z) = \alpha_0 + \Delta \alpha \cos(2\beta z + \phi), \quad (4b)$$

where n_0 is the mean refractive index, α_0 is the mean loss or gain, Δn and $\Delta \alpha$ are the perturbations from the mean, β is the propagation constant for the waveguide, and ϕ is a phase factor.

Substituting these into the wave equation and assuming a small perturbation of the traveling wave, that is, $\alpha \ll \beta$, $\Delta n \ll n$, and $\Delta \alpha \ll \alpha$, we obtain the following result for the propagation constant k of the perturbed wave:

$$k^2 = \beta^2 + 2j\beta\alpha_0 + 4\beta[k_0\Delta n + j\Delta\alpha]\cos(2\beta z + \phi), \quad (5)$$

where k_0 is the free-space propagation constant. The expression in brackets is defined as the real and imaginary parts of the grating coupling coefficient and is used to describe the interaction of the grating with the waveguide:

$$[k_0\Delta n + j\Delta\alpha] = \kappa = \kappa_r + j\kappa_i, \quad (6)$$

where κ is the complex coupling coefficient and defines the grating strength. On substitution of equation 5 back into the wave equation and assuming a trial solution of the E-field of the form:

$$E(z) = R(z)e^{-j\beta z} + S(z)e^{j\beta z}, \quad (7)$$

we arrive at the standard coupled mode representation to describe the forward $S(z)$ and backward $R(z)$ traveling waves in the cavity:

$$-dR/dz + (\alpha - j\delta)R = j\kappa S \quad (8a)$$

$$dS/dz + (\alpha - j\delta)S = j\kappa R. \quad (8b)$$

The new term δ is defined as the detuning of the main distributed feedback mode from the Bragg mode,* that is,

$$\delta = \beta - \beta_0 = 2\pi n_{\text{eff}}L(\nu - \nu_0)/c,$$

where β is the propagation constant of a distributed feedback mode, β_0 is the propagation constant of the Bragg or stop-band mode, n_{eff} is the effective refractive index, L is the cavity length, ν is the frequency, and c is the velocity of light. To simplify things further, we can neglect the influence of the facets by assuming perfect antireflective coatings. The traveling waves then become sinh functions in terms of the complex propagation constant γ :

$$R(z) = \pm \sinh\gamma(z + L/2) \quad (9a)$$

$$S(z) = \mp \sinh\gamma(z - L/2), \quad (9b)$$

where

$$\gamma^2 = (\alpha - j\delta)^2 - \kappa^2.$$

Substitution of $R(z)$ and $S(z)$ into equations 8a and 8b then produces the following general equations:

$$\gamma + (\alpha - j\delta) = j\kappa e^{\gamma L} \quad (10a)$$

$$\gamma - (\alpha - j\delta) = j\kappa e^{-\gamma L}. \quad (10b)$$

Invoking the high-gain approximation, that is, $\kappa \ll \alpha$, and inserting this into equation 10a, we obtain:

$$2(\alpha - j\delta) \approx \pm j\kappa \exp(\alpha - j\delta)L. \quad (11)$$

The condition for the resonant modes of the cavity is found by taking the phase of this expression, which gives:

$$\delta L \approx (q + 1/2)\pi + \tan^{-1}(\kappa_i/\kappa_r), \quad (12)$$

where q is the mode number. Thus, we obtain:

$$(\nu - \nu_0)(c/2n_{\text{eff}}L)^{-1} \approx q + 1/2 + (1/\pi)\tan^{-1}(\kappa_i/\kappa_r). \quad (13)$$

Neglecting the facet reflectivity and facet phase contributions, the above expressions sum up the intrinsic problem with achieving single-mode operation in distributed feedback lasers. First, the resonances are approximately $c/2n_{\text{eff}}L$ apart, which is similar to the behavior of a Fabry-

* The distributed feedback mode is the grating mode with the highest gain, the one that achieves lasing action first. The Bragg mode is the zeroth-order diffraction mode, defined by the pitch of the grating.

Perot resonator. Secondly, for the normal type of distributed feedback structure which uses predominantly index coupling (κ real), there is NO resonance at the Bragg mode. Instead, the condition is satisfied only for two degenerate modes on either side of the stop band. To achieve single-mode operation the phase condition should equal π when $q = 0$. This occurs for pure imaginary coupling when $\kappa_i > 0$ and $\kappa_r = 0$ or when a $\lambda/4$ phase shift is introduced into the grating.

Temperature Performance

The temperature performance of a distributed feedback laser is governed by several factors, including the detuning of the distributed feedback mode from the material gain peak, the strength of the grating, and the temperature performance of the active region. To achieve operation from -40°C to $+85^\circ\text{C}$, we need to look at the fundamental picture of what happens to the laser gain spectrum and its relative strength and position with respect to the distributed feedback mode.

The SDH standards require a spectral purity corresponding to greater than 30 dB of side mode suppression, which requires a minimum modal gain difference of $\sim 5 \text{ cm}^{-1}$ between the dominant Fabry-Perot mode* and the distributed feedback mode. This is normally achieved by coating the facets of the distributed feedback laser to increase the mirror loss of the Fabry-Perot mode relative to the distributed feedback mode. We can understand this from the gain condition equation of a distributed feedback laser. The general gain condition for a distributed feedback laser is derived from the coupled mode equations by inserting the correct boundary conditions at the facets. We then obtain:

$$\frac{(r(\beta) - \sqrt{R_1} e^{j\phi_1}) (r(\beta) - \sqrt{R_2} e^{j\phi_2})}{(1 - r(\beta) \sqrt{R_1} e^{j\phi_1}) (1 - r(\beta) \sqrt{R_2} e^{j\phi_2})} \exp(\alpha L) = 1, \quad (14)$$

where $r(\beta)$ is the grating reflectivity, which depends on propagation constant β , R_1 and R_2 are the facet reflectivities, ϕ_1 and ϕ_2 are the facet phases, α is the gain difference, and L is the cavity length. When $r(\beta) = 0$ we regain the Fabry-Perot laser round-trip gain condition as expected. Rearranging this into a more common form:

* Fabry-Perot modes are resonant cavity modes of the laser, i.e., $\lambda = 2m\pi/L$.

$$G = \alpha_i + \frac{1}{L} \ln \left[\frac{(1 - r(k) \sqrt{R_1} e^{j\phi_1})(1 - r(k) \sqrt{R_2} e^{j\phi_2})}{(r(k) - \sqrt{R_1} e^{j\phi_1})(r(k) - \sqrt{R_2} e^{j\phi_2})} \right], \quad (15)$$

where α_i is the internal absorption loss. We see that, essentially, the grating introduces an effective difference in the mirror loss between the distributed feedback mode and the Fabry-Perot mode that is a function of cavity length, facet reflectivity, facet phase, and grating coupling strength. This can now be used to estimate the gain difference between the Fabry-Perot mode and the distributed feedback mode (see **Figure 11**) as a function of the grating strength. To get a handle on the lower limit of grating strength to make the distributed feedback laser operate over the required temperature range we need to measure the modal gain spectrum for the various operating temperatures at fixed proportions of threshold and then compare this with mirror loss curves of the distributed feedback laser.

The modal gain of a 7-well 1300-nm Fabry-Perot laser was measured using the Hakki-Poali method¹² for three temperatures (see **Figure 12**). It shows that the three gain spectra cross at the short-wavelength side about 15 nm from the 25°C gain peak, and 20 cm^{-1} down from the gain

Figure 11

Calculated gain difference between the distributed feedback mode and the Fabry-Perot mode as a function of grating coupling coefficient κ . The Fabry-Perot mirror loss is 68 cm^{-1} and for a distributed feedback laser to work over the full temperature range a minimum κ of 60 cm^{-1} is required.

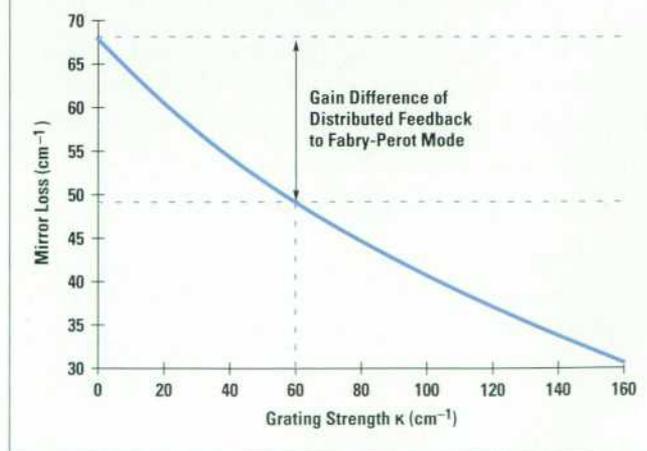
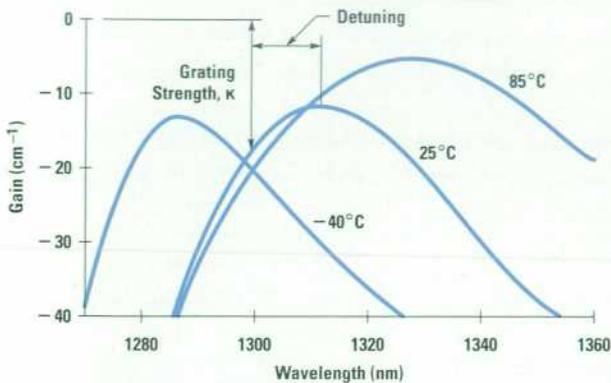


Figure 12

Measured gain spectra for a Fabry-Perot laser from -40°C to $+85^{\circ}\text{C}$ using the Hakki-Poali method. The three curves cross approximately 20 cm^{-1} from the gain equals loss condition ($G=0$) and 15 nm from the 25°C gain peak. This is used to estimate the minimum grating coupling strength required for full-temperature-range operation.



equals loss condition (i.e., $G=0$). The curve closest to the gain condition is the 85°C spectrum which has a peak of -5 cm^{-1} , similar to the value needed for 30-dB side mode suppression. This indicates that for the distributed feedback mode to lase over this temperature range we have to impose a mirror loss difference of 20 cm^{-1} , which equates to a minimum grating coupling coefficient κ of 60 cm^{-1} . In conventional index-guided distributed feedback lasers the facet reflectivity and facet phase introduce some asymmetry in the gain of the two degenerate distributed feedback modes, producing single-mode devices. Previous analysis of the single-mode yield by many groups has shown that it drops off rapidly as the grating coupling strength increases, primarily because of the reduction in gain difference between the two stop-band modes as the facet contribution gets smaller.¹³

To overcome this problem we need to introduce some gain coupling to enhance the single-mode yield. The phase condition, equation 12, shows that there is a contribution of $\tan^{-1}(\kappa_i/\kappa_r)$. When κ_i is zero the phase condition implies that the distributed feedback laser has two modes, but the introduction of pure gain coupling, that is, $\kappa_i > 0$ and $\kappa_r = 0$, allows the phase matching condition to be satisfied at the Bragg mode, which has the highest gain to propagate in the structure. Hence to achieve full-temperature-range operation of a distributed feedback

laser, we must use a high coupling coefficient and introduce gain coupling to enhance the single-mode yield.

Device Results

The planar structure for the active region was based on our Fabry-Perot strained quantum well design. The grating was n^{++} doped to block the current periodically and thereby induce a current modulation in the active region. First-order gratings were defined by holography and dry etching. The devices were overgrown with approximately 300 nm of InP and then mesas were defined and wet chemically etched. Two further overgrowths were performed for realization of the standard current blocking structure. After metallization the chips were coated with either a combination of an antireflection coating and a high-reflection coating or only one facet was coated with an antireflection coating. The chips were tested from -40°C to $+85^{\circ}\text{C}$ in submodule packages with angle-ferrule connectors to minimize the optical return loss into the chip. The threshold current and spectra were characterized. The low threshold current achieved at 85°C was 26 mA for a full-temperature-range gain-coupled device (see Figure 13).

Investigation of the detuning of the distributed feedback mode with respect to the gain peak at 25°C showed that if the device was detuned to the longer-wavelength side by more than 10 to 15 nm, low-temperature operation was not achievable (see Figure 14), even though low threshold currents of only 18 mA could be achieved. Estimates from

Figure 13

Light-current characteristic for a gain-coupled distributed feedback laser showing the excellent temperature characteristics that can be produced.

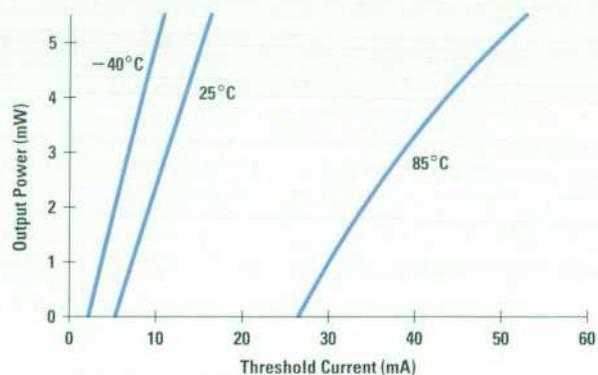
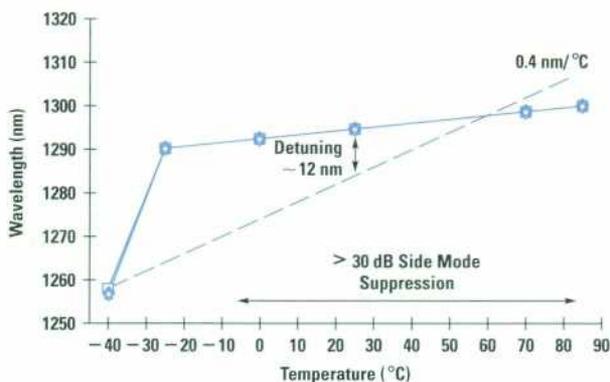


Figure 14

Wavelength temperature measurement of a positively detuned distributed feedback laser showing that the side mode suppression ratio drops below 30 dB at temperatures lower than -10°C and the laser becomes a Fabry-Perot laser at -40°C . This shows a limit on the positive detuning for successful low-temperature operation.



the detuning experiment revealed that effectively there was a 10-nm window for full-temperature-range operation with devices requiring 0-nm to -10 -nm detuning to operate over the required range.

A major concern with distributed feedback lasers is the sensitivity to external reflections. Telecommunications distributed feedback lasers operating at 1550 nm normally employ an optical isolator in the package to prevent reflections that would perturb the optical signal. The main reason for this enhanced sensitivity compared to Fabry-Perot lasers is the antireflection coating on the front facet. With the antireflection coating, reflected light can easily perturb the internal optical field profile of the laser, causing either double-mode behavior, reduced side mode suppression, or noise on the data 1s (i.e., the on state).

To achieve low cost, the package should not use an optical isolator. Instead, we designed the grating structure to reduce the sensitivity. The SONET/SDH specifications require a transmitter device to operate with a maximum discrete reflectance of -25 dB and a maximum return loss of -20 dB from the optical plant or system. A mated connector should give around -40 dB return loss, while a connector to air gives -14 dB, a worst-case scenario. We varied the grating coupling strengths to calibrate the reflection sensitivity of our buried heterostructure distributed feedback lasers. The reflection sensitivity measurements showed that at relatively moderate grating strengths

of 50 cm^{-1} , the distributed feedback spectrum was perturbed to optical return losses of more than -20 dB. For grating strengths (κ) of 90 cm^{-1} and 120 cm^{-1} the spectrum remained unperturbed even up to -14 dB optical return loss. This allowed us to gauge the minimum strength required in our design.

This need for high coupling coefficients for full-temperature-range operation and immunity to reflections is useful from a manufacturing standpoint because prohibitively expensive integral optical isolators are no longer required. The only problem that remains is the single-mode yield, which can be improved by the introduction of gain coupling.

Typically, the single-mode yield in our strained quantum well distributed feedback laser depends approximately on the number of wells in the structure. Extremely high single-mode yields of 70% to 80% have been achieved for 7-well devices with grating coupling coefficients (κL) of 2.5 to 3.5. These are extremely good results; typically, single-mode yields are much smaller than this at such high coupling coefficients. We observed a decrease in the yield as more quantum wells were introduced into the structure, with the single-mode yield being only 38% to 44% for 9-well structures for equivalent grating strengths. The final design required a trade-off between well number and single-mode yield to provide both good packaged temperature performance and high power.

Conclusion

We have discussed aspects of device design necessary to produce uncooled strained quantum well lasers for SONET/SDH/ATM systems. To cover long-haul, intermediate, and short-haul specifications, two distinct types of laser structure are required: distributed feedback and Fabry-Perot lasers. The short-haul and intermediate link applications can be covered by Fabry-Perot lasers which are simple to fabricate with a high manufacturing yield and are extremely reliable.

The long-haul market at 622 Mbits/s and 2.488 Gbits/s is addressed by our uncooled distributed feedback lasers. The basic difference is the mode selectivity introduced by the grating. The main problem for distributed feedback lasers is achieving operation over the temperature range of -40°C to $+85^{\circ}\text{C}$. We showed how this can be done by the use of high grating coupling coefficients and gain coupling to increase the single-mode yield. The successful

operation of our distributed feedback lasers from -40°C to $+85^{\circ}\text{C}$ with extremely good threshold and power characteristics shows these will play a key role in the successful deployment of FTTH and FTTC.

Acknowledgments

The authors would like to thank all the members of the device technology group at Ipswich for their support and hard work. Special thanks should be given to Herbert Lage for his contribution on the strained quantum well material growth and valuable discussions, to Paul Charles and Andrew McKee for the transition of the Fabry-Perot quantum well laser into production, to Ian Smith for all his hours working on characterization of the Fabry-Perot laser, and to Antony Hawkrigde for reliability predictions and qualification data. Special thanks to Conrad Langton for the development of the semicircular high-speed buried heterostructure trenching process, which led to a reduction in stress effects, allowing the push forward to 2.488 Gbit/s operation. Finally, thanks to Richard Ash and Sean Amos for their support through both programs.

References

1. ITU-T Recommendation G.957, ITU-T Study Group XV, 1988-93.
2. A.R. Adams, "Band Structure Engineering for Low Threshold High Efficiency Semiconductor Lasers," *Electronics Letters*, Vol. 22, 1986, pp. 249-250.
3. E. Yablonovich and E.O. Kane, "Reduction of the Lasing Threshold Current Density by Lowering the Valence Band Effective Mass," *Journal of Lightwave Technology*, Vol. LT-4, May 1986, pp. 504-506.
4. J.W. Mathew and A.E. Blakeslee, "Defects in Epitaxial Multilayers," *Journal of Crystal Growth*, Vol. 27, 1974, pp. 118-125.
5. P.J.A. Thijs, F.L. Tiemeijer, J.J.M. Binsma and T. van Dongen, "Progress in Long Wavelength Strained-Layer InGaAs(P) Quantum Well Semiconductor Lasers and Amplifiers," *IEEE Journal of Quantum Electronics*, Vol. 30, no. 2, 1994, pp. 477-499.
6. W.S. Ring, H. Lage, A.J. Taylor, I.S. Smith, and R. M. Ash, "Optimisation of highly efficient uncoated strained 1300nm InGaAsP MQW Lasers for uncooled high temperature operation," paper ThG4, *Conference on Optical Fiber Communication, Technical Digest*, Vol. 2, 1996, p. 230.
7. V. Mikhaelashvili, N. Tessler, R. Nagar, G. Eisenstein, A.G. Dentai, S. Chandrasakhar, and C.H. Joyner, "Temperature Dependent Loss and Overflow Effects in Quantum Well Lasers," *IEEE Photonics Technology Letters*, Vol. 6, no. 11, 1994, p. 1293.
8. C.E. Zah, R. Bhat, B.N. Pathak, F. Favire, W. Lin, M.C. Wang, N.C. Andreadakis, D.M. Hwang, M.A. Koza, T.P. Lee, Z. Wang, D. Darby, D. Flanders, and J.J. Hsieh, "High-Performance Uncooled 1.3- μm $\text{Al}_x\text{Ga}_y\text{In}_{1-x-y}\text{As}/\text{InP}$ Strained Layer Quantum Well Lasers for Subscriber Loop Applications," *IEEE Journal of Quantum Electronics*, Vol. 30, no. 2, 1994, pp. 511-523.
9. C.G. Van De Walle, "Band Line Ups and Deformation Potentials in the Model Solid Theory," *Physical Review B*, Vol. 39, no. 3, 1989, pp. 1871-1883.
10. J.R. Biard, W.N. Carr, and B.S. Reed, "Analysis of a GaAs Laser," *Transactions of the Metallurgical Society of the AIME*, Vol. 230, 1964, pp. 286-290.
11. J.I. Pankove, *IEEE Journal of Quantum Electronics*, Vol. QE-4, no. 4, 1968, p. 119.
12. B.W. Hakki and T.L. Poali, "Gain Spectra in GaAs Double Heterostructure Injection Lasers," *Journal of Applied Physics*, Vol. 46, March 1975, pp. 1299-1306.
13. K. David, G. Morthier, P. Vankwikelberge, R.G. Baets, T. Wolf and B. Borchert, *IEEE Journal of Quantum Electronics*, Vol. QE-27, 1991, p. 1714.



William S. Ring

Bill Ring is a principal engineer and project leader with HP's Ipswich,

England Components Operation, responsible for development of 1300-nm strained quantum well and uncooled distributed feedback laser chips. He is an expert in III-V semiconductor laser theory and his 1992 PhD dissertation was on the efficiency of strained and unstrained light-emitting devices. Before joining HP in 1994 he was a research officer with Optronics Ireland.



Simon J. Wrathall

Simon Wrathall is a senior technician presently working on device char-

acterization, test jigs, and test software at HP's Ipswich, England Components Operation. He holds a BSc degree in physics and astrophysics and is a native of Ipswich.



Adrian J. Taylor

Adrian Taylor is a senior engineer in charge of the epitaxial growth area at

HP's Ipswich, England Components Operation. He holds an MSc degree in instrumentation and analytical science (1986) from the University of Manchester Institute of Science and Technology. Born in Redhill, Surrey, UK, he is married and has a daughter. Before joining HP he was with GEC for seven years.

Modeled Optimization and Experimental Verification of a Low-Dispersion Source for Long-Haul 2.488-Gbit/s Systems

Stephen M. Gee

Herbert Lage

Chris Park

Kevin A. Williams

Richard V. Penty

Ian H. White

Joseph A. Barnard

This paper describes microwave, laser, and fiber models that were used in the development of the HP LSC2500 2.488-Gbit/s laser diode module. Knowledge of the modeled behavior of the laser diode as a function of the input electrical pulse shape has led to deliberately shaping the input pulse to give the minimum wavelength excursion during direct modulation, and therefore a high yield of low-dispersion-penalty laser diodes. These devices can be successfully used for transmission distances in excess of 200 km.

