

A Technology Platform for Providing Broadband Communications Services

G. David Bergland
James J. Ferenc
David A. Morano
Walter M. Pittio
Donald D. Shugard
Thomas L. Smith

Electrical and optical components are being created that can be flexibly interconnected to meet the needs of a wide range of emerging business customers. The electrical components are based on a versatile, high-speed electronic-bus architecture that can simultaneously transport voice, data, and video signals, synchronously or asynchronously, at multi-gigabit-per-second data rates. Port cards can *reserve* bus bandwidth for synchronous circuit connections, and can *contend* for asynchronous bandwidth to send packets.

A high-speed optical network interconnects these electronic switches. As data-rate requirements increase to multi-gigabit-per-second transfer rates, optical switches can be used to form distributed broadband networks that span buildings, campuses, and corporations. Inherent data rate and protocol transparency makes optical switches ideal for networking a wide variety of information sources.

Introduction

Our goal has been to define and implement a number of reusable electronic and optical components and interconnections that can economically provide a wide range of broadband services. These building blocks, and their interconnections, constitute a general-purpose, broadband-technology platform that can be configured to meet a wide variety of evolving customer needs.

In the work environment, projected broadband networking applications include multimedia messaging, networked computing, visual communications, real-time multimedia group collaboration, and interactive visualization.¹ The bandwidth requirements for these applications can range from kilobits to megabits to gigabits per second—per endpoint. Future user bandwidth requirements could easily span five orders of magnitude.

Networking standards, such as asynchronous transfer mode (ATM), are, by themselves, not enough to meet customer needs. They must be coupled with the means to deliver true customer satisfaction. Our customers have told us they want a broadband network that is simple to install, operate, and

expand. The network must make it easy to provide reliable services over a broad range of bit rates, and to meet both traffic needs and latency requirements. Customers want a system that removes the effects of this variability, not one that adds to it.

The architecture of a multimedia network is strongly influenced by the transmission medium connecting the desk top to the network. The medium of choice for the last 100 meters to the desk top is still unshielded twisted pair copper wire. The bandwidth available on unshielded twisted pair has been increasing yearly. Rates of 100 to 155 megabits per second are now state of the art, and 622 megabits per second may soon be within reach. The 100-meter distance limit for high-bandwidth, unshielded twisted pair puts a physical limit on the number of desk tops that can be connected to a single wiring closet or hub. In standard office buildings, only about 40 to 80 desks can be reached from a single wiring closet or hub.

Thus, the multimedia network architecture of choice consists of small wiring-closet *access hubs* interconnected with fiber to higher-bandwidth *backbone hubs*. The

Panel 1. Acronyms and Terms Used in This Paper

ATM	— Asynchronous transfer mode
BLAST	— Balanced logic asynchronous synchronous transceiver chip
BASE	— Broadband ATM/STM electrical dual bus
CMOS	— Complementary metal-oxide-semiconductor technology
LAN	— Local area network
Octet	— BASE bus consisting of eight transmit and eight receive data signals, plus clock and control signals
SONET	— Synchronous optical network
ORT	— Optical receiver transmitter board
OSB	— Optical switch board

backbone hubs are interconnected hierarchically, or as a mesh, to give broadband connectivity to a larger number of people, over a larger physical area. See Figure 1 for a view of the architecture. Following these principles, access hubs and backbone hubs are being implemented to form the building blocks of a broadband network architecture that can support multimedia, visual, and computing endpoints that are distributed in buildings and campuses, and in private or corporate networks spread over a wide geographic area.

A Broadband ATM/STM Platform. Our goal was to design a small, wall-mountable hub that would have high enough bandwidth to handle a large variety of user applications, endpoints, links, and networks. The architecture also has to provide a smooth evolutionary path from the customer's present-day environment to the broadband multi-media networks of the future. The architecture must handle today's synchronous transfer mode (STM) interfaces for voice and telephony, and applications such as video, local area network (LAN) and packet networking, as well as the ATM links and applications of the future.