As digital modulation speeds increase to meet increasing demands on capacity, and as transmission lengths become longer with the availability of narrow-linewidth sources and optical amplifiers, a requirement for low-dispersion-penalty laser diodes has emerged. *Dispersion penalty* is caused by the wavelength dependence of propagation speed in optical fiber, resulting in the different wavelength components of the optical pulse traveling at different speeds along the fiber and arriving at the receiver with variable delay. In extreme cases, optical power generated during one bit period can arrive at the receiver during the adjacent bit period, causing errors, a problem exacerbated in high-frequency systems with short bit periods. Even in truly single-mode laser diodes, when modulated directly by applying an electrical signal to the laser bias, the optical signal can be composed of a range of wavelengths because the wavelength itself is bias dependent, a property known as *laser chirp*. More advanced laser diodes with integral modulation sections are becoming available, but even these have some wavelength modulation caused by optical effects within the diode.

Since the penalty caused by laser chirp is dependent on drive conditions, packaging, and laser chip properties, we have developed a model that can take the input electrical data and predict the electrical signal reaching the laser diode, taking into account package effects. The model then predicts the output optical pulse from the laser based on laser chip parameters. A fiber model is added to predict the optical pulse shape after a known length of fiber.

This model has been used to develop the HP LSC2500, an optically isolated, 1550-nm, 2.488-Gbit/s, distributed feedback (DFB) laser module. The model has helped with the development of both the packaging and the laser chip.

The model was developed jointly by the HP Ipswich Components Operation (ICO), Bristol University, and Barnard Microsystems Limited. The experimental comparison has been carried out at ICO. Models have been developed for other laser structures such as electroabsorption modulators (EADFB). These models will help with the design of components for higher frequencies (10 Gbits/s).

This paper includes predicted results from the modeling and actual results measured on the finished product. Sections in this paper introduce the microwave model, the laser model, and the fiber model. A description is given of the laser diode. Experimental results compare the measured and predicted microwave s-parameters and also the measured and predicted optical pulses both before and after the fiber. Dispersion penalty measurements are included, showing that directly modulated DFB lasers with low chirp can, in practice, operate over fiber transmission distances in excess of 200 km.

DFB Laser Chip Design

The design of the distributed feedback (DFB) laser diode is based upon ICO's proven ridge waveguide technology. The laser has a strained multiple quantum well active region grown by metalorganic vapor-phase epitaxy and a first-order DFB grating. The front and back facets are anti-reflection and high-reflection coated, respectively. The chip is mounted with the active region uppermost on a silicon carbide heat sink.

The quaternary material system consisting of combinations of the four elements indium (In), gallium (Ga), arsenic (As) and phosphorus (P) is used for the growth of the epitaxial layers. The number of quantum wells in the active region and the strain—that is, the mismatch of the

lattice constant of the quantum wells with respect to the InP substrate—have been optimized for high-speed performance and fast response to changes in drive current. In particular, the shift of the emission wavelength under modulation (chirp) has been minimized to reduce the dispersion penalty in long-haul fiber-optic links. The compressive strain of the quantum wells is compensated by opposite strain in the barriers to avoid any impact of the strain on the reliability of the devices. The thickness of each strained layer is below the critical limit at which relaxation occurs.

Mode selectivity is achieved by a uniform first-order grating with a coupling coefficient of approximately 30 to 40 cm^{-1} . This relatively modest coupling efficiency was chosen to keep the photon density uniform inside the cavity and to prevent distortion of the light-versus-current characteristic. The antireflection and high-reflection coatings are applied to improve the suppression of side modes and to enhance the front-facet efficiency of the devices.

Typical device parameters are summarized in **Table I**. The low internal loss and the high gain of the strained multiple quantum well active region results in average threshold currents of only 16 mA. The front-facet efficiency of 0.35 mW/mA is an excellent achievement for DFB lasers in the 1550-nm wavelength range, where lasers typically suffer from low internal quantum efficiencies. The linewidth is measured under 2.5-Gbit/s modulation with 12% extinction ratio (power in the off state divided by power in the on state) and found to be around 0.4 nm (with the reading taken 30 dB below the peak). This shows the good wavelength stability and therefore low chirp of the device under intensity modulation and proves its suitability

Table I
Laser Diode Parameters

Parameter	Symbol	Value
Length	L	350 μm
Threshold Current	I_{th}	16 mA
Front-Facet Efficiency	η	0.35W/A
30-dB Linewidth without Modulation		0.38 nm
30-dB Linewidth with Modulation		0.42 nm

for high-speed long-haul data transmission systems. The small broadening of the linewidth under modulation is in agreement with a direct measurement of the chirp.

Laser Driver Circuit Model

The 2.488-Gbit/s pulses into the laser diode transmitter package can be generated by a signal generator, which can be either a pseudorandom binary sequence (PRBS) nonreturn-to-zero (NRZ) waveform generator test set or a laser driver chip. The laser driver chip can be based on a fast silicon bipolar technology such as the Hewlett-Packard HP-25 process, or on a GaAs FET technology.

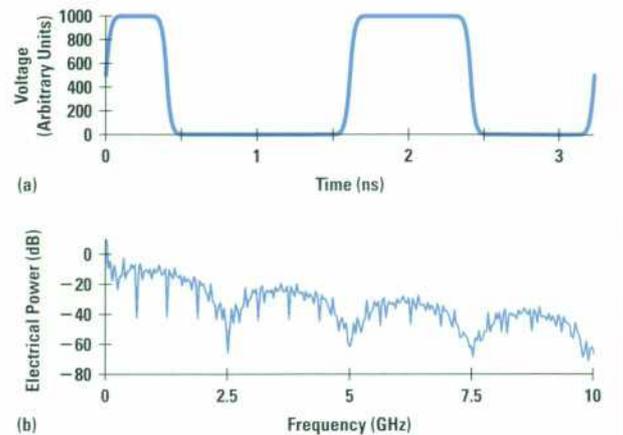
Initially we need to know the frequency range that we should be concerned with. We describe the 2.488-Gbit/s waveform with a 75-ps 10%-to-90% rise time through the use of a Fermi function to define the rising and falling pulse edges. The use of this function, more usually applied in the area of semiconductor device physics, gives a reasonable approximation to the true waveform from real signal generators. We create a 128-bit PRBS waveform and take the fast Fourier transform (FFT) of the waveform to get an idea of the spectral distribution of the energy in the waveform.

A 2.488-Gbit/s signal was obtained from an HP 70841B PRBS signal generator and its spectrum was monitored on an HP 71400 lightwave signal analyzer. **Figure 1a** shows the predicted waveform and **Figure 1b** shows the predicted spectrum. The spectrum agreed well with our predictions. From the spectrum, we see that we should principally be concerned with the energy from very low frequencies up to about 2 GHz.

It is important to know the attributes of the driver circuit to account for any multiple reflections between the signal generator and impedance discontinuities in the laser package. If the signal generator is a PRBS test set, then the equivalent circuit for the source is well-approximated by a voltage source in series with a 50-ohm resistance. If a laser driver chip is used, the characterization of the source impedance is more complex, and can be divided into two cases. In the first case, in which we have all the information on the driver circuit (including small-signal and large-signal circuit models), the extraction of a relatively simple circuit describing the source impedance characteristics of the driver is straightforward. In the case in which we have very little or no information about the driver source impedance, we need to measure two sets of

Figure 1

(a) The predicted 2.488-Gbit/s signal obtained from an HP 70841B PRBS signal generator. (b) The spectrum predicted for the 2.488-Gbit/s signal.



s_{11} microwave scattering parameters looking into the driver output, using a microwave vector network analyzer such as the HP 8510. For both sets, the driver output bias level is set to the value to be used in practice. In the first set, we have the output modulation current (which is added to the bias current) set to zero, and in the second set, we have the output modulation current at the high value as used in practice. We can fit two equivalent circuits to the two sets of s-parameters. Finally, we merge the two circuit topologies to end up with one description for the source impedance of the driver.

From inspection of a Smith chart plot of s_{11} , one can tell whether the driver circuit is based on silicon bipolar or GaAs FET technology because of the telltale dispersion in the output conductance of GaAs FETs. In this example, we have fitted the equivalent circuit shown in **Figure 2** to describe the source impedance of a GaAs FET-based laser diode driver. The parallel RC (PRC) combination models the nonzero output conductance (finite output resistance) of the output transistors in the driver and the output capacitance of the transistors. The series RC (SRC) circuit combination describes the dispersion in the output conductance of the GaAs FETs as a function of frequency. The inductance (IND) models the bond wire from the chip to the carrier. The ideal transmission line (TLIN) models any transmission line between the measurement reference plane and the location of the bond wire. This last element

is important since one often encounters s-parameter measurement data that has not been completely deembedded.

Using a small-signal microwave circuit simulator with optimization capability, one can find the best set of parameters for the circuit elements that gives the best agreement between the measured s_{11} data and the s_{11} values predicted by the circuit simulator. Typical values for the circuit of Figure 2 for a GaAs FET-based driver, derived through the use of linear circuit optimization, are as follows:

$R_2 = 349\Omega$, $R_1 = 216\Omega$, $L_1 = 0.757$ nH, $C_2 = 1.45$ pF, $C_1 = 15.0$ pF, $Z_1 = 139\Omega$, $E_1 = 7.09$ degrees.

The inverse of the R_1C_1 product defines the frequency in radians at which the output resistance of the transistor changes from being R_2 to being equal to the parallel combination of R_1 and R_2 . In this example, this break frequency is just above 49 MHz. At frequencies well below 49 MHz, the output resistance is 349 ohms (R_2), whereas at frequencies well above 49 MHz, the output resistance of the source FETs is about 133 ohms. This change in output resistance of the driver transistors is referred to as the dispersion in the transistor output conductance (or resistance). The value of 7.09 degrees at 1.0 GHz for the TLIN indicates that there was some residual transmission line length that had not been removed during the microwave s-parameter measurement procedure.

In practice, the transistor can be considered to consist of a current source in parallel with the PRC circuit element in Figure 2. We can use either a large-signal, time-domain circuit simulator or a small-signal, frequency-domain circuit simulator to predict the shape of the driver output waveform. The manner in which a time-domain circuit

simulator can be used is relatively straightforward, so we will illustrate the use of a frequency-domain circuit simulator. If we use a standard 50-ohm characteristic impedance to predict the s-parameters for the driver circuit, we can replace the current source in parallel with R_2 by a voltage source in series with R_2 . The reason for replacing the current source with a voltage source is that the s-parameters are voltage ratio parameters, so we must use a voltage source. Since we use the standard 50-ohm resistance in series with the voltage source, we need only have $(R_2 - 50)$ ohms of series resistance.

We can define a driving PRBS waveform for the voltage source and use the FFT algorithm to transform the time-domain representation of the waveform into the frequency domain. We can use a linear microwave simulator to predict the real and imaginary parts of the voltage transfer ratio s_{21} and convolve the predicted s_{21} response with the frequency-domain representation of the voltage source waveform. Finally, we use the inverse FFT to transform the frequency-domain representation of the output pulse back into the time domain. We know from measurements using an HP 54120B sampling oscilloscope that the 10%-to-90% rise and fall time for the pulse from the driver circuit is 130 ps. Consequently, we can engineer the rise and fall time of the voltage source pulse such that the pulse obtained from the driver circuit model corresponds to that measured from the real driver circuit.

Figure 3 shows the driver circuit model obtained in this way along with the voltage source waveform and the driver circuit model output waveform. The waveform time axis extends from zero to 3.2 ns. The 10%-to-90% rise and fall time for the voltage source signal connected in series with a 50-ohm resistance to port 1 is 75 ps as generated by the PRBS test set. The model predicts a rise and fall time for the waveform from port 2 of 127 ps, which corresponds well with a measured value of around 130 ps.

Electrical Model for the Laser Diode

The next task is to create an electrical model for the distributed feedback laser diode. We first look at the construction of a model for use with a linear circuit simulator. From the S-shaped I-V characteristics of the laser diode, we find that in a manner analogous to that used in the Gummel-Poon model for the silicon bipolar transistor, we can simulate the dc I-V characteristics through the use of a resistance in series with two parallel diodes. We have

Figure 2

Equivalent circuit developed to describe the source impedance of a GaAs FET-based laser diode driver.

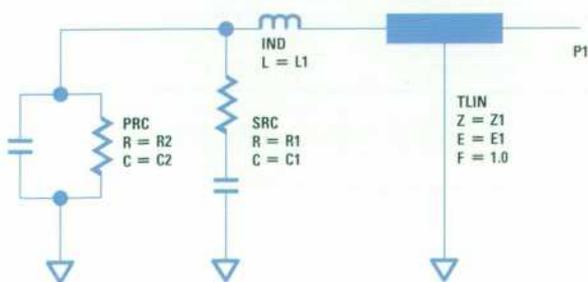
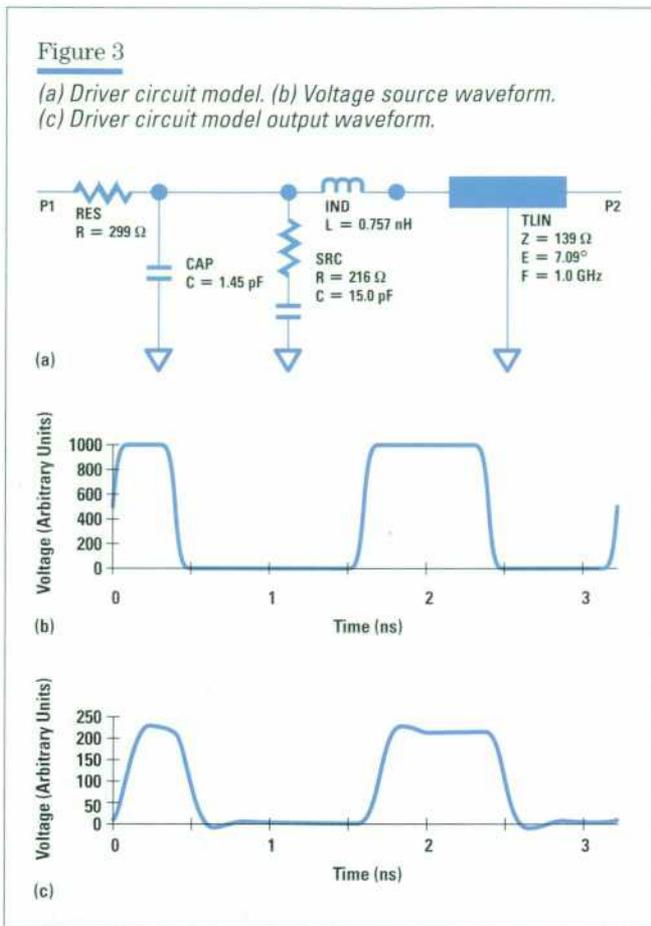


Figure 3

(a) Driver circuit model. (b) Voltage source waveform. (c) Driver circuit model output waveform.



written parameter extraction software for such a combination of diodes to find the parameters that best fit the classical exponential diode equation:

$$I = I_s(e^{aV/nkT} - 1).$$

In this example we found that the best-fit parameters are $I_{s1} = 1.1 \times 10^{-6}$ amperes and $n_1 = 3.9$ for the first diode, $I_{s2} = 1.0 \times 10^{-13}$ amperes and $n_2 = 1.18$ for the second diode, and $R_{series} = 5.5$ ohms.

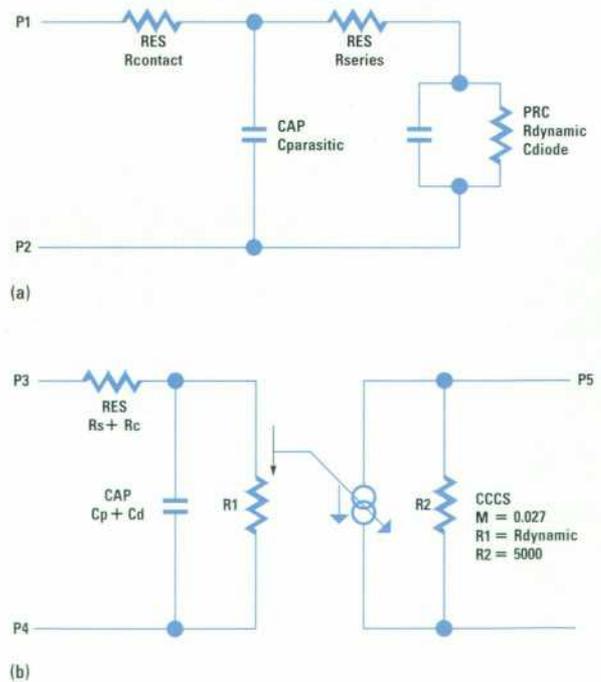
Diode 1, with $n_1 = 3.9$, dominates the diode characteristics for forward currents through the laser diode of less than 1 mA, whereas diode 2 dominates the characteristics at 1 mA and above. Since the laser diode is always biased in the on state with about 15 mA of current, we consider only the characteristics of diode 2. With the electrical characteristics of the DFB laser diode following the classical exponential diode equation, we calculate a dynamic laser diode resistance of 7.5 ohms when biased at 15 mA (the zero state close to threshold) and 6.2 ohms when

biased at 45 mA (the one state). Since the laser diode will spend about the same time in the zero state as it does in the one state, we take the arithmetic mean of the dynamic resistances to end up with 6.85 ohms.

In all modeling of physical devices, it is important to keep in mind the physics of the device being modeled. The small-signal model for the laser diode is shown in **Figure 4a**. Since the series resistance is very small in practice, we reduce the circuit complexity to that shown in **Figure 4b**, but add a current-controlled current source (CCCS) to monitor the current passing through the laser diode resistive path (that is, through the dynamic diode resistance, here shown as R1). The current amplification factor M is negative to enable us to monitor the magnitude of the current passing through the dynamic resistance. To get the voltage generated by the CCCS across the 50.0 ohms at the output port of the CCCS to equal the voltage across the

Figure 4

(a) Small-signal model for the laser diode. Since the series resistance is very small in practice, we reduce the circuit complexity to that shown in (b), but add a current-controlled current source (CCCS) to monitor the current passing through the laser diode resistive path (that is, through the dynamic diode resistance, here shown as R1).



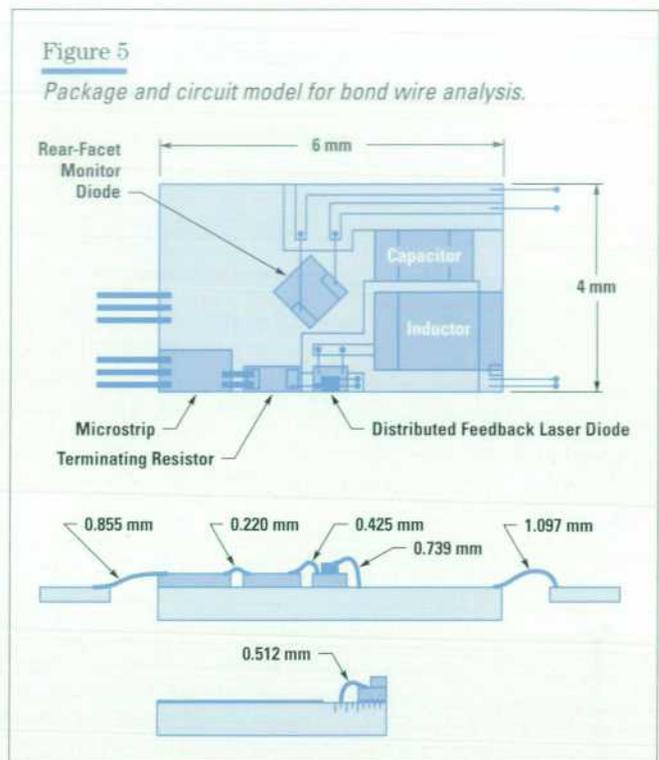
dynamic resistance ($R_{dynamic}$), we must have the magnitude of the CCCS current amplification factor equal to $R_{dynamic}/50.0$, that is, 0.027. We used the HP 8702 vector network analyzer connected to a PC via an HP-IB (IEEE 488, IEC 625) link to derive the electrical s-parameters (i.e., S_{11}) for the laser diode, and we used the HP VEE software to save the data as an industry-standard S1P one-port s-parameter device file. We then used the optimization feature in the linear circuit simulator to derive the best-fit resistance and capacitance values.

Package Model

The standard specifications for the package call for a standard impedance level of 25 ohms. The approach taken by HP is to use a submodule containing the laser diode, a chip resistance used for wide-bandwidth impedance matching, and some other optical components. This submodule is mounted on a thermoelectric cooler inside the overall industry-standard butterfly package, which also contains another chip resistor, a chip capacitor, and a chip inductor. To model the electrical response of the packages and the components within the packages accurately, we made extensive use of a vector network analyzer to derive the best equivalent circuit topologies and the best values for the elements within the equivalent circuit for the chip resistors, chip capacitor and chip inductor. The model for the chip inductance was by far the most complex, but a multisection, finite-Q (quality factor) circuit topology enabled us to predict all the measured resonances.

Additionally, we made a careful analysis of the bond wire lengths. In this analysis, we use Bezier curves to simulate the side views of the bond wires, from which we can quite accurately predict bond wire length. Remember that even at 2 GHz, a 1-mm bond wire will have an inductive reactance of about 12.6 ohms, which is significant in a 25-ohm characteristic impedance system. Therefore, attention to detail was most important, and bond wire engineering played an important role in this design. An example of this work, taken from a different example, is as shown in **Figure 5**. In this work, we have looked at the use of ball and wedge bonding to get the best performance, while keeping the costs associated with high-volume manufacturing under control.

In our development work, we came to appreciate the value of the interaction between the simulation work and the precision measurements that were made on the devices

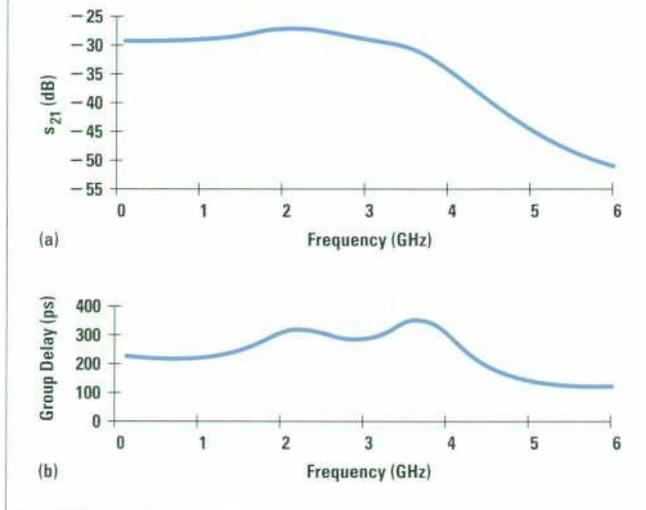


and circuits. Several special test fixtures were constructed for use in the measurement of the characteristics of the laser diode. These precision fixtures were intended for use up to 6 GHz and have standard SMA connectors. We carefully characterized the fixtures so that any component or circuit could simply be placed within the test fixture environment as a subcircuit, simplifying the extraction of the parameters for the elements within the subcircuit. We used standard optimization techniques to get the equivalent circuit response to match the measured response. It is necessary to minimize the number of optimization variables because the difficulty of finding the global minimum using an optimizer increases rapidly as the number of variables increases.

The circuit topology for the overall package is too detailed for us to describe at any length in this paper. We export the predicted s-parameter response for the circuit to a standard S2P s-parameter file. In a 50-ohm characteristic impedance system, we predict the s_{21} response as shown in **Figure 6a**, together with the group delay shown in **Figure 6b**. The input reference plane (port 1) is at the end of a microstrip line on a printed circuit board substrate 8.0 mm from the butterfly package itself. The bias and the signal are brought in through separate terminals

Figure 6

(a) Overall package model s_{21} response. (b) Package model group delay response.

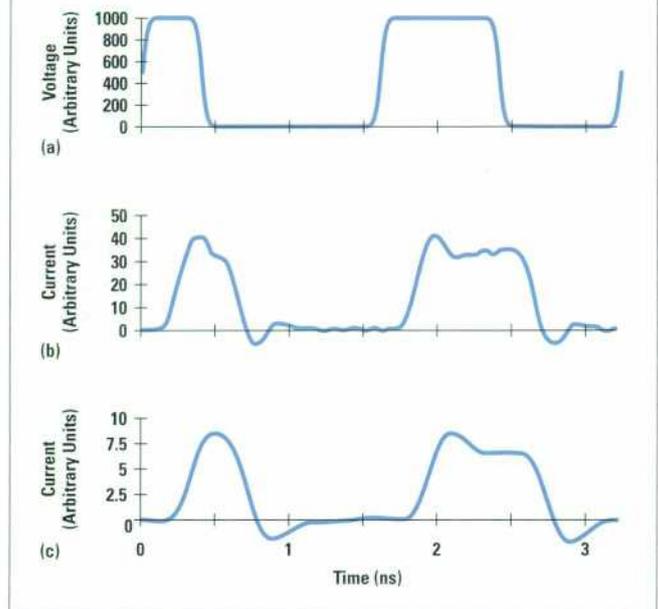


on the butterfly package. The group delay plot is important since we need to maintain the integrity of the rectangular pulses coming in, pulses with Fourier components that we must be concerned with up to at least 2 GHz. The response shown in **Figure 6a** is what we would expect to measure using a vector network analyzer in a 50-ohm environment. In practice, a laser diode driver chip is used to supply the bias current and drive the pulses into the laser diode.

Having defined the package response in an s-parameter file, we can apply the FFT technique discussed earlier to the s_{21} voltage transfer ratio to predict the voltage across the laser diode dynamic resistance, and consequently, the current through the resistive part of the laser diode. This we have done for the case in which the packaged laser diode is driven directly from a 50-ohm characteristic impedance PRBS test set, and for the case in which the packaged laser diode is driven from the laser diode driver example discussed before. In **Figure 7**, the 2.488-Gbit/s voltage waveform with a 10%-to-90% rise and fall time of 75 ps from the signal generator is shown on top, and below it the predicted current waveform into the laser diode when the packaged laser diode is driven directly from the test set. The bottom waveform represents the case in which the laser diode driver is used. Note that the laser diode driver slows down the rising and falling edges of the waveform.

Figure 7

(a) 2.488-Gbit/s voltage waveform with a 10%-to-90% rise and fall time of 75 ps from the signal generator. (b) Predicted current waveform into the laser diode when the packaged laser diode is driven directly from the test set. (c) Predicted current waveform into the laser diode for the case in which the laser diode driver is used.



Using a fast sampling oscilloscope to monitor the current through the laser diode, we then predict the waveform display for the oscilloscope triggered at twice the bit interval to look as shown in **Figure 8**. The plot resembles the eye diagram frequently used in communications engineering. The upper plot is driven directly from the fast PRBS test set, while the lower plot is with the laser diode driver present. In both cases, we notice the effects of overshoot.

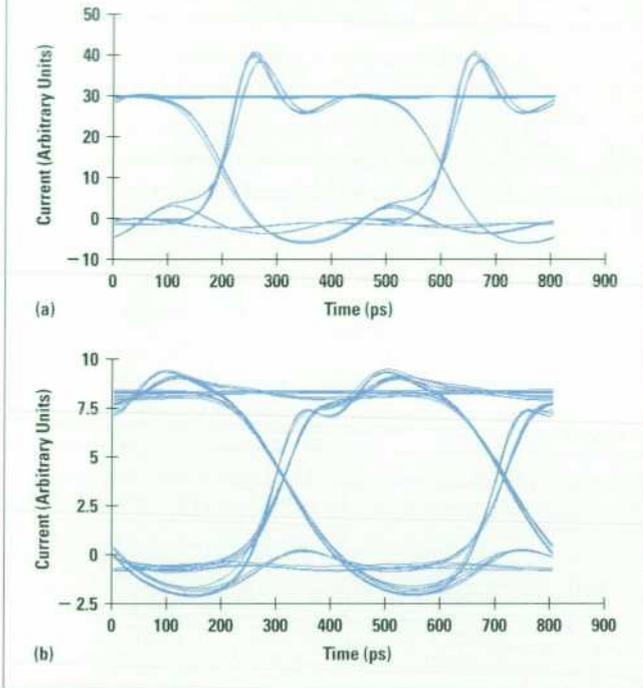
Before the predicted current waveform into the laser diode is used in the laser diode simulator to predict the amplitude and the wavelength of the light emitted from the laser diode, we rescale the amplitude of the predicted current and add the laser diode bias offset.

Laser Model

A laser model for multiple quantum well distributed feedback lasers has been developed specifically for dynamic laser modeling. Emphasis has been placed on developing fast simulations that can readily be used for design and optimization, and hence lumped spatial models have been

Figure 8

Waveforms of Figure 7 displayed on an oscilloscope triggered at twice the bit interval. The plots resemble the eye diagram frequently used in communications engineering. (a) Waveform of Figure 7b. (b) Waveform of Figure 7c.



Of key importance in the carrier dynamics is the time taken for electrons to cross the separate confinement layers and to be captured by the wells. This limits the dynamic performance of high-speed lasers. While the capture process itself may occur on subpicosecond time scales, the diffusion across the confinement layers can take tens of picoseconds. This is shown schematically in **Figure 9**. The diffusion time constant depends on waveguide dimensions and doping levels. Thermionic emission of carriers back out of the wells into the barrier and waveguide layers also occurs, again on time scales of tens of picoseconds. This emission time constant can be controlled by varying the well and barrier dimensions. It is evident that a short diffusion time constant and a long emission time constant lead to efficient carrier injection into the wells and enhanced bandwidth. Alternatively, careful design of the quantum wells suppresses unwanted high-frequency oscillations responsible for excess wavelength chirp and patterning under digital modulation. The electron rate equation for the barrier and confinement layers has therefore been introduced to account for carrier injection initially into the barrier levels, and to describe the coupling between carriers in the barriers and the wells:

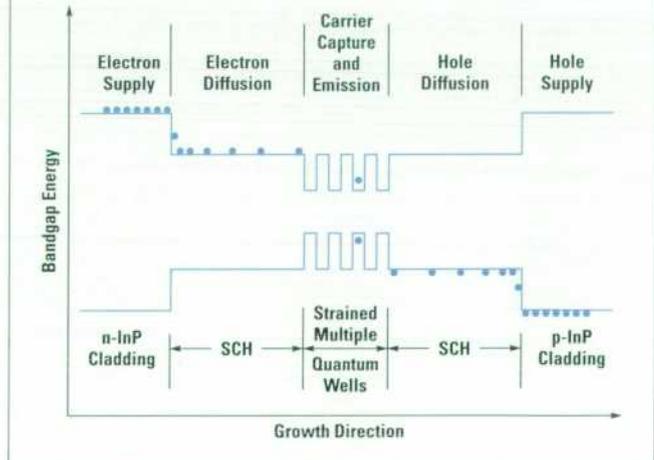
$$\frac{dN_b}{dt} = \frac{J}{qd_b} + \frac{N_w V_w}{\tau_e V_b} - \frac{N_b}{\tau_c}$$

used rather than traveling-wave models. By implementing empirical descriptions obtained from rigorous steady-state models, longitudinal-mode spatial hole burning in distributed feedback lasers could be included. Such descriptions also make it possible to take into account the effect of the coupling strength of the internal grating. To model the high levels of hole burning observed both in phase-shifted gratings and in antireflection and high-reflection coated lasers, the laser is subdivided into three sections along the cavity to account for the resulting inhomogeneity. The three sections are coupled to each other and the photon population through empirical relations which are derived from traveling wave models.

The model also includes nonlinear effects resulting from carrier transport in quantum well lasers through the implementation of a barrier-level carrier rate equation. The barrier carrier population is coupled to the well carrier population through capture and escape processes. These processes describe the carrier bottleneck, which can significantly reduce the laser bandwidth.

Figure 9

Diffusion across the confinement layers in a multiple quantum well distributed feedback laser. (SCH is the waveguide, or separate confinement heterostructure.)



Here N_b and N_w represent the carrier densities in the barriers and wells, respectively. V_b and V_w are the volumes in the barriers and wells. d_b accounts for the depth of the separate confinement layers. The current density J is injected into the separate confinement layers. q is the electron charge. The carriers are depleted with a capture lifetime τ_c . This lifetime is dominated by the diffusion across the layers and is of the order of tens of picoseconds. Carriers are also reinjected from the wells into the barrier and confinement layers with an associated time constant τ_e , which is determined by the thermionic emission. This can be of the order of a hundred picoseconds, and has significant temperature dependence.

The injection of carriers into the quantum well is therefore determined by the ratio of the diffusion and capture rate time constant to the thermionic emission rate. This creates a carrier bottleneck of considerable importance for high-bandwidth lasers. The rate equation for carriers in the wells is therefore written in terms of the capture and emission rates:

$$\frac{dN_w}{dt} = \frac{N_b V_b}{\tau_c V_w} - \frac{N_w}{\tau_e} - v_g GP - \frac{N_w}{\tau_s}$$

The above equation also takes into account carrier depletion through stimulated emission, which is described by the term $v_g GP$, where v_g is the group velocity of light in the laser cavity, G is the optical gain, and P is the photon density. Both radiative and nonradiative recombination mechanisms are described through the carrier lifetime τ_s . The gain term includes a linear dependence on the carrier density and a nonlinear gain compression term ϵ :

$$G = \frac{(dg/dN)(N_w - N_t)}{(1 + \epsilon P)}$$

where dg/dN is the differential gain and N_t is the transparency carrier density. In the simulation, the parameters used in the modeling have been derived from steady-state measurements wherever feasible. The gain has been modeled by assuming a linear differential gain dg/dN of $6 \times 10^{-16} \text{ cm}^2$, a transparency carrier density N_t of $1.5 \times 10^{18} \text{ cm}^{-3}$, and a gain suppression factor ϵ of $3 \times 10^{17} \text{ cm}^3$. The gain suppression factor accounts for nonlinearities such as spectral hole burning, carrier heating, and transverse-mode spatial hole burning.

The rate equation for the photons is coupled only to the carrier population in the wells:

$$\frac{dP}{dt} = \Gamma \beta B N_w^2 + \Gamma v_g GP - \frac{P}{\tau_p}$$

The term $v_g GP$, which depletes carriers in the wells (N_w), feeds photons into the guided laser beam. The spontaneous emission is described by the term $\Gamma \beta B N_w^2$, where Γ describes the overlap of the optical mode with the quantum wells, β is the spontaneous emission coupling factor, and B is the bimolecular recombination coefficient. τ_p is the photon lifetime, which is a function of the grating coupling strength.

The cavity refractive index is perturbed by fluctuations in the carrier density to chirp the operating wavelength λ . Because wavelength chirp limits device performance as described earlier, the linewidth enhancement factor α_H is introduced to describe chirp for a given laser structure. The linewidth enhancement factor is proportional to the quotient of the differential refractive index dependence on carrier density $d\mu/dN$ and the differential gain dg/dN .

$$\alpha_H = \frac{4\pi d\mu/dN}{\lambda dg/dN}$$

The lower the value of α_H the lower the wavelength chirp. The linewidth enhancement factor for a strained quantum well InGaAsP laser is approximately 2.5.

The transport effects lead to a significant carrier population in the barriers and the confinement layers. Because a significantly larger proportion of the optical mode overlaps with the confinement layers when compared with the wells, this population also perturbs the effective refractive index of the cavity and therefore leads to enhanced wavelength chirp. The overall effective refractive index perturbation can therefore be summarized in terms of contributions from the barriers and the wells:

$$\mu_{\text{eff}} = \Gamma_w \frac{d\mu}{dN_w} N_w + \Gamma_b \frac{d\mu}{dN_b} N_b$$

The rate equations are solved using Runge-Kutta algorithms.

The laser model is written using C++ and runs on a PC and is fast enough to be suitable for interactive use. The outputs are described in terms of the time-resolved power and wavelength.

Optical Fiber Model

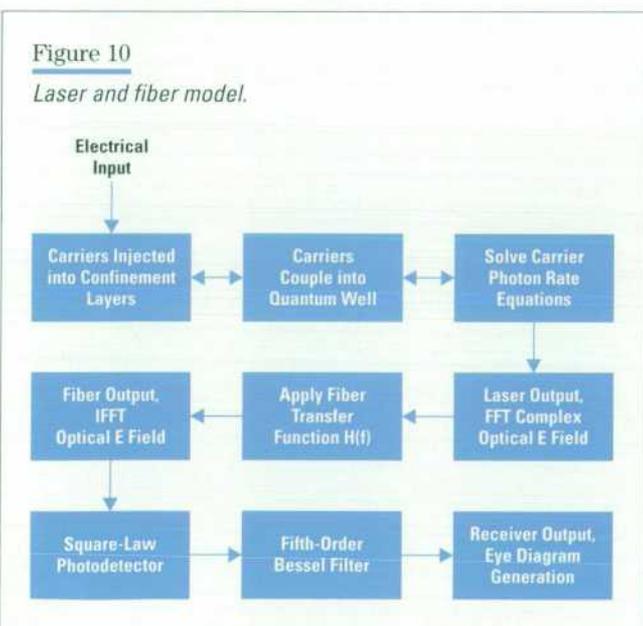
This model takes the complex optical field generated by the laser model to determine the signal propagating along the fiber. The optical field experiences a frequency dependent retardation as it propagates along the fiber. The model accounts for dispersion using the transfer function $H(f)$:

$$H(f) = \exp\left(-j \frac{\pi D \lambda^2 L}{c} f^2\right).$$

In the simulation, D is the fiber chromatic dispersion, λ is the center wavelength, L is the fiber length, f is the base-band frequency, and c is the speed of light. The dispersion D is assumed to be $17 \text{ ps} \cdot \text{nm}^{-1} \text{km}^{-1}$. The model assumes that first-order dispersion dominates and neglects nonlinearities, a justifiable assumption for low-power radiation ($< 10 \text{ dBm}$) at 1550 nm in standard-dispersion fiber. The output optical pulse after any desired length of fiber can be predicted.

The fiber output power is incident on a square-law detector, the output of which is subsequently filtered in the frequency domain by a fifth-order Bessel filter with a bandwidth of 0.7 times the bit rate. Superposition of bit periods gives eye diagrams, which can be compared directly with experiment.

A schematic drawing of the elements of the laser and fiber model is shown in **Figure 10**.



Comparison of s-Parameter Results

This section presents a comparison of the modeled and measured s_{11} and s_{21} parameters for an HP LSC2500 2.488-Gbit/s DFB laser diode developed at ICO. The s_{11} measurement determines the reflected power from the 25Ω -terminated laser diode into the 50Ω HP 85047A s-parameter test set. The s_{21} measurement determines the AM response of the laser diode. This is measured at high bias currents (where the intrinsic relaxation resonance of the chip is out of the frequency band of interest) to compare the measured result with the purely electrical microwave model.

Using WaveMaker software, a complete circuit description of the modulation signal path, from the RF pin on the butterfly package through the laser diode and back out to the two ground pins of the outer package, was constructed in a netlist format. Electrically, the laser diode was modeled as a 7Ω resistor (determined by experimental measurement) in parallel with a 4-pF capacitor (estimated from the structural layout).

s_{11} Measurements. The netlist was written to display the modeled and measured s_{11} results on the same plot. The circuit element values can then be adjusted, maintaining realistic values, to best match the measured data over the frequency range of interest.

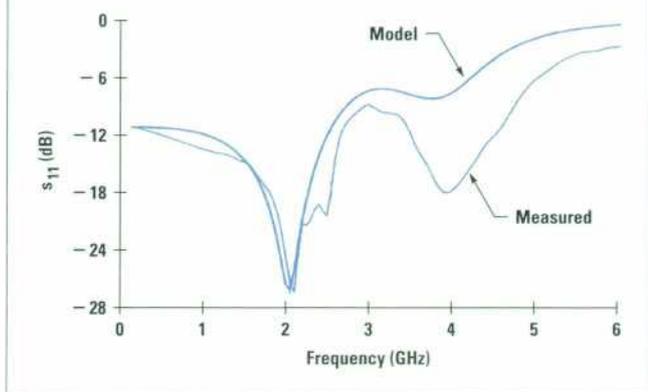
The model was used to determine the effect of changing component values on the s-parameter response. This strategy was used to determine component values that will optimize the product performance and to set tolerances on component values and bond wire lengths.

s_{21} Measurements. The s_{21} measurement is an electro-optical measurement that cannot be truly accounted for in WaveMaker because of the relaxation oscillation of the laser chip. However, by using a high drive current for the laser diode to push any intrinsic resonance effects of the laser diode well out of the 6-GHz bandwidth of interest, it is possible to compare the modeled and measured data. This purely electrical model then predicts the measured s_{21} parameter except in magnitude, since no account can be taken of the laser diode's electrooptic conversion efficiency.

The simulation results in **Figure 11** and **Figure 12** show the best achieved s_{11} and s_{21} fits, respectively, for a 2.488-Gbit/s DFB laser that has an impedance matching

Figure 11

A plot comparing the measured and modeled s_{11} parameters.



resistor and a 3.9-pF capacitor in the package. The capacitor is placed in parallel with the laser diode, thereby establishing a resonant circuit between the capacitor, the resistance of the laser chip, and the inductance of the gold wire bonds. The response of the resonant circuit smooths out the rising edge of the electrical pulse, preventing overshoot and hence minimizing wavelength chirp.

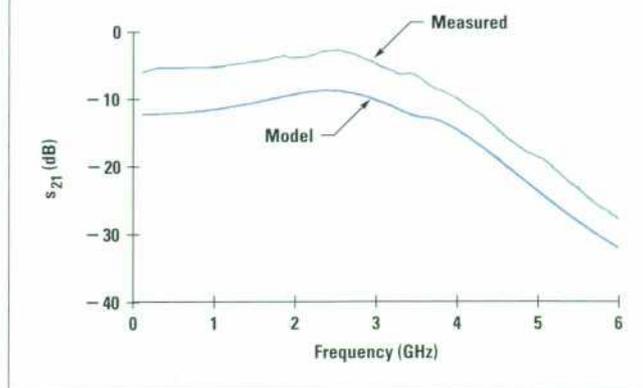
The bond wires in the package were modeled as pure inductances using the well-established value of 1 nH of inductance per millimeter of bond wire for 17- μ m-diameter bond wire. All of the connections on the RF path were made using two bond wires in parallel, which effectively halves the inductance value. For example, 0.7 nH represents two bond wires 1.4 mm long. It was found experimentally that additional bond wires do not reduce the inductance significantly more because of the inevitable closer proximity of these wires on the bond pads.

Comparison of Measured and Modeled s_{11} . Figure 11 shows a plot of the measured and modeled s_{11} parameters overlaid. The low-frequency (< 3 GHz) fit of the model is quite respectable. It predicts the resonance resulting from the transition between the 50-ohm transmission line of the test set and the 25-ohm microstrip lines of the package to within 0.1 GHz and it also predicts the magnitude of the measured data generally to within better than 3 dB.

The resonance at 3.9 GHz caused by the interaction between the bond wire inductance, the capacitor, and the chip resistance is also predicted to within 0.2 GHz. However the depth of this resonance, or the Q factor of the

Figure 12

A plot comparing the measured and modeled s_{21} parameters.



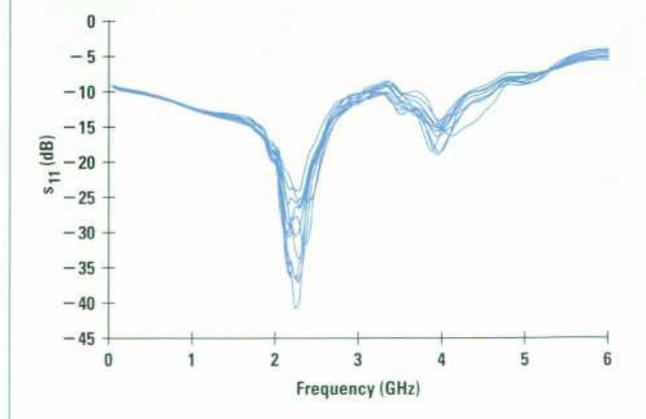
oscillation, is not so well-simulated. The reason for this is not fully understood.

A feature of the measured results in Figure 11 is sharp spikes in the s_{11} response at around 2.4 GHz. These have been found to be a resonance that occurs because the submodule base is not a true ground.

Figure 13 shows s_{11} measurements on ten different modules, showing the reproducibility. Variations in the 3.9-GHz resonance are attributed to tolerance variations in the capacitor used and bond wire length process variations. All of these measured results meet the s_{11} requirement for a 2.488-Gbit/s laser diode.

Figure 13

A plot of s_{11} measurements on ten different modules.



Comparison of Measured and Modeled s_{21} . Figure 12 is a plot displaying the measured and modeled s_{21} parameters of a device. The low-pass filtering action seen in the measured s_{21} response is also modeled in WaveMaker. This filtering prevents high-frequency electrical oscillations from reaching the chip, which would cause it to produce unwanted optical oscillations. Comparing the shapes of these curves, the modeled results predict 3.4 dB of ripple with a resonant peak at 2.4 GHz while the measured results show 2.6 dB of ripple with a resonance at 2.5 GHz. The roll-off of the modeled data is 9 dB/GHz, in good agreement with the measured data.

The 3-dB bandwidth requirement for a 2.488-Gbit/s laser is a 3.5-GHz minimum. The internal filter has been designed to give enough bandwidth and to limit the speed of the rising edge. Figure 14 displays s_{21} measurements on ten modules, showing that typical modules have a 4-GHz 3-dB bandwidth.

Microwave Model Output. The modeled s_{21} parameter file is convolved with an ideal, user-defined pulse train in the frequency domain and then inverse Fast Fourier transformed back to give a realistic electrical pulse train in the time domain, which is then offset to set an extinction ratio. This offset electrical pulse can then be used as the input to a model developed by Bristol University, which takes an electrical current pulse input and, using parameters to describe the laser diode in detail, models the optical output from the device. This model also includes a

description of fiber parameters and can predict the optical pulses after lengths of fiber.