Bus Architecture. The broadband ATM/STM electrical (BASE) bus designed for our hub has a dual-bus architecture, in which all port boards transmit on a transmit bus, and receive from a separate receive bus. The transmit bus is looped back onto the receive bus through a loop-back circuit located at the far end of the bus, as shown in Figure 2. The advantage of using two unidirectional buses is that the clock and data signals propagate together, minimizing the skew between them. This

enables standard bus operations to be performed at much higher speeds. A BASE bus consisting of eight transmit data signals and eight receive data signals, plus the clocks and control signals, is referred to as an *octet*. All BASE-bus signals are differentially driven by balanced logic asynchronous synchronous transceiver (BLAST) chips. Figure 2 shows a four-octet BASE-bus system configuration, with one of the octets (indicated as the 4th octet) shown in detail. Each of the four octets of the BASE bus is a physically independent, pluggable unit that can be installed and replaced in the field.

The port-board interface to the bus is through a BLAST chip bus-interface controller, which is implemented in standard CMOS technology. The bus-interface controller effectively isolates the speed of the bus from the speed of the port board. Therefore, the logic on the port board needs to be only as fast as the application requires.

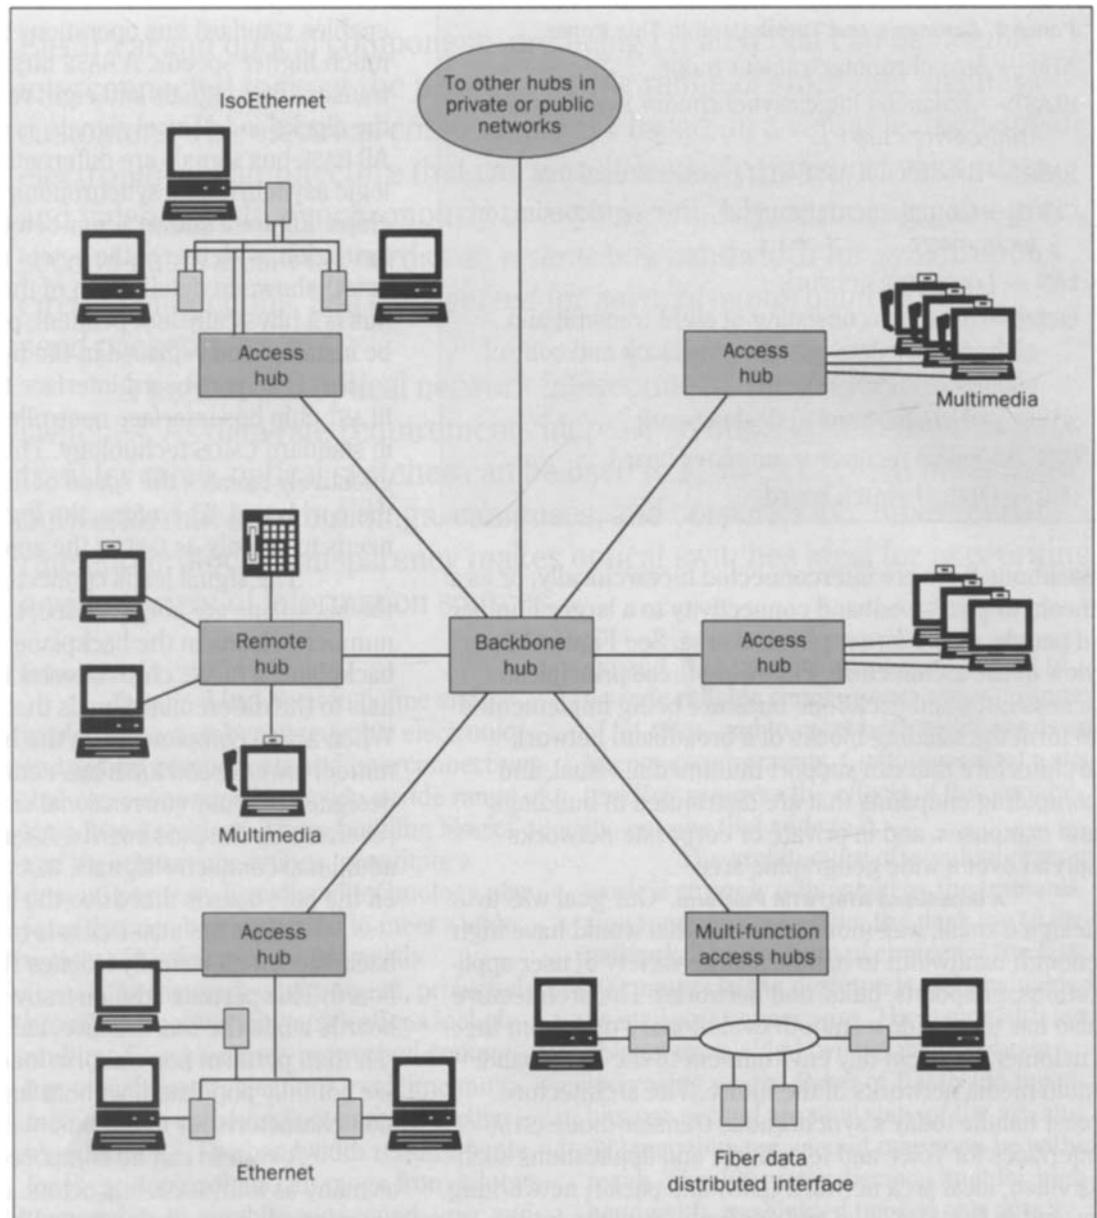
The signal leads connecting the port board to the backplane are single ended, thereby reducing the number of leads in the backplane connector. Once on the backplane, a BLAST chip converts these single-ended signals to the differential signals that appear on the bus. When active components on the backplane are not permitted, lower-speed BASE-bus configurations can be designed that use conventional single-ended drivers and receivers on the port boards. Alternately, at the cost of additional connector signals, BLAST chips can be placed on the port boards that drive the backplane differentially.