Comparison of Optical Pulse Shapes

The modeling software developed by the University of Bristol was designed to calculate the optical output of quantum well DFB laser diodes grown using the InGaAsP-InP material system. The model can be run using a time-domain description of an electrical pulse train as the input. This has been created at ICO using WaveMaker software, as discussed in the previous section. The Bristol software then models the laser's response to this large-signal modulation. The reason for installing these combined models at ICO was not only so that we could optimize the s-parameter response of the 2.5-Gbit/s laser but also so that we could predict the effects on the system performance of changing the package components and bond wire lengths by looking at the pulses at the receiver after the fiber. The unfiltered optical output can be displayed as pulses in the time domain before and after an arbitrary length of fiber, the dispersion properties of which can be specified.

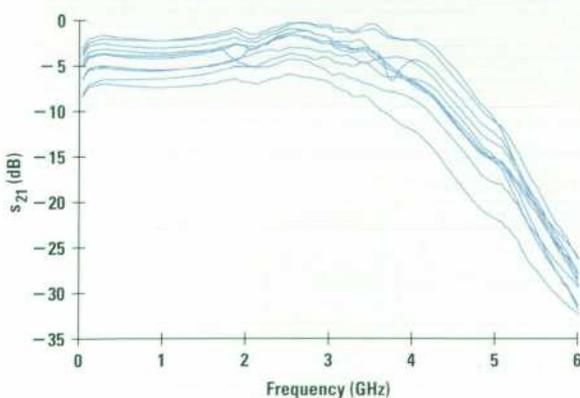
The software also displays the wavelength deviation of the laser output during the pulse train, calculated from the known laser parameters. Depending on the timing of the wavelength deviation or chirp during a bit period, quite different effects can be seen when the signal is transmitted over fiber. At 1550 nm, short wavelengths travel faster through the fiber than long wavelengths. An optical pulse that starts with a short wavelength and ends with a long wavelength will spread apart over the fiber, while a long wavelength at the start and a short one at the end will cause the pulse to be compressed over the fiber.

This model has been used in conjunction with WaveMaker to model a system using an HP 2.5-Gbit/s directly modulated DFB laser in a 14-pin butterfly package as the transmitter, and over 106 km of standard single-mode fiber with 17 ps/nm/km of dispersion.

To test the accuracy of this combined model, the modeled results were compared with experimental results acquired by modulating the laser with an HP 70841B pattern generator. The output after the fiber was then fed into an HP 11982A lightwave converter with a 15-GHz bandwidth and the resulting pulses were displayed on an HP 54120B digital sampling oscilloscope. For this paper, a 0101 0011 pulse train was used to compare modeled with experimental results.

Figure 14

A plot of s_{21} measurements on ten modules.



Back-to-Back (No Fiber) Results. Figures 15a and 15b show the comparison between modeled and experimental results measured back to back for a 2.5-Gbit/s laser. The pulse shape of single and double ones can be seen to be in good agreement with the measured data. It is important that the pulse widths for single and double ones are in both cases 400 ps and 800 ps, respectively.

The other features of the pulse train predicted by the model are the oscillations in both the double zeros and the double ones. The first oscillation in the double zero, caused by an electrical resonance between a 3.9-pF parallel plate capacitor, the bond wire inductance, and the laser resistance and capacitance, is predicted with correct magnitude. A second oscillation appears which is damped with respect to the first. The frequency of oscillation in the ones (again because of the electrical oscillations) is 2.5 GHz, in very good agreement.

Results over 106 km. Figures 16a and 16b display modeled and experimental data after transmitting the pulse through 106 km of fiber, equivalent to a dispersion of 1800 ps/nm at 1550 nm. As can be seen in both cases the pulses are getting narrower and the widths are still in reasonable agreement. It is the narrowing of the pulse width that has been found to cause an improvement in bit error rate over fiber. This can be explained by realizing that as the pulse compresses, the power at the decision point is raised, thus improving the signal-to-noise ratio at the receiver.

The features of the modeled pulses include the sharpening and subsequent buildup of optical power in the front of the pulses and the progression into the zeros of the small optical output that was induced by the resonance.

In summary, for both back-to-back transmission and transmission over 106 km of fiber, the model has predicted both

Figure 15

- (a) Modeled 2.5-Gbit/s pulse train back to back.
- (b) Measured 2.5-Gbit/s pulse train back to back.

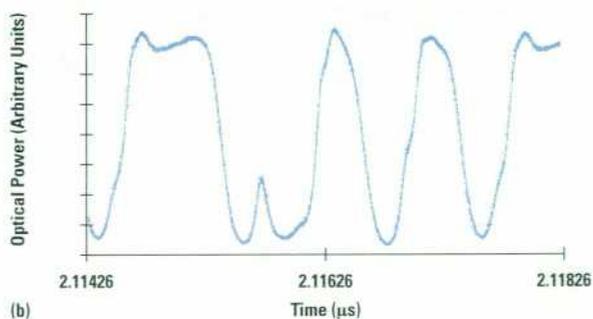
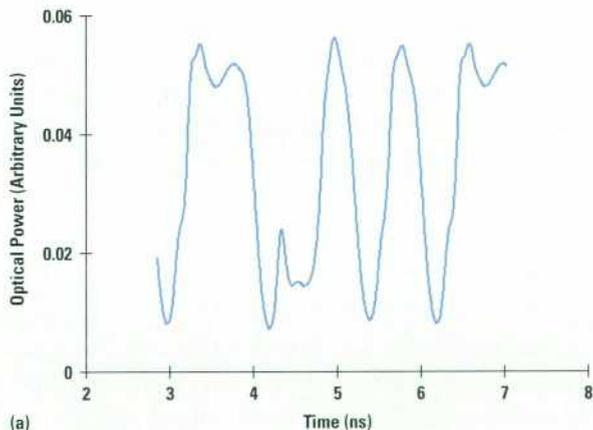
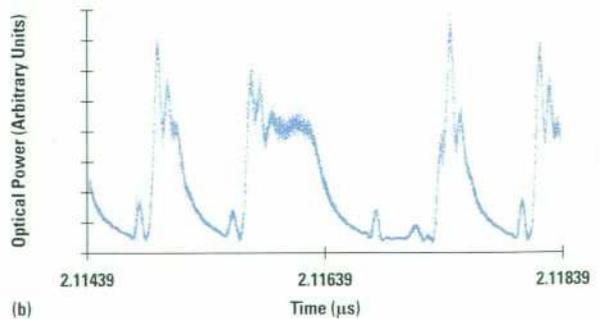
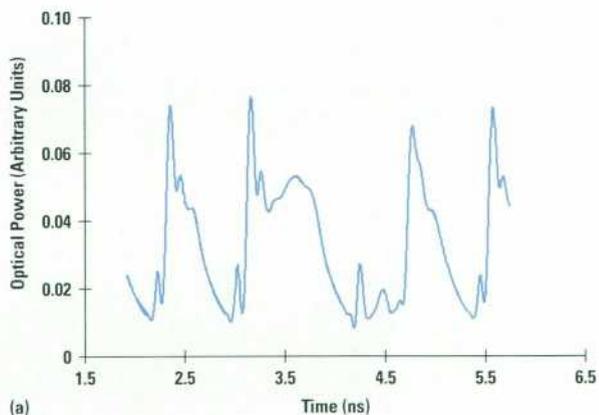


Figure 16

- (a) Modeled 2.5-Gbit/s pulse train after 106 km of fiber.
- (b) Measured 2.5-Gbit/s pulse train after 106 km of fiber.



pulse widths and pulse shapes with a reasonable degree of accuracy. The model was also successful in predicting specific features of the pulse train, such as oscillations and the sharpening of the rising edge.

Dispersion Penalty Measurements

The combined models discussed in this paper were used to produce good s-parameter results and good pulse shapes over fiber. The experimental RF test to confirm this data transmission was the bit error rate test, the experimental arrangement of which is shown in **Figure 17**. A pseudorandom $2^{23} - 1$ bit stream is produced by the HP 70841B pattern generator at 2.488 Gbits/s. This data modulates the laser diode, which has a dc bias provided by a separate current source. The laser's optical output is then directed straight into an optical attenuator for back-to-back measurements or into the fiber, an amplifier (if required), and an attenuator for measurements over fiber. The purpose of the attenuator is to vary the power level reaching the receiver. The output of the attenuator is then fed into an HP 83446A STM 16 lightwave clock and data receiver, the output of which is fed into an HP 70842B error detector, which compares the transmitted and received optical pulses, measures the total number of erroneous bits received during a time period, and calculates the bit error rate.

Bit error rate curves are plotted in **Figure 18**. These are plots of bit error rate versus received power in dBm for

Figure 17

Schematic drawing of the bit error rate test setup.

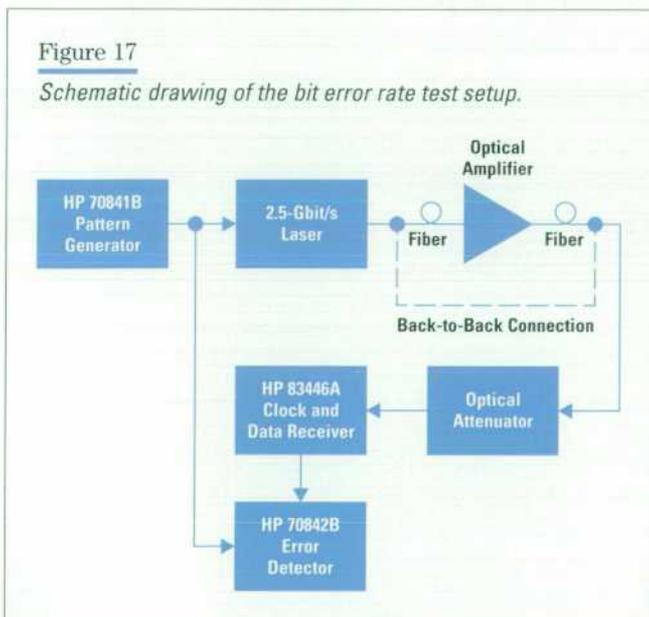
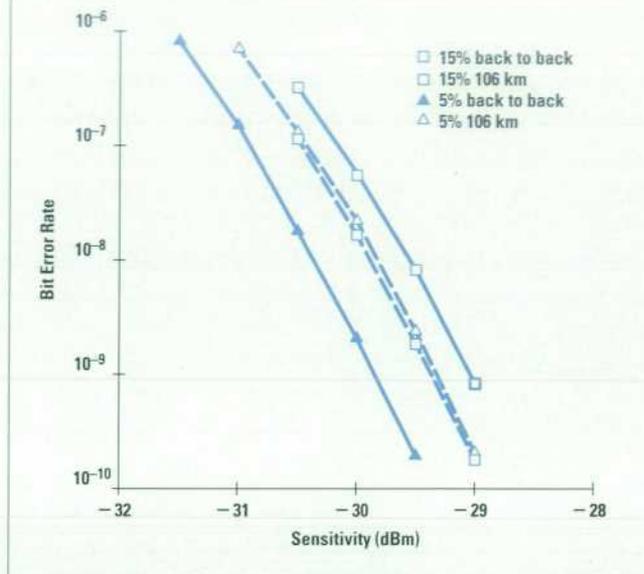


Figure 18

Bit error rate curves at two different extinction ratios.



varying extinction ratios, measured back to back and over fiber. As the power at the receiver increases, the bit error rate drops. The dispersion penalty indicates how much the power at the receiver must change to maintain the same bit error rate for measurements back to back and over fiber. The dispersion penalty can be seen as the horizontal displacement between the two bit error rate curves and is measured in dB. The value can be positive or negative as explained later.

Figure 18 shows how these dispersion penalties depend on another parameter—the extinction ratio. The extinction ratio is a measure of how close to its threshold current or off state the laser is modulated. It is in the region closest to the threshold current that the largest changes in the device's carrier density, refractive index, and hence emission wavelength occur. To transmit data over fiber we intuitively want the wavelength to vary as little as possible between the on and off states. The extinction ratio is defined as the power in the off state divided by the power in the on state. The greater the extinction ratio the less the wavelength shift, but for maximum sensitivity it is desirable to have as little power as possible in the off state. A compromise value must be used. The measurements in **Figure 18** were taken at extinction ratios of 5%

and 15%. The respective dispersion penalty measurements over 106 km of fiber are 0.5 dB and -0.33 dB.

Figure 19 shows bit error rate results for a device measured back to back, at 530 ps/nm (31 km), 1800 ps/nm (106 km), 2400 ps/nm (141 km), and 3600 ps/nm (211 km). Here it can clearly be seen that the dispersion penalty can improve over fiber. The improvement in the dispersion penalty is attributed to the pulse compression that takes place over the fiber when pulses consist of long wavelengths at the beginning and short wavelengths at the end. Over 211 km, a double pulse compresses from 800 ps to

624 ps, which must be close to the limit of useful compression if the sampling is halfway through each bit.

In this example, the error rate is seen to get worse over shorter distances of fiber before improving over longer distances. This can be explained by power leaving the center of the pulse before being replaced by the faster traveling tail. This confirms that it is not only the magnitude of the chirp but also the timing of it that is important.

Conclusion

We have described the microwave, laser, and fiber models that were used in the development of the HP LSC2500 2.488-Gbit/s DFB laser diode module. These models have been shown to accurately predict the microwave s-parameters of the laser module and therefore the electrical signal reaching the laser diode. Knowing the electrical signal to the laser, the optical output has been predicted, which leads to the wavelength chirp and therefore the dispersion penalty over long lengths of single-mode fiber at 1550 nm. The predicted optical pulse shapes both before and after the fiber agree closely with the experimental results.

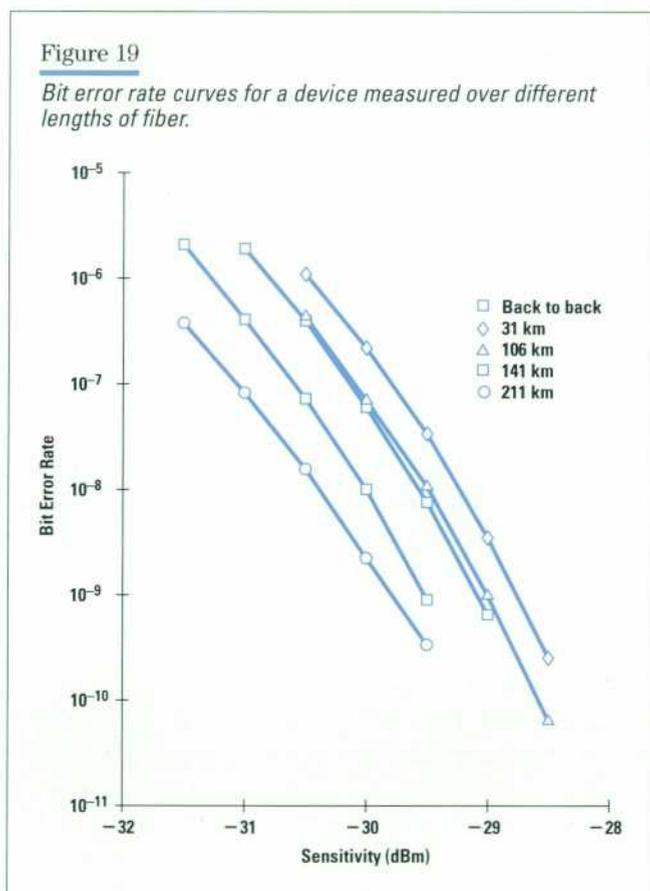
With the aid of the models, a low-chirp multiple quantum well DFB laser module has been developed. Knowledge of the modeled behavior of the laser diode as a function of the input electrical pulse shape has led to deliberately shaping the input pulse to give the minimum wavelength excursion during direct modulation, and therefore a high yield of low-dispersion-penalty laser diodes. We have also demonstrated that these devices can be successfully used for transmission distances in excess of 200 km.

Acknowledgments

The authors wish to acknowledge Dave Smith and Nigel Jowitt of HP ICO for supplying test results from the automated test equipment and John Lynn, Adrian Kendall, and Maheschandra Mistry for help with prototype work.

Figure 19

Bit error rate curves for a device measured over different lengths of fiber.





Stephen M. Gee

Steve Gee is a senior engineer at HP's Ipswich, England Components

Operation. He is currently part of the team working on a 2.5-Gbit/s electroabsorption modulator and a 2.5-Gbit/s transceiver. He joined ICO in 1995 after receiving his BSc degree in physics and physical electronics from the University of Bath. His nonwork interests include playing football and cricket and cycling.



Herbert Lage

Herbert Lage is a device engineer at HP's Ipswich Components Operation

(ICO), where he is working on the development of distributed feedback laser chips. A native of Dortmund, Germany, he received a doctorate in physics in 1992 and worked as a senior scientist at Tele Danmark Research before joining ICO in 1994. He enjoys skiing, windsurfing, and tennis in his nonwork time.



Chris Park

Chris Park is a program manager at HP's Ipswich, England Components

Operation. He is currently focusing on a wavelength division multiplexed version of the HP LSC2500 optically isolated distributed feedback laser module and on a 2.5-Gbit/s transceiver. A native of Halifax, Yorkshire, England, he is married and has three children.



Kevin A. Williams

Kevin A. Williams is a Royal Society University Research Fellow at the

University of Bristol (United Kingdom), where he teaches in the Department of Electrical and Electronic Engineering. Kevin received his doctorate in physics in 1995 from the University of Bath.



Richard V. Penty

Richard V. Penty is a Lecturer in electronic engineering at the University

of Bristol (United Kingdom). He received a doctorate from Cambridge University in 1990, specializing in nonlinear optical fibers. A native of Nottingham, England, he is married and the father of two children.



Ian H. White

Ian H. White is Professor of optical communications systems at the

University of Bristol (United Kingdom). He teaches in the Department of Electrical and Electronic Engineering. Professor White received his doctorate in electrical sciences from Cambridge University in 1984.



Joseph A. Barnard

Joseph A. Barnard is managing director of Barnard Microsystems

Limited in London, England. His expertise is in microwave optoelectronic circuit and system design software, and he has published widely on different aspects of this subject. He received a doctorate in engineering physics in 1982 from Cornell University. He is an avid tennis and chess player.

Flip-Chip Photodetector for High-Speed Communications Instrumentation

Tun S. Tan

David M. Braun

Tim L. Bagwell

Christopher P. Kocot

Joseph Straznicky

Susan R. Sloan

A family of 7-GHz-bandwidth optical receivers and a nine-channel optical receiver with a gigabit-per-second data rate per channel have been developed for multigigabit lightwave test systems for long-haul fiber-optic telecommunications links and gigabit optical interconnects for computer systems. A new micro-flip-chip process, featuring liftoff-based small-diameter solder bumps, is incorporated with HP high-speed InP p-i-n photodetectors to minimize parasitic capacitance and inductance and enhance responsivity.

Ultrahigh-performance optical receivers to convert optical signals into electrical impulses require very careful packaging of photodetector and electronic amplifier ICs. Very high-performance photodetector and amplifier ICs can be individually fabricated, but when packaged by conventional methods that employ wire bonds, the system performance falls short of the desired goal. This is a result of parasitic capacitance and inductance introduced by large bonding pads and long bond wires. A 7-GHz-bandwidth optical receiver composed of a 26-GHz InP photodetector and a 12-GHz GaAs amplifier requires package-related parasitic capacitance and inductance of less than 50 fF and 10 pH, respectively. To minimize the parasitic components introduced by packaging, we have developed a hybrid flip-chip integration approach that bonds a photodetector directly onto a GaAs or Si receiver IC by means of very small-diameter solder bumps.

The packaging of high-speed optoelectronic components requires coupling of light into a small photosensitive area of the photodetectors, 10 to 15 μm in diameter. This is usually accomplished by precise alignment of the optical fiber or the transmitter (laser) to the photosensitive area of the detector. The flip-chip die attachment technique allows precise placement of the detector and the transmitter on the prealigned and prefabricated solder bump pads. To further ease the alignment tolerance and reduce the packaging cost, we have

integrated a microlens with the detector to increase its effective light acceptance aperture. The combination of low parasitics and self-alignment of the flip-chip die attachment with an integrated microlens has allowed us to realize a family of 7-GHz optical receivers for lightwave test systems and a gigabit nine-channel optical receiver for optical interconnect applications.

Flip-Chip Photodetector with Microlens

Designing a photodetector that simultaneously satisfies the criteria of high bandwidth, high responsivity, small parasitic capacitance, and high optical coupling efficiency is a challenging task. A p-i-n photodetector with a thin i-layer has a very high bandwidth but fails to satisfy the responsivity and low parasitic capacitance requirements. It is possible to meet the bandwidth, low parasitic capacitance, and responsivity criteria by designing a detector with a small geometry and a moderate i-layer thickness. However, an expensive package that employs precision optics is required to meet the coupling efficiency requirement.

By a combination of the flip-chip packaging concept and device geometry optimization we can meet the required system goal. The flip-chip die attachment technique provides the following advantages to the design and performance of a high-speed photodetector:

1. The p-i-n photodetector in a flip-chip configuration, as shown in **Figure 1**, allows the incident optical signal to traverse twice, or make a double pass through the light-absorbing i-layer. This is made possible by the close proximity of the reflecting front ohmic contact to the i-layer and by allowing most of the reflected signal to be converted again to the electrical signal. This gives the designer an extra degree of freedom to satisfy the conflicting requirements of high bandwidth and high responsivity.

2. In a flip-chip configuration, all the electrical contacts are made through the solder bumps, which are located on the front side of the detector. The backside of the detector, where the incident optical signal enters, is free of any metallized contacts. This allows an integral microlens¹ to be fabricated on the back surface of the photodetector. The microlens increases the effective acceptance aperture, thereby increasing the optical coupling efficiency while enabling the use of small-geometry devices. This significantly reduces parasitic capacitance.

3. A micro-flip-chip die attachment technology can be used. Very small-diameter solder bumps allow precision placement of a photodetector directly on top of a receiver IC while adding negligible parasitic capacitance and inductance.

Micro-Flip-Chip Technology

The conventional hybrid integration and packaging approach, which wire-bonds active components to a common substrate, requires bonding pads at least 0.003 inch by 0.003 inch in size. The parasitic capacitance of a bonding pad alone amounts to 80 fF, and the wire-bond inductance is typically a few hundred nH.

There are several integration approaches, both monolithic and hybrid, that can reduce the extent of parasitic component loading. Optoelectronic integration technologies that monolithically combine the photodetector and electronic transistors, although the most effective, are still in their infancy, and are not the most cost-effective methods today. Flip-chip bonding the photodiode and the receiver IC to a thin-film substrate using conventional solder bump technology eliminates the parasitic capacitance and inductance of the wire bonds but still suffers from interconnect trace loading between the photodetector and the amplifier. The total capacitance due to the interconnect and the bump pads amounts to 100 to 200 fF. Flip-chip bonding a photodiode directly onto a receiver IC eliminates nearly all unwanted parasitics. However, the size of the solder bumps must be kept small to meet the 50-fF goal. Simulations indicate that solder bumps less than 40 μm in diameter are necessary to meet the performance goal.

Figure 1
Flip-chip photodetector cross section.

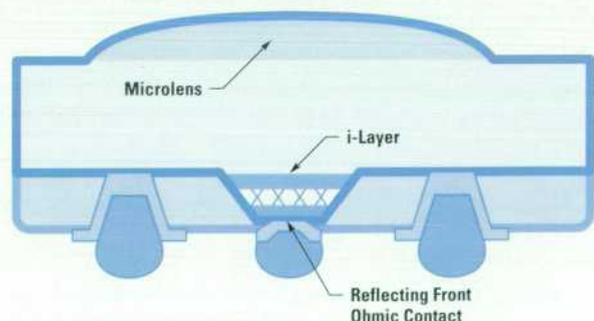
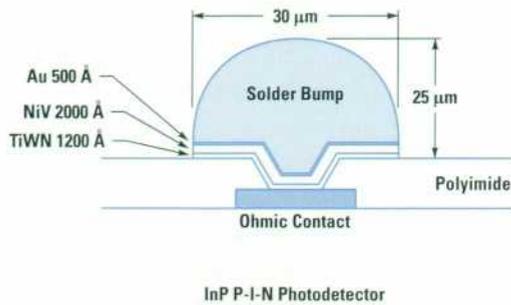


Figure 2

Micro-solder bump cross section.



Based on our system performance goals, we have chosen to develop a micro-flip-chip technology that features solder bumps 30 μm in diameter. During the course of our effort, we have addressed the following technical challenges:

- A process for depositing and defining 30- μm -diameter solder bumps
- A solder bump process that does not induce significant stress or contaminants on the photodetector to cause degradation of its performance
- An underfill process that makes the micro-flip-chip die attachment mechanically rugged
- Precision placement of very small die (350 μm by 350 μm) onto a 1-mm-by-1-mm integrated circuit.

Fabrication Process

The fabrication of a flip-chip photodetector with an integrated microlens involves the following process steps. Undoped InGaAs optical absorption layers and p⁺ doped InP epitaxial layers are grown on the S-doped n⁺ InP substrate using an organometallic vapor phase epitaxial (OMVPE) system. The active region of the detector is defined by mesa etching. The p-ohmic and n-ohmic contacts, which consist of Ti/Pt/Au layers, are fabricated on the front side of the diode. The device is then passivated with polyimide. A microlens is defined on the polished backside by reflowing a circular photoresist pattern at 250°C to form a hemispherical surface, which is subsequently transferred into the InP substrate by means of ion milling. The result is a smooth lens surface with an acceptance aperture of 90 μm and a radius of curvature of 115 μm . A Si₃N₄ film, which serves as an antireflective coating for

the lens, is deposited using PECVD (plasma enhanced chemical vapor deposition). Reactive ion etching opens up the vias in the polyimide to provide electrical contact to the ohmic metal underneath. The base metal stack is deposited by magnetron sputtering following a sputter etch cleaning step. The base metal stack consists of 1200Å of TiWN, 2000Å of NiV, and 500Å of Au.² **Figure 2** shows the cross section of the base metal after deposition and patterning.

The TiWN layer promotes adhesion and acts as a metallurgical barrier to solder. Since the interconnect metal on the InP photodetectors is gold-based, it is imperative to have a metallurgical barrier. Otherwise, upon reflow, the solder will consume the interconnect metal on the photodiode. Cr is another possible choice as the base metal; however, early experiments indicated that Cr etchant attacks InP substrates. NiV provides good solderability, while Au preserves the solderability of Ni when exposed to air. After deposition, the base metal pattern of 30 dots is defined on top of the vias with photoresist patterning and wet etching.

To create 30- μm -diameter, 25- μm -high hemispherical bumps after reflow, we have chosen to define solder pancakes. The tight alignment tolerance and small aperture required make the conventional shadow mask-based evaporation unsuitable. This necessitated the development of a liftoff-based solder deposition and definition process. In the micro-flip-chip process, 40- μm -diameter holes are photolithographically defined in a 14- μm -thick layer of positive photoresist. The exposure and the development time of the resist are controlled so that the sidewalls of the 40- μm -diameter openings are vertical, which is paramount to a successful liftoff. The solder layer is 60wtPb/40wtSn and is deposited by evaporation. This composition of solder was chosen to ensure that the reflow temperature is sufficiently low that it will not induce any thermal stress on the p-i-n photodetector. The diameter of the bumps is set at 40 μm and their height at 14 μm after the liftoff, which gives the optimum aspect ratio of the bumps after the solder is reflowed. **Figure 3** shows an SEM micrograph of the solder pancakes after liftoff.

Before solder reflow, the bumps are coated with no-clean water-based flux. Once fluxed, the bumps are reflowed at 240°C for 45 seconds. The diameter of the reflowed bumps is 30 μm . In contrast to the conventional shadow

Figure 3

SEM micrograph of solder pancakes after liftoff.

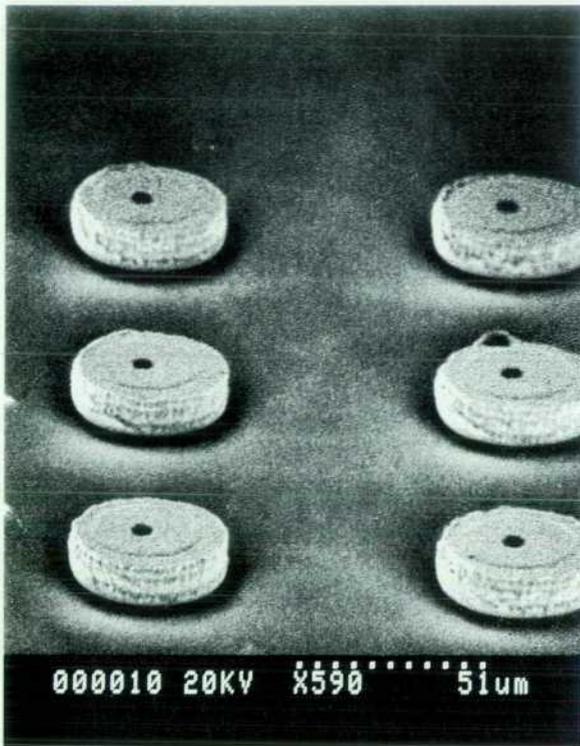
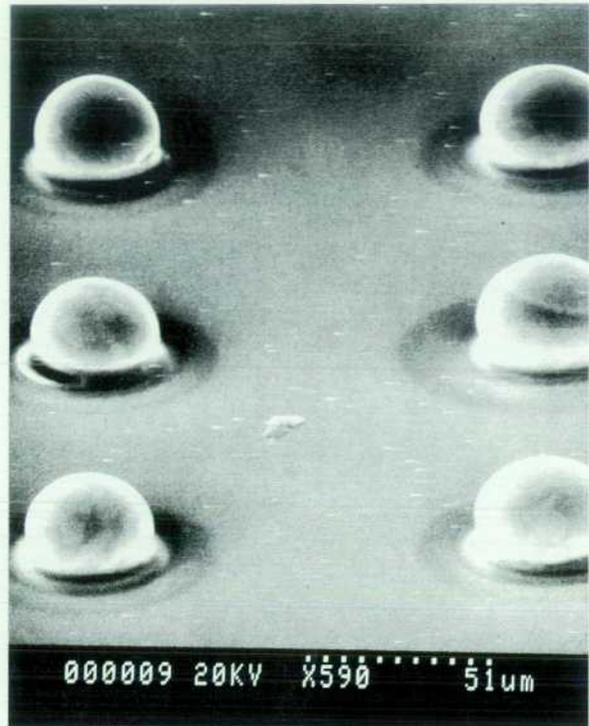


Figure 4

SEM micrograph of microbumps after reflow.



mask technique, our liftoff process allows uniform and precise definition of solder bumps with diameters less than 50 μm . It also eliminates the halo ring normally observed with shadow masks. **Figure 4** is an SEM micrograph of reflowed solder bumps.

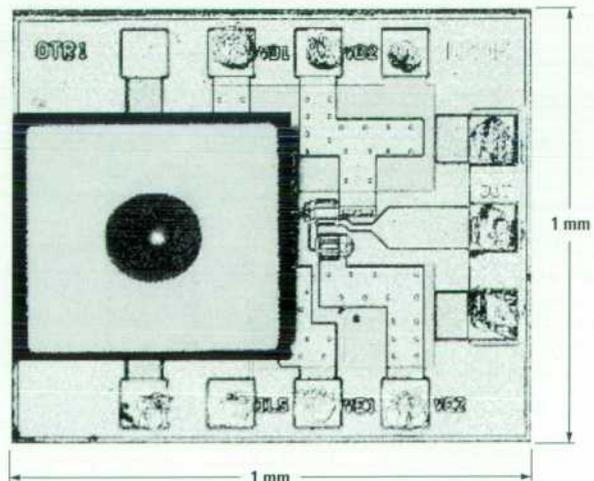
Flip-Chip Assembly

To reduce parasitic capacitance, the dimensions of both the die and the detector active area must be minimized. A typical high-speed InP photodetector has a 10- μm diameter and its die dimensions are 350 by 350 μm . This detector is die-attached to a transimpedance amplifier to realize an optical receiver as shown in **Figure 5**.

The major elements of the die-attach process consist of fluxing the solder bumps, aligning, and reflowing the solder. A precisely controlled amount of flux is applied to the solder bumps of the detector using a microsyringe. It is important to limit the volume of the flux to approximately 1 nl to facilitate the alignment and to minimize the amount of solid residue left on the surface of the detector and the

Figure 5

Optical receiver consisting of a flip-chip photodetector attached to a transimpedance amplifier.



circuit. To increase the mechanical robustness of the attachment, an epoxy underfill composed of 65%wt silica is used to fill the gap between the die and the substrate. The solder bumps on the detector are aligned to the bump pads of the amplifier. Because of the small dimensions of both the chips and the solder bumps, a high-precision flip-chip aligner is used for the alignment. The aligner allows simultaneous viewing of both its lower and upper chucks by means of split optics. Before alignment, the amplifier is loaded on the lower chuck and the detector on the upper chuck. The lower chuck is moved until the image of the bumps and the image of the amplifier contacts are superimposed. A sharp image of both chips, necessary for precise alignment, can be obtained only when the layer of flux coating the solder bumps is very thin. After the alignment both chips are brought into proximity and tacked with a precisely controlled force. Because of the small number and size of the bumps the tacking force should not exceed 10 grams; otherwise, the bumps become deformed and the surface tension is not strong enough to lift the chip during the reflow.

The solder bump interconnection is completed by solder reflow at 240°C in a nitrogen atmosphere using a clamshell type oven. The oven provides very stable and repeatable annealing conditions from run to run. The reflow time is 45 seconds, which includes 15 seconds of ramp time and 30 seconds of dwell time.

Mechanical Characterization

The mechanical strength of the microbumps without the underfill has been tested in shake and shear force tests. Several optical receivers were mounted on a chuck, and shear force applied to the detector was steadily increased until the bumps failed. An average shear force that caused complete die separation was 6 grams, or 2 grams per bump, which corresponds to an equivalent of 25,000g acceleration. The bumps sheared in the middle, which is the optimum situation. The measured shear force agrees well with the predicted value. The shake test subjected our optical receivers to accelerations of 1000g in three axes for a total of 18 hits. All samples have passed this test. When the gap between the die and the substrate is filled with the epoxy underfill material, the attachment becomes so strong that it is impossible to separate the die from the substrate without destroying one or both of them.

Electrical Characterization

One of the key features of our flip-chip detector is its low total capacitance. **Figure 6** shows the measured capacitance for standard and flip-chip photodetectors. FCDxx and PDxx designate the flip-chip and standard devices and xx denotes the diameter of the photosensitive area of the photodiode. FCD15 and PD14 have nearly identical photosensitive areas, while the total active area of the former is one-third that of the latter. As a result, the total capacitance of FCD15 is one-half that of PD14. This is made possible by the micro-solder-bump technology, which allows a reduction in the size of the p-ohmic contact.

Dc and frequency-domain measurements are made by solder bumping individual photodiodes onto specially designed GaAs-based chip carriers whose interconnect patterns maintain a characteristic impedance of 50 ohms. An optical heterodyne system³ illuminates the lensed surface of the packaged device (**Figure 5**), and the converted signal is measured by a 50-GHz network analyzer.

Figure 7 shows a typical frequency response of a 7- μ m-diameter flip-chip detector. Its -3-dB electrical bandwidth is in excess of 50 GHz, which is in good agreement with our simulations. This state-of-the-art frequency performance is the result of careful minimization of the parasitic capacitance and inductance made possible by the micro-flip-chip technology.

Figure 6
Capacitance of p-i-n photodetectors.

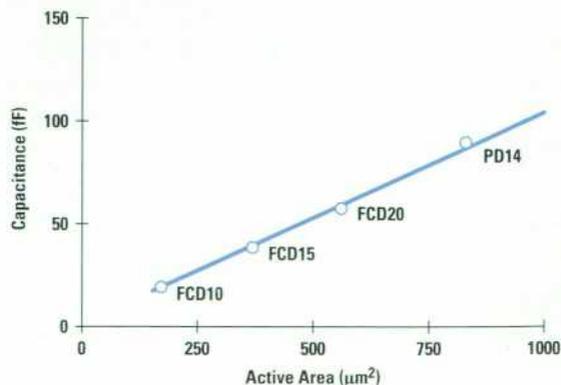
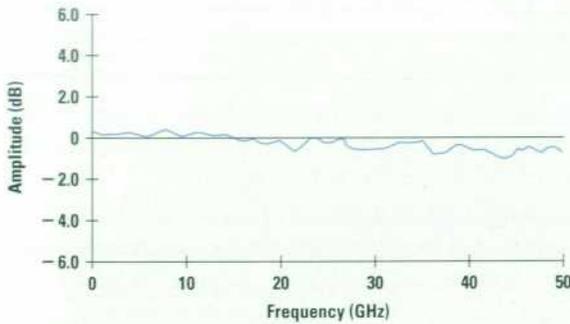


Figure 7

Frequency response of a 7- μm -diameter flip-chip photodetector.



Experiments were done to eliminate concern that the stress induced on the detectors in the flip-chip bonding process might degrade the performance of the photodetectors. The dark current of the InP p-i-n photodetector is one of the parameters that is most sensitive to stress-induced damage. To determine the effect of induced

stress, the dark current of a very large population of flip-chip-mounted photodetectors was monitored. On the average, the increase in dark current is insignificant and under 200 pA. The typical final dark current of a flip-chip-mounted photodetector is under 1 nA.

Typical dc responsivity of a 10- μm -diameter flip-chip detector with a 0.7- μm -thick i-layer or active layer is 0.8A/W. In comparison, the responsivity of standard (non-flip-chip) detectors with the same i-layer thickness is 0.55A/W. The enhanced responsivity of the flip-chip detector is a result of the incident light beam bouncing off the top metal contact and passing twice through the absorbing i-layer.

The InP microlens integrated at the back of the flip-chip detector significantly increases the effective light acceptance area of the detector. **Figure 8** shows the normalized optical responses of lensed and unlensed 25- μm flip-chip photodetectors. The lensed photodetector has an effective collection surface that is four times larger than the unlensed device. The advantage of having a microlens becomes more evident for smaller diameters.

Figure 8

Responses of unlensed (left) and lensed (right) 25- μm -diameter flip-chip photodetectors.

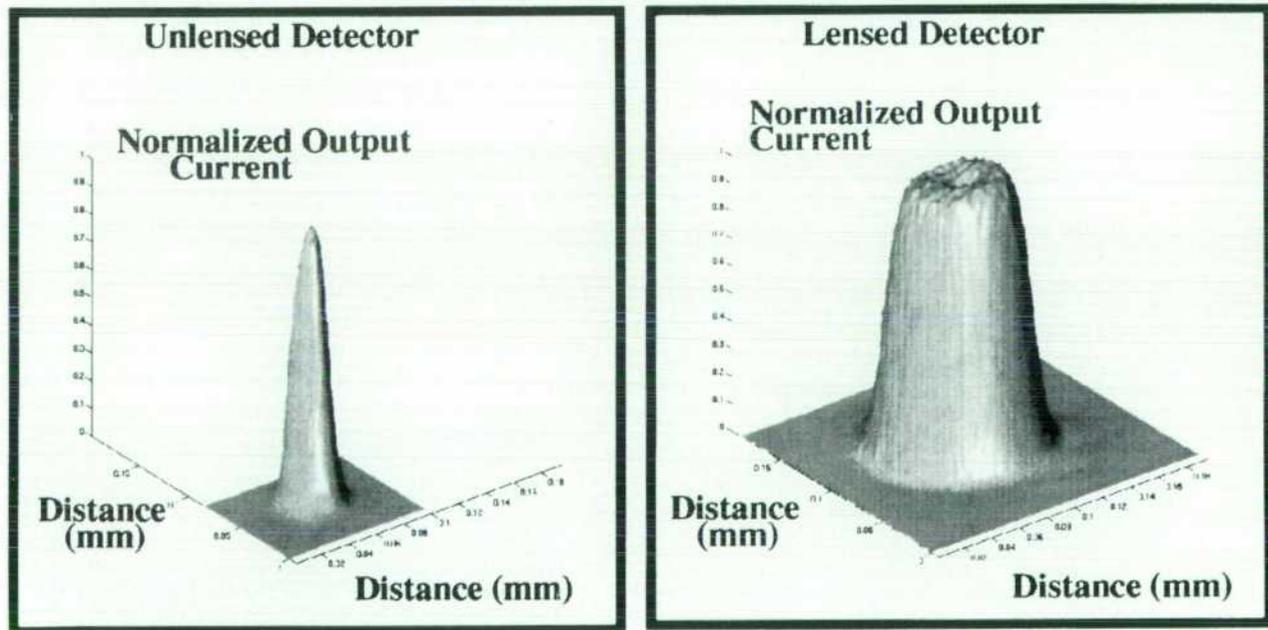
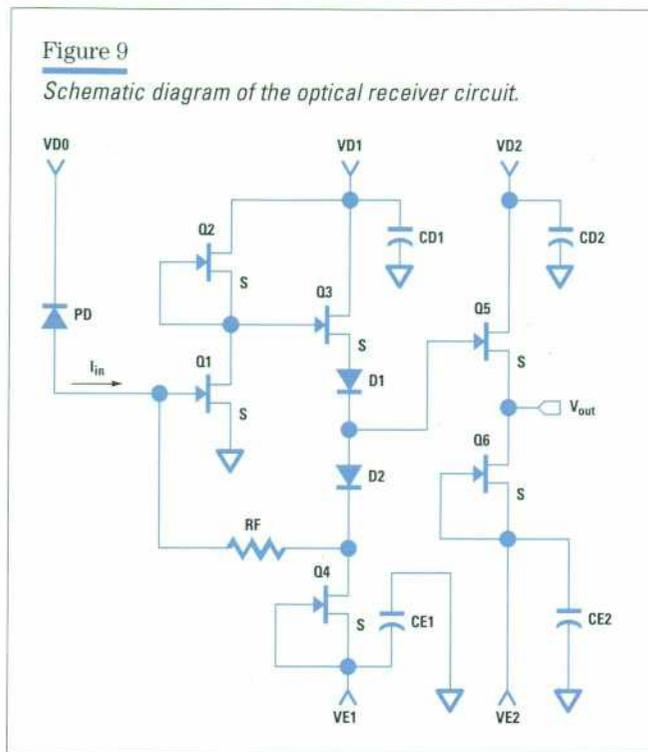


Figure 9

Schematic diagram of the optical receiver circuit.

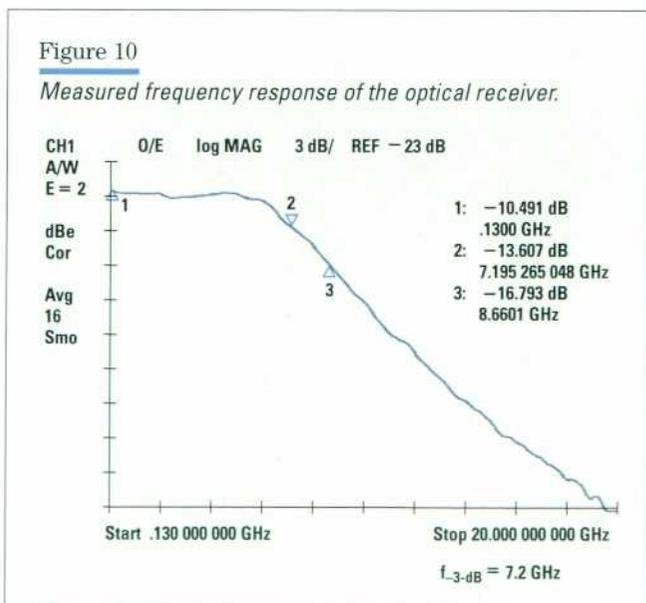


Optical Receiver Performance

An optical receiver was designed and assembled employing a transimpedance amplifier based on a GaAs MODIC (modulation doped integrated circuit) and a 15- μm flip-chip photodiode as shown in **Figure 5**. The schematic diagram of the receiver circuit is shown in **Figure 9**. This

Figure 10

Measured frequency response of the optical receiver.



transimpedance design was chosen to focus on the effect of die attaching techniques on the performance of the receiver. The measured frequency response of the receiver is shown in **Figure 10**. The -3-dB electrical bandwidth of the flip-chip optical receiver was 7.2 GHz and the conversion gain was 560V/W. The input referred noise current spectral density was 10 pA/Hz^{1/2}. In comparison, a similar optical receiver composed of the same GaAs MODIC-based amplifier and a wire-bonded 15- μm photodiode had a bandwidth of 4.2 GHz and an input referred noise current spectral density of 20 pA/Hz^{1/2}. The performance of the receiver is consistent with the improved bandwidth, smaller capacitance, and higher responsivity of the flip-chip detector.

Nine-Channel Optical Receiver

The flip-chip die attachment provides a capability for precise placement of the chip on predefined bump pads. During the reflow process, the surface tension of the molten solder will pull the chip to within a few micrometers of the predefined pads. This feature is very useful in the assembly of a multichannel optical receiver, which requires each detector element to be placed precisely at a predefined location to facilitate optical coupling and reduce cross talk.

We have designed and fabricated a nine-channel silicon bipolar receiver IC that has nine sets of solder bump pads with nine photodetectors flip-chip bonded on them as shown in **Figure 11**. The electrical performance of each channel is tested by supplying a 1-Gbit/s NRZ pseudo-random optical input signal. The electrical output, as represented by the eye diagram shown in **Figure 12**, indicates 1-Gbit/s performance per channel.