When the BLAST chip is mounted directly on the backplane, it electrically isolates the bus from the port board. This permits administrative personnel to plug in boards while the bus is active, called a hot plug-in. They can then perform non-uniform loading of systems that are not fully populated, without impacting the transmission characteristics of the bus.

A system can be composed of as few as one, or as many as four, BASE-bus octets. Four is a physical limit, not a logical limit. Each octet within a carrier is approximately two inches high and, typically, between six and 18 inches long. Multiple carriers can be interconnected using active repeaters. Port boards can access one or more octets in parallel, if the port boards require very high bandwidths, or require redundant buses for reliability.

Universal Port Slots. Each port slot on the bus provides access to the whole bandwidth of the bus. In practice, the bandwidth used may vary from 64 kilobits per second to 622 megabits per second, or higher. Any

Figure 1. Desktops are connected directly, or through local area networks (LANs), to access hubs in wiring closets, using unshielded twisted pair copper wire. Access hubs are connected, using fiber, to one or more backbone hubs. Backbone hubs use fiber to interconnect with each other and with the public network. Remote access hubs can be connected with the rest of the network by using either optical fiber or public wide area network (WAN) facilities.



port card can be plugged into any slot on the bus, with the following exception: Interconnect port cards, which connect hubs, should be placed near the loop-around circuit to take full advantage of the bandwidth available in both the transmit and receive buses. This is discussed further in the section, "Dual Bus Operation."

The BLAST Chip. The BLAST chip on each port slot on the BASE octet bus is a quad transceiver. The chip

performs all necessary electrical and timing functions to interface port card signals to the BASE bus. Eight BLAST chips are necessary per slot (two per bus) to accommodate all of the signals on a BASE-bus octet. The bus side of the BLAST chip uses a proprietary differential-balanced driver that uses reduced voltage swings. These features provide superior noise immunity, while minimizing radiation and power consumption. The BLAST chip also has