Figure 11

Nine-channel optical receiver.

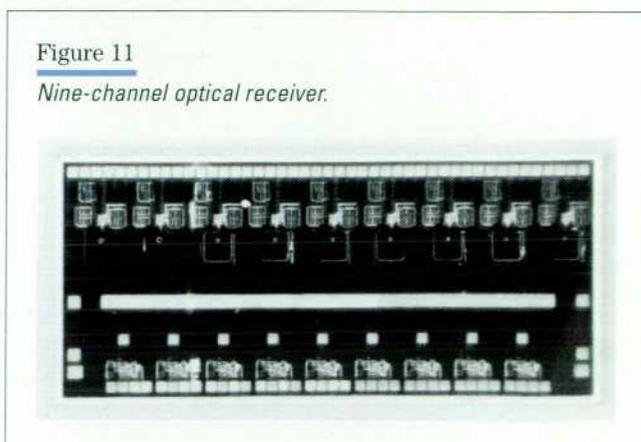
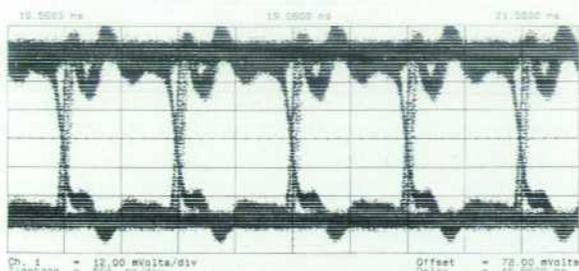


Figure 12

Output eye diagram of one channel of the nine-channel optical receiver with a 1-Gbit/s input signal.



Conclusion

We have successfully incorporated a micro-flip-chip technology with our optoelectronic components. In the flip-chip die attached form, a p-i-n photodetector with an i-layer thickness of 0.75 μm has an average responsivity of 0.8A/W at 1300 nm and a dark current of less than 1 nA. A 7- μm -diameter flip-chip photodetector has a -3-dB electrical bandwidth in excess of 50 GHz. The flip-chip optical receiver with a 15- μm detector has a measured -3-dB electrical bandwidth of 7 GHz while a wire-bonded version has a -3-dB bandwidth of 3.9 GHz. The self-aligning feature of the flip-chip die attachment is very useful in the assembly of a nine-channel optical receiver that operates at speeds in excess of 1 Gbit/s per channel. The same flip-chip technology can be easily combined with a vertical-cavity surface emitting laser (VCSEL) transmitter to realize a flip-chip optical transceiver.

Acknowledgments

The authors wish to thank Bill Imler, Ken Scholz, and Bill Wun for valuable technical discussions, Marie Amistoso, Jeanne Glenn, Yogesh Desai, Kevin Tibbs, Don Cropper, and Marilyn Planting for their technical assistance, Forrest Kellert for epitaxial material, and Ron Kaneshiro for his leadership and support.

References

1. O. Wada, "Ion Beam Etching of InP and Its Application to the Fabrication of High Radiance InGaAs/InP Light Emitting Diodes," *Journal of the Electrochemical Society: Solid-State Science and Technology*, October 1984.
2. B. Imler, K. Scholz, M. Cobarruviaz, R. Haitz, V.K. Nagesh, and C. Chao, "Precision Flip-Chip Solder Bump Interconnects for Optical Packaging," *IEEE Electronic Components and Technology Conference*, 1992.
3. T.S. Tan, R.L. Jungerman, and S.S. Elliott, "Optical Receiver and Modulator Frequency Response Measurement with a Nd:YAG Ring Laser Heterodyne Technique," *IEEE Transactions on Microwave Theory and Techniques*, Vol. MTT-30, 1989, pp. 1217-1222.



Tun S. Tan

Tun S. Tan has been a member of the technical staff of the Solid-State

Technology Laboratory of HP Laboratories since 1989. His recent projects include high-speed optoelectronic devices, high-flux LEDs, and diffractive optics. He received his PhD degree in electrical engineering from Stanford University in 1977. He is married, has three children, and enjoys reading, tennis, and family travel.



David M. Braun

David Braun received his MSEE degree in 1980 from the University of

Wisconsin, specializing in electronic devices. He is a hardware design engineer with HP's Lightwave Division, responsible for semiconductor device design, processing, and testing and optical-mechanical instrument design. He is married, has four children, and enjoys sailing and cycling.



Tim L. Bagwell

A member of the technical staff at HP's Lightwave Division, Tim Bagwell is responsible for optical device develop-

ment and testing, including work on semiconductor lasers for communications applications. He received a BSEE degree in 1980 from the University of Illinois. He is married, has two children, and enjoys outdoor activities and piano.



Christopher P. Kocot

Chris Kocot is a member of the technical staff of HP Laboratories working

in the area of device physics, mainly LEDs and photodetectors. He received his PhD degree in 1979 from Warsaw University and Stanford University in a student exchange program and joined HP in 1980. He is married, has two daughters, and enjoys auto mechanics and swimming.



Susan R. Sloan

Susan Sloan is an R&D project manager for fiber-optic test instru-

mentation with HP's Lightwave Division. She has also managed flip-chip development and other lightwave technology. Her MS degree in physics (1984) is from the University of Colorado. She is married, has two children, and enjoys reading and writing fiction.

Joseph Straznicky

Author's biography appears on page 61.

A 2.488-Gbit/s Silicon Bipolar Clock and Data Recovery Circuit for SONET Fiber-Optic Communications Networks

Richard Walker

Cheryl Stout

Chu-Sun Yen

Lewis R. Dove

Adjustment-free clock and data recovery for 2.488-Gbit/s SONET applications is provided by a 1.77W, $3.45 \times 3.45\text{-mm}^2$ chip implemented in a 25-GHz f_T silicon bipolar process. The chip has an on-chip VCO and operates from 2 to 3 Gbits/s over process, voltage, and temperature variations with a single off-chip filter capacitor. For network monitoring, a highly reliable loss-of-signal detector is provided. For good mechanical, thermal, and RF performance, a custom package was developed using HP's fine-line hybrid process.

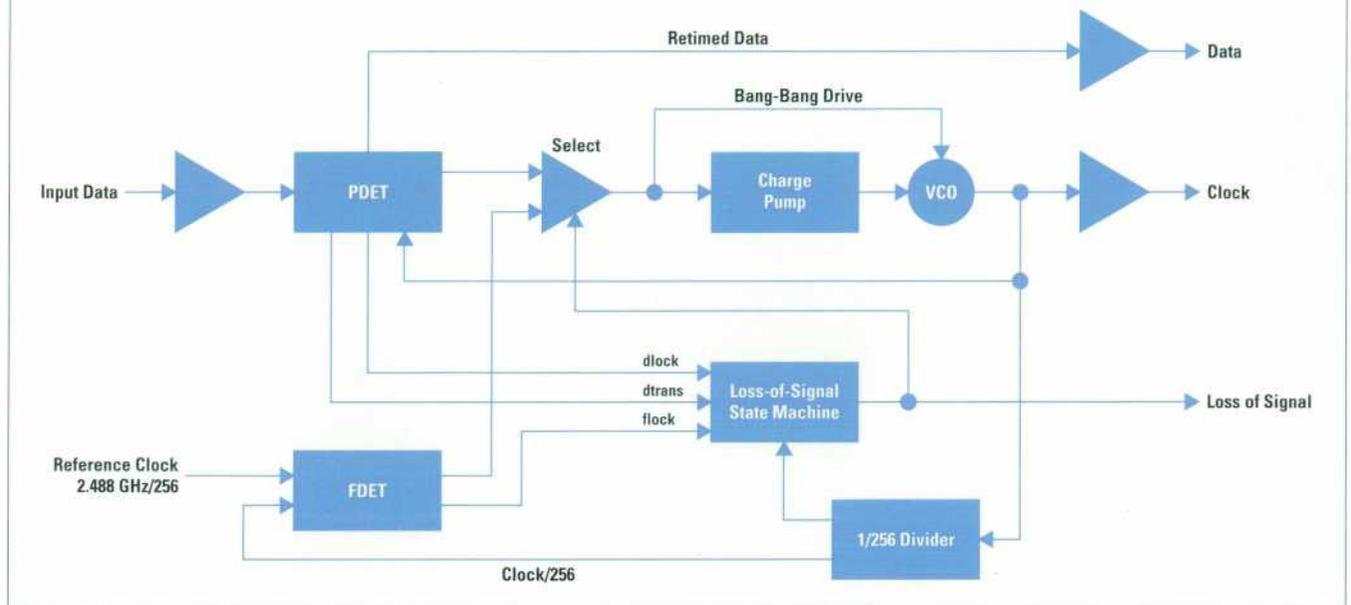
SONET 2.488-Gbit/s transmission and switching systems, network backbones, and video transmission are among the many applications benefiting from inexpensive and robust clock and data recovery circuits. Clock and data recovery circuits are used in high-speed communications systems, typically long-span, single-mode fiber-optic links. Their job is to regenerate clean clock and data signals from arbitrary scrambled data inputs that have been corrupted by jitter and intersymbol interference.

To provide highly reliable clock and data recovery for 2.488-Gbit/s SONET data transmission, the HP CDR 2500 clock and data recovery circuit has been developed. Previous commercial solutions for this application have required multiple chips and GaAs processes.¹ The CDR 2500 was designed in HP's 25-GHz f_T silicon bipolar process,² and incorporates several new circuit ideas developed at HP Laboratories, namely an arbitrary-data phase detector, a reliable loss-of-lock detector, and a monolithic wide-range VCO circuit.

Building on techniques developed for the HP G-Link data communications chipset,³ the CDR 2500 has been optimized for telecommunications needs. The major differences between these two applications are how the data is coded for

Figure 1

Clock and data recovery chip block diagram. FDET = frequency detector. PDET = phase detector.



transmission and how the circuit responds to jitter on the incoming data. In the G-Link application, a specific line code was used that guaranteed periodic transitions in the data. Such a code simplifies the clock recovery job. The telecommunications circuit, by contrast, must operate on scrambled data with random transitions.

Data communications links are typically much shorter than telecommunications links. The extra signal-to-noise ratio in a short link means that time jitter is much less of a problem than in a telecommunications link. Because of this, the HP CDR 2500 design required extra care to meet the more stringent jitter requirements of long-haul SONET applications.

Chip Block Diagram

A simplified block diagram of the chip is shown in **Figure 1**. **Table I** summarizes the chip characteristics. The input amplifier can be driven either single-ended or differentially and is ac-coupled and terminated on the associated thick-film hybrid. A Cherry-Hooper wideband amplifier⁴ at the input minimizes pulse width distortion with single-ended input signals.

Table I

Summary of Chip Characteristics

Parameter	Value
Guaranteed Frequency Range	2 to 3 Gbits/s
Supply Voltage	4.5V to 5.5V
Supply Current (nominal)	340 mA
Power Dissipation	1.77W
Case Temperature Range	0 to 60°C
Die Size (gate array)	3.45 × 3.45 mm ²
Number of Active Devices	3606
IC Process	25-GHz f_T Si Bipolar

The phase-locked-loop portion of the circuit includes two different phase detector blocks. To ensure that the on-chip VCO will operate at 2.488 GHz over variations in processing parameters, it has been designed with a 3:1 frequency tuning range. To avoid locking onto harmonics of the data signal, it is important that the VCO be kept close to the

desired operating frequency. The first frequency/phase detector, labeled FDET in **Figure 1**, is used to frequency-lock the VCO initially to an external reference equal to 1/256 the desired bit frequency. The FDET block also produces a frequency lock error signal (*flock*), which is used by the state machine to sense when the VCO has reached the correct frequency. The second phase detector, labeled PDET, operates on random data and produces three signals: data transition detect (*dtrans*), data lock detect (*dlock*), and a tri-state bang-bang phase error signal.

The loss-of-signal (LOS) state machine block monitors the status outputs of the two detectors. On power-up, the state machine enables the frequency detector (FDET). When the VCO is frequency-locked to the reference and data transitions are detected on the data input, the phase-locked loop is switched to use the bang-bang phase detector. If the data lock detector stabilizes within an allowable time, then the loop remains phase-locked to the data. Otherwise, the VCO is retrained to the reference clock and the sequence repeats.

If the input data is interrupted or corrupted, the VCO needs to be retrained to the reference frequency input. Switching to the reference loop destroys the phase-locked condition between the VCO and the incoming data, causing errors in the retimed data output. In a telecommunications application, these errors could potentially interrupt service for thousands of telephone users. Therefore, a highly reliable algorithm must be used for determining loss of signal, so that the VCO is retrained to the reference only under high bit error rate conditions.

Loss-of-Signal Detection

The robust LOS detector is built by statistically processing the PDET data lock detector output. The data lock detector operates by monitoring the phase relationship between data zero crossings and the VCO output. Any rising or falling edge occurring more than ± 135 degrees away from the nominal location is flagged as a raw phase error.

For typical noisy input signals, large input data phase excursions occur with approximately the same probability as the bit error rate (BER). Random data has half as many transitions as actual bits, so the raw phase error rate can be approximated as BER/2.

If we simply used raw phase errors to determine when to restart the link by relocking the VCO to the reference,

then we would get a very unreliable and sloppy relationship between mean time to restart (MTTR) and BER. For a simple LOS algorithm in which the link is restarted on an isolated phase error event, the link would restart an average of once per second at a BER of 10^{-9} and every 10 seconds at a BER of 10^{-10} . This is a 1:1 relationship between BER and MTTR. In this simple system, an unrealistic BER of 10^{-16} would be required to ensure an MTTR of greater than one year between restarts. For system reliability reasons, we chose as a design target that the link should restart within 1/10 of a second at 10^{-5} BER, and should assert LOS less than once per year at 10^{-7} BER.

A way to steepen the LOS/BER relationship is to process phase error events in multibit bins and require that multiple consecutive bins each contain at least one error before asserting LOS. The current design groups phase error events into one of four programmable bin sizes (N), each requiring a different number of consecutive errored bins (M) before asserting LOS. Assuming that the phase error rate is approximately equal to half the BER, then if T_{bit} is the bit time,

$$\text{MTTR} \approx \frac{NT_{\text{bit}}}{(1 - (1 - \text{BER}/2)^N)^M}$$

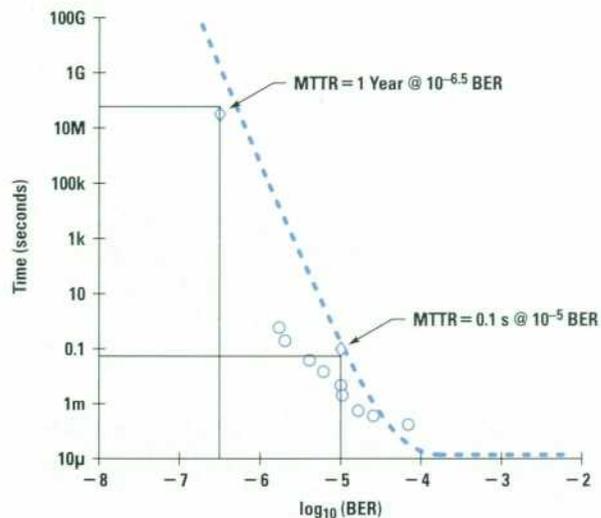
The exponent M in the denominator of the MTTR expression sets the slope relation between phase error and MTTR. Setting M to 7 requires only one decade change in BER to change from an MTTR of one second to one year. Selection of one of the four possible states is programmable by two digital bond options, providing an LOS threshold between approximately 10^{-4} and 10^{-6} BER. The exact BER varies according to the application because of the system dependent relationship between phase error rate and BER. **Figure 2** shows measured MTTR versus calculated MTTR results for a system in which $M = 7$ and $N = 2^{16}$.

VCO

The VCO is composed of a cascade of two variable-delay blocks as shown in **Figure 3**. Each variable-delay block is composed of an amplifier/delay cell and an interpolator. A second interpolation section extends the tuning range while minimizing interpolation jitter. The shaded interpolator in **Figure 3** interpolates between a nondelayed input and an input delayed by the shaded amplifier/delay cell. The shaded elements are recursively used as a delay cell for the unshaded interpolator (with the bang-bang

Figure 2

Calculated and measured relationship between BER and mean time to link restart. Parameters used were $N = 2^{16}$, $M = 7$. The dashed line is a calculation assuming the raw phase error rate is approximately equal to BER/2. The circles are measured data points.



input). This recursive connection minimizes the phase difference between the inputs of the interpolators.

The VCO outputs are taken from noncritical, highly buffered portions of the loop to minimize loading. The 3:1 frequency range of the circuit guarantees 2-to-3-GHz operation across process, temperature, and power supply variations. The main tuning input is prefiltered by two 100-MHz poles to reduce sensitivity to power supply noise. The wide-bandwidth bang-bang tuning input has 500 times less gain than the main tuning input, and is implemented by injecting small currents into the interpolation cell.

The phase-locked loop uses a bang-bang phase detector and a positive feedback charge pump similar to those described in references 3 and 5. The two-tuning-input VCO architecture, in conjunction with a binary-quantized phase detector, results in a VCO drive voltage equivalent to a first-order sigma-delta conversion of the instantaneous loop frequency error. By adjusting the relative gain of the two VCO tuning inputs, a trade-off can be made between jitter generation and jitter accommodation. Unlike a traditional analog loop, such a system cannot be characterized

by a loop bandwidth. The ability of the loop to track incoming phase jitter is a slew-rate limited process, with an effective jitter bandwidth proportional to jitter amplitude. In practice, this behavior is ideal for the input of a SONET regenerator, for which a wide input jitter-tracking bandwidth is desirable to minimize sampling errors, and in which the overall system jitter transfer function will be set by a separate narrowband transmitter phase-locked loop. For this design, the bang-bang amplitude and charge pump time constant have been set to meet the SONET jitter tolerance specification. The resulting jitter generation, as calculated from the jitter spectrum in Figure 4, is 0.0049 UI rms (UI = unit interval = T_{bit}).

Figure 3

(a) Ring oscillator VCO block diagram showing cascaded variable-delay blocks. A basic variable-delay block consists of an amplifier/delay cell and an interpolator (shown shaded). A second interpolator extends the tuning range of each block. (b) Interpolator operation.

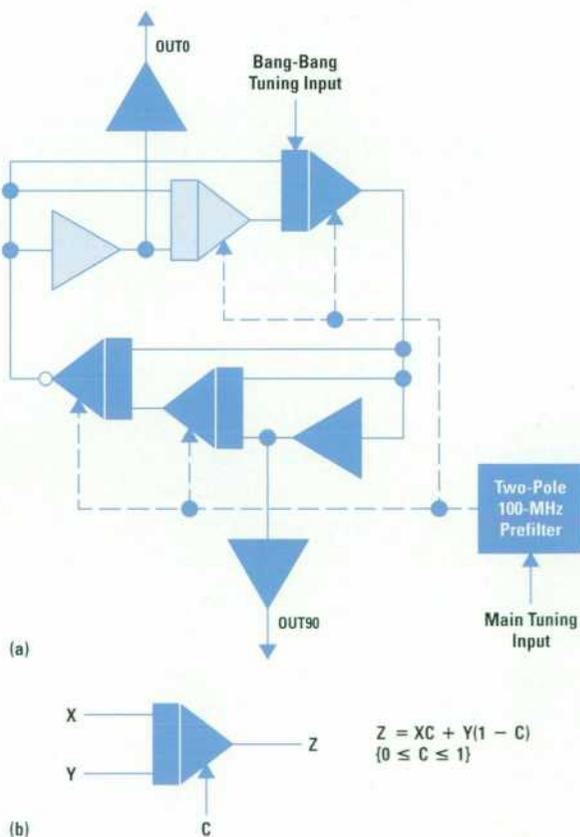


Figure 4

Clock phase noise measured under open-loop conditions and when locked to a $2^{23} - 1$ pseudorandom binary sequence (PRBS). Under locked conditions, the integrated value of the phase noise, filtered by a single-pole 5-kHz high-pass filter and a 100-MHz low-pass filter is 0.0046 UI, or 1.6 degrees rms.

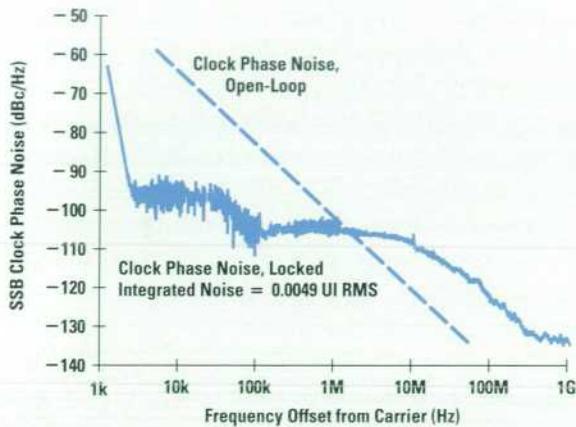


Figure 6

Clock and data recovery chip.

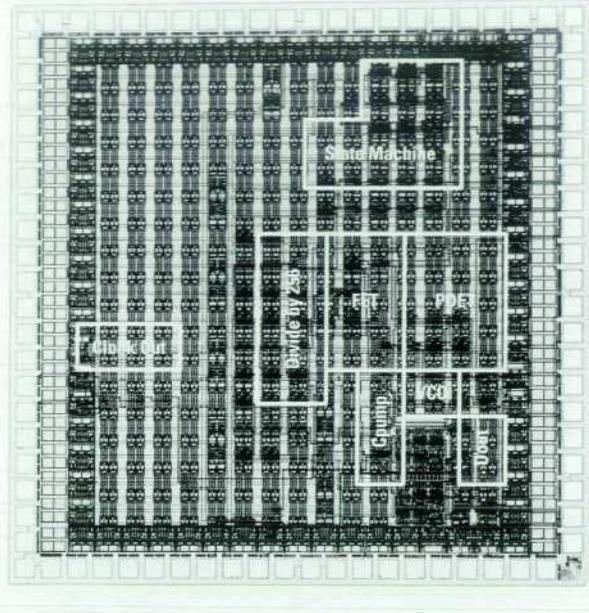


Figure 5 shows the recovered clock and data in the presence of noisy data. 3-GHz white noise was added to a $(2^{23} - 1)$ -bit pseudorandom binary sequence (PRBS) signal until the system was operating at 10^{-4} BER. The

noisy input signal, the recovered data, and the recovered clock are shown in the top, middle, and bottom traces, respectively.

Figure 5

Top trace: Input $2^{23} - 1$ PRBS data with broadband noise added to achieve 10^{-4} BER. Middle trace: Recovered data eye. Bottom trace: Recovered clock. All signals are triggered from the BER tester clock input.