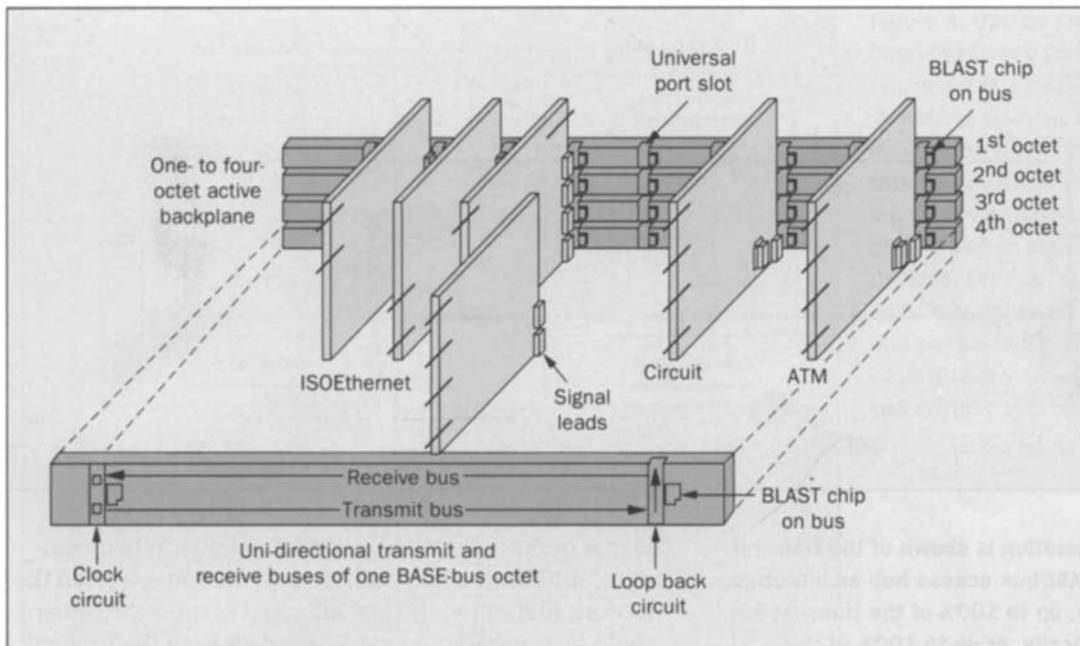


Figure 2. A four-octet BASE-bus system is shown with the balanced logic asynchronous synchronous transceiver (BLAST) chips on the backplane. Each octet transports differentially driven multi-gigabit-per-second signals on both the transmit and receive buses. The BLAST chip electrically isolates the bus from the port boards, so that backplane bus signals remain balanced under all conditions, including hot plug-in.

internal registers and an internal clock-phase adjustment circuit that allows very precise timing control of the BASE-bus signals.

By placing the BLAST chips on the backplane, the BASE-bus signal leads are electrically isolated from the port cards. This totally eliminates variable-loading problems, which plague passive buses. In passive buses, the transmission characteristics of the bus signal leads change with the number and position of plugged-in port cards. This effect requires that the bus be designed around the full range of possible transmission characteristics, and usually requires a reduction in transmission speeds or bus length. Placing the BLAST chips directly on the bus further increases potential bus speeds by reducing the length of the transmission-line taps seen by the signal leads at each port slot. This active backplane scheme, using BLAST chips, creates an excellent transmission environment for transmitting signals that are only a few nanoseconds wide.

The BLAST chip design uses a novel differential signal driver, and an unusual termination design that maintains balanced differential bus signals under all bus conditions. The BLAST chip drivers perform a differential "wired-OR" logic function when two or more BLAST chips are driving the bus simultaneously. This feature eliminates the need for single-ended signals for functions, such as arbitration and control, thereby eliminating

one of the main potential bandwidth bottlenecks of a bus design. Finally, the BLAST chip addresses a problem that conventional driver circuits often cause, propagating unbalanced signals when coming out of the isolation state, called the tri-state. The design of the BLAST-driver circuit maintains balanced operation in both the drive and the isolation state.

In applications where active backplanes cannot be used, the BLAST chip can be placed on the port cards to drive the backplane leads through a connector. This approach takes advantage of many of the BLAST chip-design features. However, it requires a slower bus rate to compensate for the effects of variable loading, hot plug-in, and long electrical taps through the backplane connectors.

Controller Chip and Port Board. The bus-interface controller connects a port board to the BASE bus and allows the board to use any amount of bus bandwidth, up to the full speed of the bus. The controller contains enough expandable buffering on the chip to isolate the bandwidth of the port-board application from the bandwidth of the bus. This isolation feature provides maximum flexibility for customizing each port board to its application and for minimizing cost. Ongoing cost reductions can be done by concentrating solely on minimizing per-line cost on the port board. New high-density and low-density