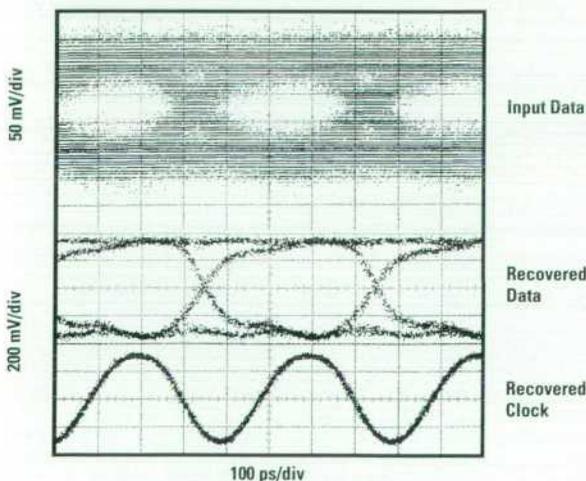


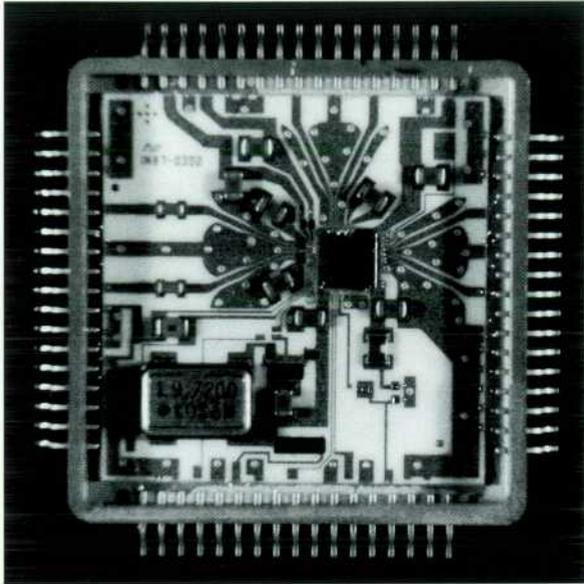
Figure 6 shows the die photo of the chip. The layout was done using a gate array methodology with fully differential ECL cells. Less than half of the array capacity was used in this design.

Package and Substrate

Figure 7 shows a picture of the assembled CDR 2500 multichip module. It consists of the HP-25 clock and data recovery IC attached to the 96% alumina substrate along with two diodes, a discrete transistor, 12 chip capacitors, and a chip resistor. The large packaged device in the lower left corner of the substrate is the reference oscillator. This is a commercially available surface mount crystal oscillator operating at $2.488 \text{ GHz}/256 = 9.71875 \text{ MHz}$. All devices are attached to the substrate using conductive epoxy. The two transistors visible on the hybrid are used to level shift the loss-of-signal output to be compatible with 5V CMOS. 0.001-inch gold wire bonding is used to connect to the die. The 1-in-by-1-in substrate is attached to a 68-pin Kovar package using conductive epoxy. The package leads are formed into a gull-wing configuration so that the package

Figure 7

Clock and data recovery multichip module.



can be soldered to a printed circuit board. The module is tested at speed (2.488 Gbits/s).

The substrate is fabricated using proprietary HP enhanced thick-film processes. These processes make the high-speed performance of the packaged clock and data recovery circuit possible. They provide high-frequency performance and density rivaling what can be obtained using thin-film processing at a fraction of thin-film's cost.

Clock and data recovery circuits require excellent input return loss to prevent intersymbol interference caused by reflections at impedance discontinuities. The target for the CDR 2500 was to achieve better than 15 dB return loss up to 2.5 GHz. However, the discontinuities caused by the package's gull-wing lead and by the wire bonds that connect the substrate input and output lines to the package pins eat up much of the available return loss specification. Therefore, the return loss of the thick-film transmission lines and the input 50-ohm termination resistors had to be even better. That meant that tight control of the transmission line impedances was critical. To minimize high-frequency parasitics, it was important to use physically small, tightly controlled termination resistors.

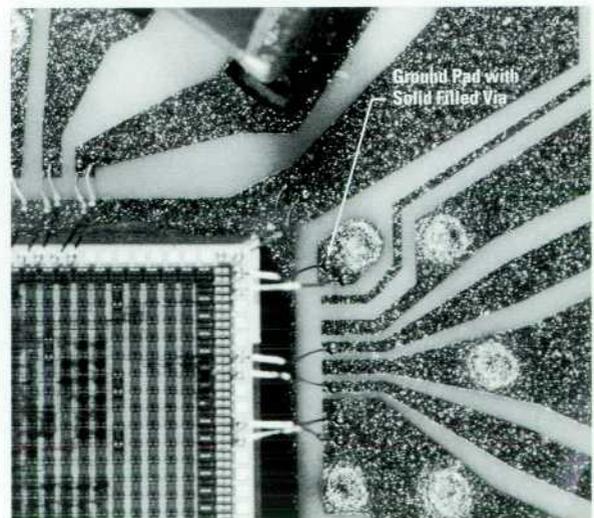
The three HP proprietary thick-film processes used in fabricating the HP CDR 2500 are:

- Fine-line etched gold conductors
- Solid filled vias
- Small thick-film resistors.

The fine-line etched gold is the predominant pattern covering the substrate in **Figure 7**. Gold thick-film paste is printed across the substrate and is etched using a photolithographic process to produce the conductor pattern. The fine-line gold has several advantages over standard thick-film printing. It resolves very small dimensions—down to 0.002-inch lines and spaces. This can be seen in **Figure 8**, which shows a closeup of the area of the substrate around the pad for the HP-25 die. Several of the gaps between conductors leading to the die are 0.002 inch wide. This permits high pattern densities near the IC, making it possible to minimize wire-bond length and bringing the bypassing capacitors very close to the chip. Both of these factors help to minimize noise in the circuit. Having short, consistent wire-bond lengths is critical for achieving good return loss. Also, the fine-line etch process permits much better control of conductor line widths and spaces, needed to maintain accurate 50-ohm transmission line impedances. The etched fine-line conductors also have straight, well-defined side walls, which offer good high-frequency performance.

Figure 8

A close-up of the clock and data recovery module thick-film circuitry near the HP-25 die pad.

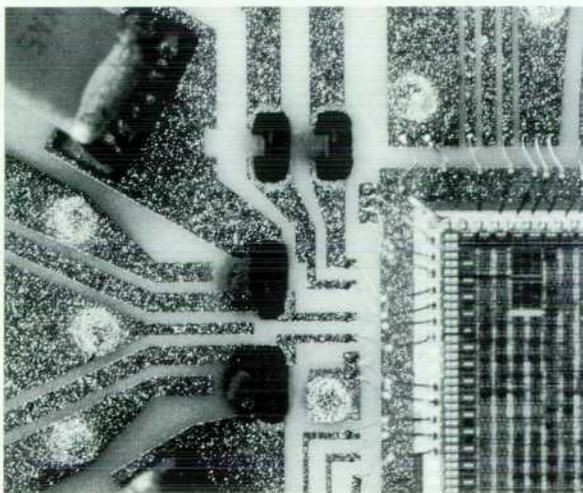


Solid filled vias can be seen in the ground regions in **Figure 8**. The vias are used to connect ground pads on top of the substrate to the backside ground plane. The solid filled via is formed by printing a solid plug of paste into the hole, filling the hole completely. This provides a very consistent, low-resistance ($\ll 25$ milliohms), low-inductance path to ground at RF frequencies. In **Figure 8**, there is a ground pad that has a solid filled via near the upper right corner of the pad for the IC. Note how small the gold pad is—only about 0.002 inch larger than the via itself. Two bond wires are bonded within 0.002 inch of the via.

A closeup of the input section near the IC pad is shown in **Figure 9**. The four black elements are thick-film resistors. These are the small HP resistors, which can be as small as 0.010 inch by 0.010 inch, very small for thick-film technology. The two lower resistors are 0.010-by-0.012-inch, 50-ohm termination resistors for the circuit's input lines. They are laser trimmed to $\pm 0.5\%$ tolerance. Their small size, which minimizes their high-frequency parasitics, and their tight tolerance make them excellent terminations. The two resistors above the terminations are 0.010-by-0.012-inch 1-kilohm bias resistors. They are made small to maximize circuit density around the chip and keep bypass circuitry as close to the die as possible.

Figure 9

A close-up of the clock and data recovery module's thick-film input section showing small resistors (black elements).



Substrate Technology Selection

Why is the HP proprietary thick-film on ceramic technology used for this application? Why couldn't a laminate technology, which is generally slightly less expensive, be used instead? Primarily, because it would be difficult to achieve the electrical performance. A laminate would have difficulty holding the tight transmission line tolerances that are required. If a laminate were used, the ten resistors now implemented using thick-film technology couldn't be integrated into the substrate. In general, this would increase cost and reduce the circuit density around the chip, thereby moving bypassing farther from the IC and degrading the electrical performance. For the input termination resistors in particular, using add-on resistors would add parasitics, making the return loss specification even harder to meet. Keeping the size of the circuit as small as possible is important, and using add-on resistors would probably necessitate an increase in the size of the device.

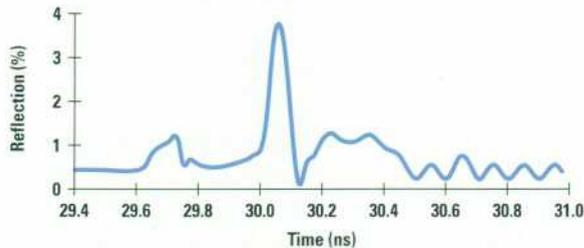
The HP-25 die dissipates 1.8W. For the original application of the HP CDR 2500, the circuit sits in a Kovar package in an environment with no air flow and no room to attach a heat spreader or a heat sink. Thus, the thermal performance of the device is important, and using a laminate instead of ceramic for the substrate would substantially increase the die's junction temperature. The die must be back-biased to $-5.2V$, so cutting a hole in the substrate and attaching the die directly to a pedestal brazed to the package is not an option. Attaching it to an electrically insulating heat slug is possible, but the added cost of this approach would negate some of the cost savings of the laminate.

Design

To determine that thick-film fine-line transmission lines could be controlled well enough to meet the CDR 2500's return loss specifications, a Monte Carlo analysis was run on the HP Microwave Design System (HP MDS). This analyzed the change in transmission line characteristic impedance as a function of the inherent variations in HP's fine-line conductor process. Changes in conductor line width, substrate dielectric constant, and substrate thickness are the primary factors that affect characteristic impedance. HP MDS varies these parameters, and the result

Figure 10

A time-domain reflectometer measurement of one of the input lines of the clock and data recovery multichip module.



is a probability density function and a standard deviation for the resulting changes in characteristic impedance. The changes in characteristic impedance were put into a SPICE simulation of the clock and data recovery circuit's input. This confirmed that a fine-line solution could produce acceptable results in a manufacturing environment.

HP MDS was used to design the multichip module's critical transmission lines, after previously correlating the MDS models to measured data of the HP processes. The design of the HP CDR 2500 multichip module worked the first time it was prototyped. **Figure 10** is a time-domain reflectometer (TDR) display of one of the prototype input transmission lines, measured using a 160-ps pulse. The vertical scale is 2% reflection per division. The first inductive discontinuity (at 29.7280 ns) is caused by the test fixture used to probe the package. The second and largest inductive discontinuity is the wire bond leading from the package pin to the substrate. The region immediately following that (for nearly two time divisions) is the microstrip and grounded coplanar waveguide transmission line. It shows very little reflection, and there is no discontinuity at the 0.010-inch 50-ohm termination resistor.

The production version of the HP CDR 2500 replaces the package-to-substrate wirebonds with 0.001-by-0.005-inch ribbon bonds. This lowers the inductive discontinuity and further improves the module return loss.

Trends in Serial Communication

Long-haul transmission capacity has traditionally grown by expansion of bit rates. The last 20 years has seen data rates jump by over three orders of magnitude. The need for higher rates is motivated by the potential capacity of

existing fibers and by the high cost of installing new fiber-optic cables.

Recent development of fiber broadband optical amplifiers and workable wavelength-domain multiplexing (WDM) techniques have somewhat changed this trend. Because WDM allows multiple high-speed channels to share the same fiber, there is less of a need for heroic efforts to extend the serial data rate of each channel. Pulse dispersion at high bit rates and the limitations of electrical circuits make it likely that individual link speeds will top out at 10 to 40 Gbits/s for the foreseeable future.

High-speed clock and data recovery circuits will still be needed for each wavelength channel, and the market demand for such circuits will continue to grow in relation to world data transmission needs.

As the industry becomes more mature, we will see increasing pressure on vendors to increase the integration level of their link ICs by including limiting amplifiers, demultiplexers, and payload processing circuitry on their chips. This is starting now at lower data rates (e.g., 155 Mbits/s), and will proceed up to higher rates as IC processes improve.

The relatively uncritical data communications needs will likely be met adequately with future CMOS processes. Gigabit-rate clock and data recovery circuits can already be built using multiphase sampling in current CMOS processes. Such architectures achieve high speed at the cost of systematic jitter caused by mismatched clock phases. For jitter-sensitive telecommunications applications, it is likely that bipolar processes will be needed for at least the next five to ten years. At the end of that time, bipolar implementations will likely be less desirable, the reason being that, although bipolar processes provide superior front-end clock and data recovery performance, they are not capable of the extremely high integration levels required for payload processing. A compromise approach would be to add relatively low-performance CMOS devices onto an existing high-speed bipolar process. Such a process would then be an ideal compromise between the precision high-speed, low-jitter front-end clock and data recovery requirements and the need for large-scale digital payload processing on the same chip.

Conclusion

The CDR 2500 is a step forward towards higher integration density for SONET 2.488-Gbit/s systems. HP's 25-GHz f_T bipolar process provides good jitter performance at a lower cost than previous multichip GaAs implementations for this data rate. The HP CDR 2500 takes advantage of proprietary HP thick-film technology that includes etched fine-line gold conductors, solid filled vias, and small resistors to provide excellent RF and high-speed digital performance at a lower cost than thin-film technology.

Acknowledgments

The CDR 2500 clock and data recovery multichip module was a collaborative effort between HP Laboratories, the HP Colorado Springs Technology Center (CSTC), the HP Fiber-Optic Components Operation (FCO), and the HP Optical Communications Division (OCD). Dr. Soo-Young Chai of Samsung (formerly of AT&T) provided valuable customer input for the design. The authors thank Jean Norman for her support in hybrid microassembly. Pat Petruno, Richard Dugan, and Virginia Graham of OCD were invaluable in allowing us to make prototype fab runs

on the OCD Masterslice HP-25 gate array. We also wish to thank Teik Goh of FCO and Jack Casey, Steve Meyer, Jim Taylor, Steve Chaffee, Randy Fuchs, Don Griffith, and Albert Yeh of CSTC for their help in transferring this design into production.

References

1. H. Ransijn and P. O'Connor, "A PLL-Based 2.5-Gbit/s GaAs Clock and Data Regenerator IC," *IEEE Journal of Solid State Circuits*, Vol. 26, no. 10, October 1991, pp. 1345-1353.
2. W.M. Huang, et al, "A high-speed bipolar technology featuring self-aligned single-poly base and submicrometer emitter contacts," *IEEE Electron Device Letters*, Vol. 11, no. 9, September 1990, pp. 412-414.
3. C.S. Yen, R. Walker, P. Petruno, C. Stout, B. Lai, and W. McFarland, "G-Link: A Chipset for Gigabit-Rate Data Communication," *Hewlett-Packard Journal*, Vol. 43, no. 5, October 1992, pp. 103-116.
4. D. Faulkner, "A Wide-Band Limiting Amplifier for Optical Fiber Repeaters," *IEEE Journal of Solid State Circuits*, Vol. 18, no. 3, June 1983, pp. 333-340.
5. B. Lai and R.C. Walker, "A Monolithic 622Mb/s Clock Extraction Data Retiming Circuit," *International Solid State Circuits Conference Digest*, 1991, pp. 144-145.



Richard Walker

Rick Walker joined HP Laboratories in 1981, and is currently a principal project engineer. He develops and designs integrated circuits for gigabit-rate serial data transmission. He received a BS degree from the California Institute of Technology in 1982 and an MS degree from the California State University at Chico in 1992. His interests include tai'chi chuan, playing guitar, and raising exotic plants.



Cheryl Stout

Cheryl Stout is a hardware design engineer at HP Laboratories. She currently develops CMOS clock and data-recovery circuits. In her 14 years with HP Laboratories, she has designed gallium arsenide and silicon integrated circuits for gigabit-rate serial data transmission. She received a BSEE degree from the California State University at San Jose in 1979 and an MSEE degree from the University of California at Berkeley in 1983. Her interests include mountaineering and natural history.



Chu-Sun Yen

Chu-Sun Yen has done research and development at HP since 1961. He was project manager for both the HP CDR 2500 clock and data recovery circuits and the G-link chipset, and has authored 20 papers on circuits, phase-locked loops, instruments, and communications links. He holds a doctorate from Stanford University.



Lewis R. Dove

Lew Dove is an engineer-scientist at HP's Colorado Springs Technology Center, where he works on thick-film packages for high-frequency multichip modules. He has BSEE and MSEE degrees from the University of Arizona. He is married and has two sons, and his hobbies include flying, running, and skiing.

Testing Erbium-Doped Fiber Amplifiers

James R. Stimple

EDFAs can overcome losses in long fiber-optic links independent of the digital bit rate, and can amplify multiple signals in a wavelength-division multiplexed (WDM) system architecture. As more and more EDFAs are deployed, designers add new features, creating a need for more sophisticated testing. This paper provides a brief survey of the tests required to characterize EDFAs.

A small child in California calls her grandmother in Germany. To her, the call seems effortless and Grandma sounds as if she is right there. A teenager in Arkansas is surfing the Net to find out about the latest in skateboard parks for a trip planned for the upcoming summer. A retired auto worker in Japan checks e-mail to see if a message from an old coworker has arrived confirming their upcoming fishing trip. These common, everyday events are only possible because of the vast deployment of capacity on the global communications network.

Fiber-optics technology is the heart of this "information superhighway" which spans the globe. Its tremendous capacity for carrying information in digital form is revolutionizing the way people live and work. Linked with cellular phone systems and computers around the world, this network provides access to incomprehensible amounts of information and vast connectivity. The most recent element adding to the success of these systems is the fiber-optic amplifier.

Erbium-Doped Fiber Amplifiers

The development of erbium-doped fiber amplifier (EDFA) technology has greatly changed the design methodology of fiber-optic system designers. Traditional fiber-optic systems used regenerative repeaters to boost the signals, as shown in **Figure 1a**. When the length of the link exceeded the practical single-span passive limit, these regenerative repeaters detected the signal and retransmitted it with a laser, restoring the signal level as well as the signal

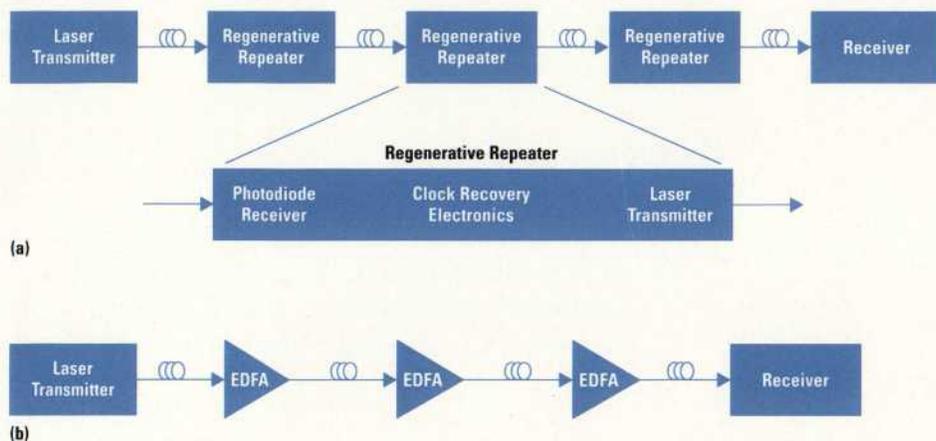


James R. Stimple

Jim Stimple is a project manager for erbium-doped fiber amplifier systems at the HP Lightwave Division. He has worked for HP since 1978. He received his BSEE degree from Northwestern University in Evanston, Illinois in 1974.

Figure 1

(a) Lightwave system with regenerative repeaters. Gain is provided by the electronics and each regenerative repeater is matched to the data rate of the system. (b) Lightwave system with erbium-doped fiber amplifiers (EDFAs). The amplifiers boost the signal independent of the data rate and allow multiple wavelengths.



fidelity. Although these regenerative repeater systems worked well, they were very expensive, and once installed, the capacity of the link was fixed.

With the development of EDFAs, the link losses could be overcome by amplification as shown in **Figure 1b**. Unlike the regenerative repeater systems, these "transparent" amplified systems are independent of the digital bit rate. This feature allows an upgrade path to higher bit rates as solutions to other limiting factors such as chromatic and polarization-mode dispersion become available. EDFAs are also able to amplify multiple signals in a wavelength-division multiplexed (WDM) system architecture, adding another dimension to the capacity equation.

Many point-to-point terrestrial links are being upgraded from regenerative repeaters to amplified links because of the high cost of laying more fibers in the ground. In many cases, adding sections of dispersion-compensating fiber with each amplifier can allow upgrades in bit rate as well as the possibility of WDM. Amplified systems will soon use WDM not only for increased capacity, but as a means for information routing, eliminating the need for expensive high-speed demultiplexing and remultiplexing at each optical node.

To date, the major emphasis of EDFAs has been in the high-capacity portions of the network, but as the cost of EDFAs comes down, they will also be deployed in the

subscriber loop. Here the emphasis will be on WDM to allow single users high-speed access to the network. Thus, the EDFA is rapidly becoming the workhorse of the system of fibers that spans our globe.

How the EDFA Works

Figure 2 shows the simplest block diagram of an EDFA. The erbium-doped fiber is a silica-based fiber waveguide with a high concentration of erbium atoms. The presence of the erbium provides for ionic transitions leading to photon emission in the 1530-to-1570-nanometer wavelength range. Pump lasers at one or more of the absorption wavelengths (typically 1480 nm or 980 nm) provide

Figure 2
EDFA block diagram.