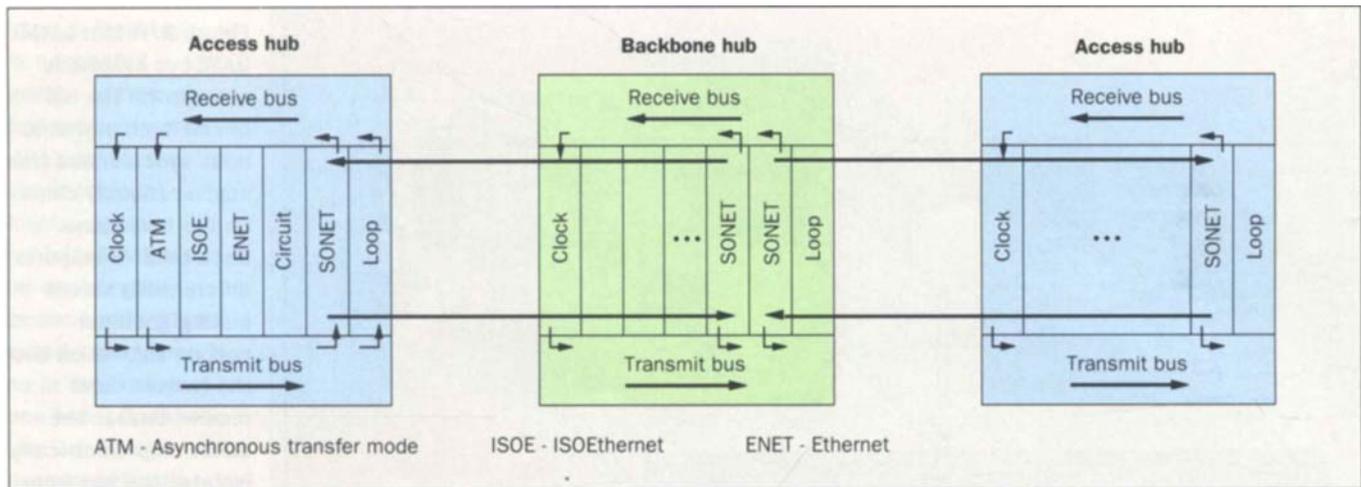


Figure 3. The functional operation is shown of the transmit and receive buses in the BASE-bus access hub architecture. In this dual-shared bus hub, up to 100% of the transmit-bus bandwidth can be routed locally, or up to 100% of this transmit-bus bandwidth can be sent to and received from a backbone hub or another access hub.

port boards can be freely mixed in the system.

The bus interface controller is implemented in standard 3.3-volt CMOS technology. It supports four-level priority bus arbitration for ATM or packet cells, and provides a circuit interface for loading synchronous data into allocated time slots on the bus.

Dual Bus Operation. When applications require a high degree of interconnect bandwidth among backbone hubs, a set of high-speed optical interconnect boards, including standard SONET interfaces, can be placed next to the intelligent loop-back circuit. (See Figure 3.) Transmit-bus traffic can be read by the optical interconnect board, just before it reaches the loop-back board, and can be routed to the backbone hub. Likewise, traffic from the backbone hub can be received by the optical interconnect board and inserted into the head end of the receive bus. The loop-back circuit ignores data that has previously been read by an optical interconnect board. This means that the bandwidth of both the transmit and receive buses can be fully utilized when there is a high degree of inter-hub traffic.

The SONET boards shown in the access hubs, and the backbone hub of Figure 3, are actually identical

boards performing two different functions. When communicating with a hub higher, or at the same level in the network hierarchy, the bus interface of the SONET board would be programmed to receive data from the transmit bus and transmit data onto the receive bus. When communicating with a hub that is lower in the network hierarchy, the bus interface would be programmed to receive data from the receive bus and transmit data onto the transmit bus.

This flexibility permits full utilization of the bus bandwidth for any mix of inter- and intra-hub traffic. For video signals, this could provide a substantial benefit, especially if one assumes that a highly variable amount of video traffic would flow between hubs. Either the inter-hub or the intra-hub traffic can dynamically grow to occupy the full bandwidth of the bus.

Optical Switching Technology

In some applications, high-speed links to backbone hubs may be required to transport all of the data an access hub can generate. Typically, in current data hubs, between 20% and 30% of the traffic in a hub is inter-hub traffic between devices attached to different hubs. In a system with distributed, special-purpose servers, or in systems with a lot of video conferencing, the rate of inter-hub traffic could easily double.