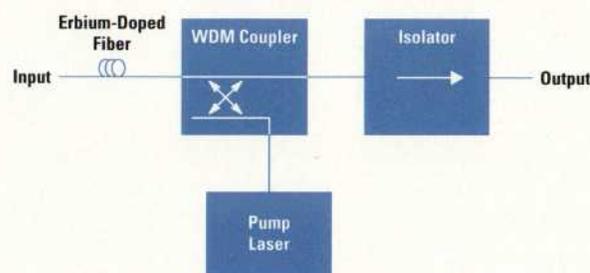
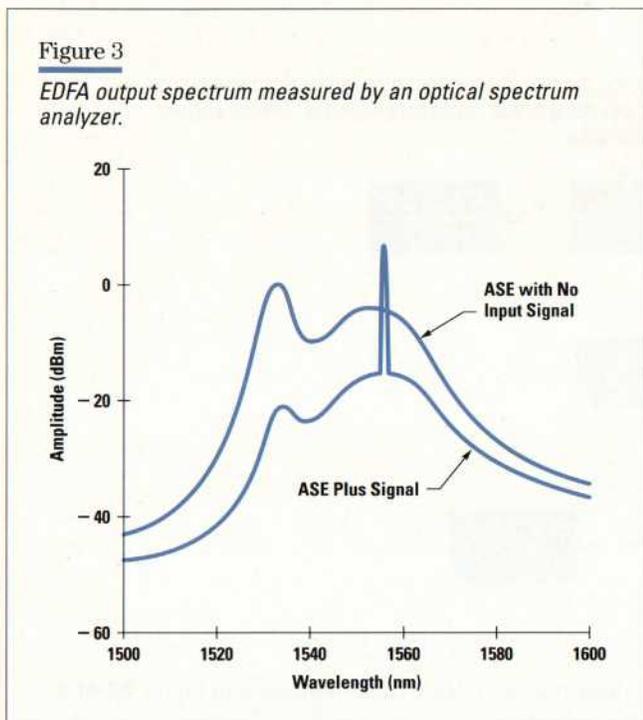


Figure 3

EDFA output spectrum measured by an optical spectrum analyzer.



the excitation for the emission process. The fiber length is typically around 70 meters and amplifier small-signal gains of 35 dB are common.

Figure 2 shows the pump laser in the counterpropagating configuration, in which the pump energy travels in the direction opposite to the signal. Amplifiers can be pumped in the copropagating direction by placing the WDM on the input side. Many amplifiers have both copropagating and counterpropagating pumps for redundancy.

Without an input signal, the EDFA is a source of spontaneously emitted photons. The wavelength spectrum of this spontaneous emission process is determined by the statistical distribution of the energy bands in the erbium atoms. Spontaneous emission photons are of random phase, random direction, and random states of polarization. As these spontaneous photons travel along the fiber, they can replicate through the process of stimulated emission. This process creates a second photon of the same wavelength, phase, polarization, and direction as the first. The total output spectrum originating from spontaneous emission photons of the EDFA is called the *amplified spontaneous emission (ASE)* and is shown in Figure 3.

With a laser signal applied at the amplifier input, many of the would-be spontaneous emission photons from both

the forward and reverse directions in the fiber are stimulated by photons from the laser. This not only provides the required amplification, but reduces the ASE as shown in Figure 3.

The parametric nature of the amplification provides a nonlinear gain characteristic. EDFAs are generally operated in the compression region. This is highly desirable for long communication links because the amplifiers provide natural leveling of the signal. An additional feature of EDFAs is their long time constants for the absorption and emission processes. These are typically from 100 to 500 microseconds. Since the lasers are modulated with 2.5-gigabit/second data or higher, there is virtually no distortion of the information being transmitted. This attribute also proves highly desirable in a WDM system because it eliminates the possibility of cross-modulation products in the amplifier, which would be devastating to the system integrity.

Testing EDFAs

As more and more EDFAs are deployed, designers are adding new features to the basic design. Amplifiers are basically custom designed to meet the particular application requirements. As the complexity of these amplifiers evolves, so does the need for more sophisticated testing. At Hewlett-Packard Laboratories and the HP Böblingen Instrument and Lightwave Divisions, significant efforts are ongoing in test method research and the development of test and measurement systems for EDFAs.

Basic Gain and Noise Figure

The basic measurement requirements for an EDFA are simply gain and noise figure. How much amplification does it provide and how much degradation of the signal-to-noise ratio might be caused by the amplifier in a given application? The diagram in Figure 4 shows the basic measurement setup. The source is generally a tunable laser to allow a variety of wavelengths and powers for the

Figure 4

Basic setup for measurements of EDFA characteristics.



input signal. The receiver is an optical spectrum analyzer to provide wavelength selection (resolution) of the output signal and ASE.

The amplifier's nonlinearity complicates the measurements, since the gain is a function of the input signal level and wavelength. The noise figure measurement is further complicated by the fact that the source has a spontaneous emission spectrum of its own in addition to the dominant laser spectral line.

To accurately measure the gain characteristics of the amplifier, one must first specify the wavelength and power of the input signal. Usually the power and wavelength are both a range of desired values. Measuring the output signal level under each of these conditions will give the gain. The gain G in each case is simply the output signal power divided by the input signal power. For low signal levels, it is important to subtract the spontaneous emission from the measured output signal to obtain the true output signal power.

The noise figure is calculated from an additional measurement of the output ASE and the equation:

$$NF = ASE/h\nu G + 1/G,$$

where ASE is expressed in watts/Hz, h is Planck's constant, and ν is the optical frequency in hertz.

The ASE measurement is difficult because the spontaneous emission spectrum of the laser is amplified and can add to the ASE of the amplifier. Elimination of this effect can be accomplished in one of three ways:

- *Interpolation source subtraction.* The level of spontaneous emission at the signal wavelength is interpolated (averaged) from measurements on either side of the signal and is subtracted from the signal.
- *Polarization extinction.* This technique uses the fact that the signal from the laser and its spontaneous emissions are polarized and the ASE in the EDFA is not polarized. By placing a polarizer in front of the optical spectrum analyzer and adjusting the polarization to null out the signal, the spontaneous emission of the tunable laser source is extinguished from the measurement, leaving ASE/2 received at the optical spectrum analyzer.
- *Time-domain extinction.* This is similar to polarization extinction in that the effects of the source are eliminated rather than subtracted. This method uses the fact that the EDFA recovery has a long time constant. Modulating

the tunable laser source on and off rapidly compared to the recovery time and measuring the ASE when the tunable laser source is off provides the extinction of source spontaneous emission.

Other measurements that are generally included along with the gain and noise figure are the output signal power (required for gain), the integrated ASE, and the total output power (signal plus the integrated ASE).

Figure 5 shows a number of measurements using the HP 83465A EDFA test system on a Fitel ErFA1313 amplifier. Figures 5a and 5b show the large-signal gain and noise figure as functions of wavelength from 1520 nm to 1570 nm at an input power level of -10 dBm. The graphs in Figure 6 show the effect of varying the input power level at a single wavelength of 1550 nm. Figure 6a demonstrates the amplifier saturation by showing the asymptotic limit of the output signal power as the input level is increased. Figure 6b shows the same effect on the gain, where there is a gradual decrease in the gain from its small-signal (linear) value. Figure 6c shows a decrease in the ASE power with increasing signal level, indicating a

Figure 5

Large-signal gain (a) and noise figure (b) of an EDFA as functions of wavelength from 1520 nm to 1570 nm at an input power level of -10 dBm.

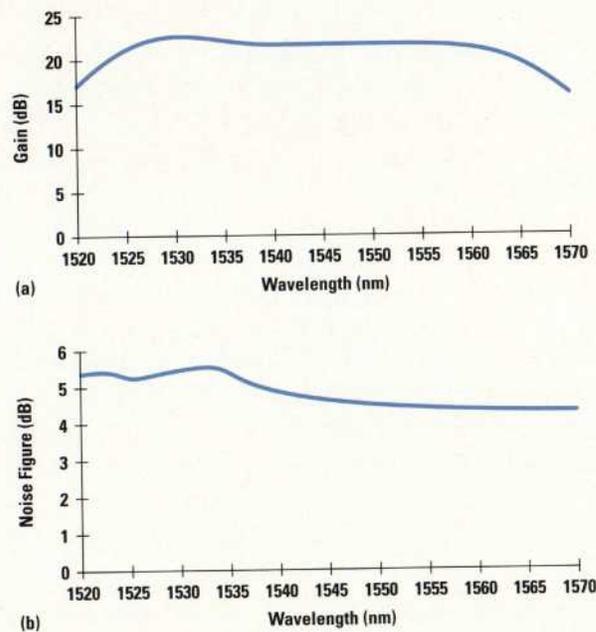
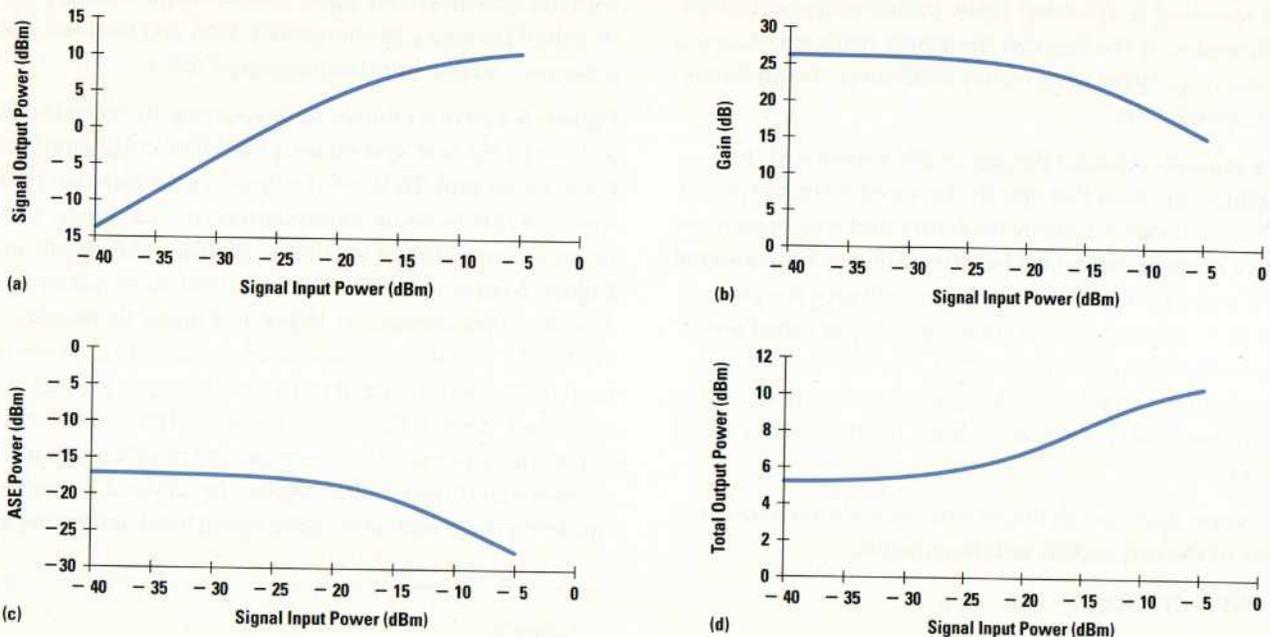


Figure 6

Effects of varying the input power level at a single wavelength of 1550 nm. (a) Signal output power, showing amplifier saturation, the asymptotic limit of the output signal power as the input level is increased. (b) The gain shows the same effect. (c) The ASE power decreases with increasing signal level, indicating a better signal-to-noise ratio. (d) The total output power shows a relatively small change (3 to 5 dB) as the signal input power changes over a very large range (35 dB).



better signal-to-noise ratio. **Figure 6d** shows the relatively small change (3 to 5 dB) in the total output power as the signal input power changes over a very large range (35 dB). This is because the energy for the signal amplification comes from the same source (the pump laser) as the ASE in the unsaturated case.

Small-Signal Gain

Because the EDFA is nonlinear, a simple measurement of the large-signal saturating laser gain may not completely characterize the amplifier. If one were to take a small-signal probe laser and measure the gain at various wavelengths in the presence of the saturating laser, this would give the small-signal gain $g(\lambda)$. This small-signal gain can also be measured with a spontaneous emission source instead of the probe laser. In this case $g(\lambda)$ is called the *noise gain profile*.

In the case of a single-channel system, $g(\lambda)$ gives the wavelength of the gain peak, which tells the designer the best operating wavelength to avoid a buildup of ASE.

In the case of a WDM system, $g(\lambda)$ tells the user how the multiple laser channels will share the available gain of the amplifier. It can also be used to characterize an effect called *spectral hole burning*, which is a depletion of gain over a narrow wavelength range caused by a large saturating laser. By measuring $g(\lambda)$ at a matrix of pump currents, saturating wavelengths, and input powers, a very complete model of the EDFA can be obtained that will predict its behavior in virtually any system environment.

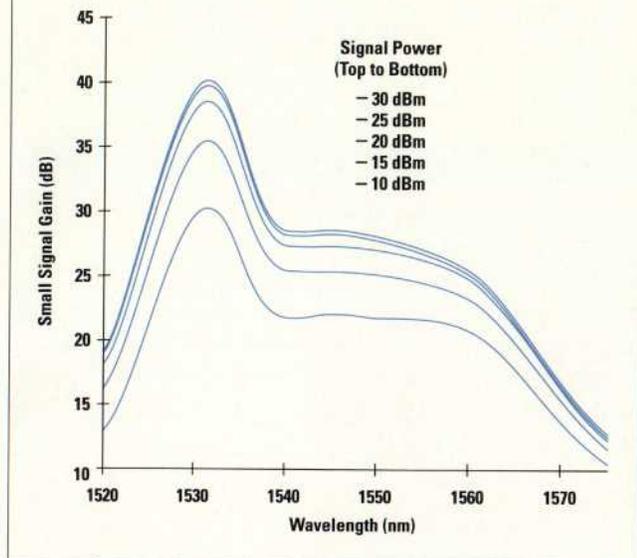
Figure 7 shows a measurement of the noise gain profile at a single saturating wavelength of 1550 nm and a variety of saturating powers. Note the difference between $g(\lambda)$ at -10 dBm shown in **Figure 7** and the large-signal gain variation with wavelength shown in **Figure 5a**. **Figure 5a** is much flatter because of the change in saturation level as the signal wavelength is tuned.

Polarization Effects

The polarization of the laser light always creates challenges in lightwave measurements. Since most systems

Figure 7

Noise gain profile at a single saturating wavelength of 1550 nm and a variety of saturating powers.



use standard circularly symmetric fibers, the polarization of the signals is not preserved. In measuring the gains previously mentioned, it is desirable to measure the *unpolarized* gains. Since the lasers used to saturate the amplifier are polarized, this can be accomplished using a technique called *polarization randomizing*. Varying the state of polarization with time and averaging the measurements over a sufficiently long time window yields a measurement that converges to the unpolarized gain.

In addition to the average gains, a user may want to know the magnitude of the changes in gain and noise figure resulting from changes in the state of polarization of either the saturating signal or the probe signal. These effects are of great importance in systems with many cascaded amplifiers. This variation is measured by simply recording the maximum and minimum gains over the time window.

Return Loss and Isolation

Since the EDFA has the potential for gain in both directions, many systems require splicing the amplifiers into the system or the use of high-return-loss angled connectors. Verification of the input and output return loss and the reverse isolation is another test requirement for EDFAs.

Return loss is measured using a fused fiber coupler with good directivity. This separates a sample of the forward propagating wave from the reverse propagating wave.

Isolation is simply a gain measurement in the reverse direction. However, since the gain is nonlinear, the simplest way to make this measurement is to measure reverse gain at low saturation (without a forward propagating wave) and correct the gain for the ratio of the saturated and unsaturated small-signal gains.

In addition, small reflections inside the EDFA can cause multipath interference. This is most apparent in cable TV applications of EDFAs. Verification of the absence of amplifier internal reflections is a very complicated procedure beyond the scope of this paper.

Monitors

Many EDFAs have features called monitors, which are simply optical signal-level taps. They detect the level of light at the input, at the output, and at the output in the reverse direction. These monitors are used for fault detection and signal-level monitoring and control. A complete measurement system will characterize and verify the operation of these features. A digital voltmeter is used as the sensing apparatus and the tunable laser source power and wavelength are varied as total output power measurements are made by the optical spectrum analyzer.

Device Setup

In many cases the test system provides the interfaces to the EDFA pumps by means of a laser diode controller. These lasers are mounted on a thermoelectric cooler system with a thermistor as the temperature sensing element. The controller regulates the laser temperature by closing the control loop to maintain the thermistor resistance constant. The controller also supplies the laser with bias current. Sometimes there is a monitor photodiode on the back facet of the laser to monitor the laser power. In this case the laser power can be leveled by monitoring the photocurrent and controlling the laser current.

Parametric Measurements

When the test system is used to set up the amplifier, the system operator is able to make parametric measurements. For example, the gain and noise figure can be measured as a function of pump laser current or pump laser temperature. This is very important for determining

the optimum parameters for best amplifier performance. Optimization is becoming more important in the case of amplifiers used for WDM applications since the management of signal-to-noise ratio is more critical because of the lower power of each channel.

A good example of optimization is in the case of an amplifier with a built-in filter to equalize or flatten the naturally rounded peak of the gain. Since the gain slope changes with both pump power and input drive level, for a fixed input drive level there is an optimum pump power for best gain flatness. The optimum pump power setting can be determined by making multiple small-signal gain measurements as a function of pump power setting. On a plot of gain slope (the derivative of small-signal gain) versus pump power, the optimum setting is the power where the slope is zero.

Conclusion

As the market for EDFAs matures, the designs are becoming more complex. Along with this complexity comes the

need for sophisticated test methods that are fast and accurate, yet able to verify all of the features of the product for utmost reliability. By supplying test systems, tools for research, and consulting for the latest in measurement techniques, HP test and measurement groups are doing their part to ensure that the manufacturers of EDFAs have the testing capacity to meet their rapidly increasing demand.

Bibliography

1. C.R. Giles and E. Desurvire, "Modeling Erbium-Doped Fiber Amplifiers," *Journal of Lightwave Technology*, Vol. 9, no. 2, February 1991.
2. D.M. Baney and J. Stimpel, "WDM EDFA Gain Characterization with a Reduced Set of Saturating Channels," *IEEE Photonics Letters*, Vol. 8, no. 12, December 1996, pp. 1615-1617.
3. H. Chou and J. Stimpel, "Inhomogeneous Gain Saturation of Erbium-Doped Fiber Amplifiers," *Conference on Optical Amplifiers and their Applications Technical Digest*, 1995, paper Th E1-1.

Hewlett-Packard Professional Books

This series is published by PTR Prentice Hall, and copies are available in bookstores worldwide. Listed below are the books published in 1997. A list of all books with summaries, bookstores, and corporate sales order information is available on the HP Press web page at URL: <http://www.hp.com/go/retailbooks>

Networks/Networking

Greenberg, J. R., *A Methodology for Developing and Deploying Internet and Intranet Solutions*, Prentice Hall PTR, December 1997, ISBN 0-13-209677-3.

Orzessek, Michael and Sommer, Peter, *ATM & MPEG-2: Integrating Digital Video into Broadband Networks*, PTR Prentice Hall, August 1997, ISBN 0-13-243700-7.

HP-UX System Administration

Poniatowski, Marty, *HP System Administration Handbook and Toolkit* (book/CD-Rom package), PTR Prentice Hall, November 1997, ISBN 0-13-905571-1.

Poniatowski, Marty, *Windows NT and HP-UX System Administrator's "How To" Book*, PTR Prentice Hall, May 1997, ISBN 0-13-861709-0.

Programming

Chew, Fred, *The Java C++ Cross Reference Handbook* (book/CD-Rom package), PTR Prentice Hall, November 1997, ISBN 0-13-848318-3

Helsel, Robert, *Visual Programming With HO Vee*, 2nd Edition, PTR Prentice Hall, January 1997, ISBN 0-13-631797-9.

Computer Security

Pipkin, Don, *Halting the Hacker*, PTR Prentice Hall, January 1997, ISBN 0-13-243718-X.

Software Development

Grady, Robert B., *Successful Software Process Improvement*, PTR Prentice Hall, May 1997, ISBN 0-13-626623-1.

Mikkelsen, Tim and Pherigo, Suzanne, *Practical Software Configuration Management: The Latenight Developer's Handbook* (book/CD-Rom package), PTR Prentice Hall, May 1997, ISBN 0-13-240854-6.

Object-Oriented

Chaudri, Akmal and Loomis, Mary, *Object Databases in Practice*, PTR Prentice Hall, December 1997, ISBN 0-13-899725-X.

Mellquist, Peter, *SNMP++: An Object-Oriented Approach to Developing Network Management Applications*, PTR Prentice Hall, August 1997, ISBN 0-13-264607-2.

Peripherals

Day, Jerry, *Color Scanning Handbook: Your Guide to Hewlett-Packard Scanjet Color Scanners*, PTR Prentice Hall, July 1997, ISBN 0-13-357211-0.

Test & Measurement

Derickson, Dennis, *Fiber-Optic Test & Measurement*, PTR Prentice Hall, October 1997, ISBN 0-13-534330-5.

Manufacturing

Eisemann, Robert, *Machinery Malfunction Diagnosis and Correction: Analysis and Trouble Shooting for Process Industries*, Prentice Hall PTR, December 1997, ISBN 0-13-240946-1.

Mahoney, R. Michael, *High-Mix Low Volume Manufacturing*, PTR Prentice Hall, March 1997, ISBN 0-13-255688-X.



The Hewlett-Packard Journal

The Hewlett-Packard Journal is published by the Hewlett-Packard Company to recognize technical contributions made by Hewlett-Packard (HP) personnel. While the information found in this publication is believed to be accurate, the Hewlett-Packard Company disclaims all warranties of merchantability and fitness for a particular purpose and all obligations and liabilities for damages, including but not limited to indirect, special, or consequential damages, attorney's and expert's fees, and court costs, arising out of or in connection with this publication.



Subscriptions

The Hewlett-Packard Journal is distributed free of charge to HP research, design, and manufacturing engineering personnel, as well as to qualified non-HP individuals, libraries, and educational institutions.

To receive an **HP employee subscription** send an e-mail message indicating your HP entity, employee number, and mailstop to: ldc_litpro@hp0000.hp.com

Qualified non-HP individuals, libraries, and educational institutions in the **U.S.** can request a subscription by going to our website and following the directions to subscribe.

To request an **International subscription** locate your nearest country representative listed on our website and contact them directly for a subscription. Free subscriptions may not be available in all countries.

Back issues of the Hewlett-Packard Journal can be ordered through our website.



Our Website

Current and recent issues are available online at <http://www.hp.com/hpj/journal.html>



Submissions

Although articles in the Hewlett-Packard Journal are primarily authored by HP employees, articles from non-HP authors dealing with HP-related research or solutions to technical problems made possible by using HP equipment are also considered for publication. Before doing any work on an article, please contact the editor by sending an e-mail message to: hp_journal@hp.com



Copyright

Copyright © 1997 Hewlett-Packard Company. All rights reserved. Permission to copy without fee all or part of this publication is hereby granted provided that 1) the copies are not made, used, displayed, or distributed for commercial advantage; 2) the Hewlett-Packard Company copyright notice and the title of the publication and date appear on the copies; and 3) a notice appears stating that the copying is by permission of the Hewlett-Packard Company.



Inquiries

Please address inquiries, submissions, and requests to:

Editor
Hewlett-Packard Journal, Mail Stop 20BH,
3000 Hanover Street
Palo Alto, CA 94304-1185 U.S.A.

HEWLETT-PACKARD Journal

DECEMBER 1997 • Volume 48, Number 5

Technical Information from the Hewlett-Packard Company