Our current BASE-bus design includes a four-octet hub that could serve as a backbone hub by transporting four times the data that can be generated by a one-octet access hub. Each backbone hub should be able to handle

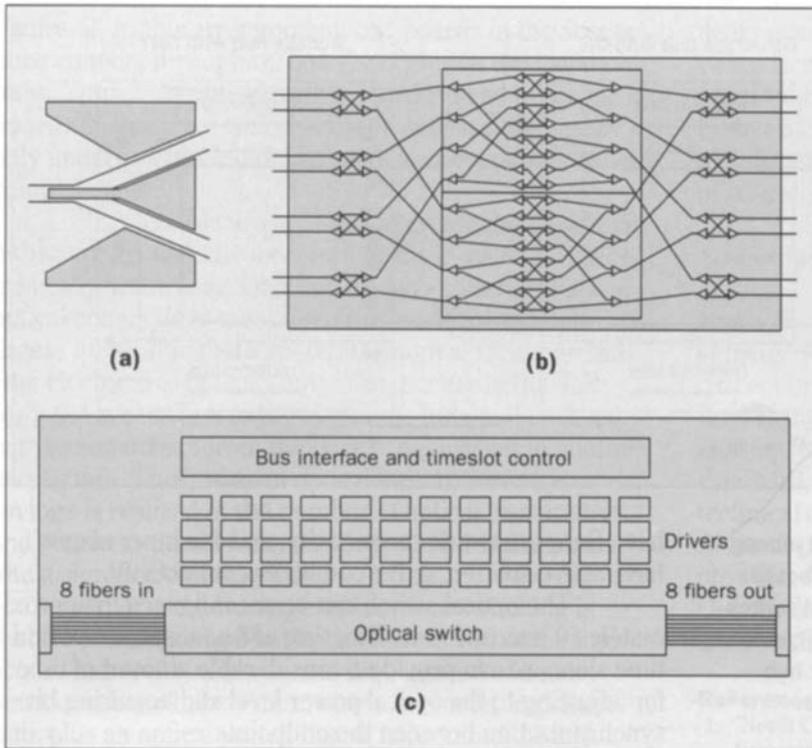


Figure 4. Optical switching technology has been developed to the point at which it can be employed in a backbone hub. (a) The optical switching element can support multi-gigabit bandwidth, and is transparent to both the bit rate and protocol used. (b) The optical switching device can be used as an 8X8 time multiplexed switch supporting gigabit-per-second devices. (c) The high-bandwidth optical switch board could transport both circuit-mode and packet-mode traffic simultaneously. No optical-to-electrical conversions are required at the switch.

at least four single-octet access hubs, and potentially a dozen or more, with various levels of concentration. Most of the backbone-hub bandwidth would be dedicated to performing this backbone-interconnect function.

The backbone function can also be performed by taking advantage of recently developed optical switching technology. This second alternative calls for plugging an optical switch (OS) board into the backbone hub and replacing the SONET boards in the access hubs with optical receiver-transmitter (ORT) boards. A single OS board can switch all of the traffic among at least eight access hubs without using up any of the backbone-hub bus bandwidth. This approach appears to be less expensive and could be extended to even higher bandwidths. We have been experimenting with the guided-wave optical switching technology for many years.^{2,3} Recent advances have made the optical switch devices well suited for product manufacture.^{4,5} Components of this technology that are of special interest are shown in Figure 3.

Guided-wave optical switches have long been admired for being able to switch high-speed digital optical signals in a data- and protocol-transparent manner.

From the point of view of the data being transmitted, the switch behaves just like an optical fiber. The only difference is that light can go in on one port of the switch and be routed to come out on any of several output ports.

In the example in Figure 4, the inputs of up to eight fibers can be routed to any of eight fiber outputs. If the switch is time multiplexed, it can be made non-blocking. The current plan is to use 256 universal time slots, which results in being able to support 2048 simultaneous connections.

In the backbone-hub application, each fiber would be connected to the BASE bus of an access hub using an ORT board. With a 2.5-gigabits-per-second LASER transmitter, more than 2 gigabits per second could be passed between each pair of access hubs homing on an optical backbone hub. Using time division multiplexing, interconnection bandwidth could be dynamically allocated to provide whatever connectivity is needed.

An 8X8 optical switch array and its drivers could be mounted on one port board. It would be connectorized, and installed or replaced in the field. The switch array could be duplicated for reliability, or be installed in

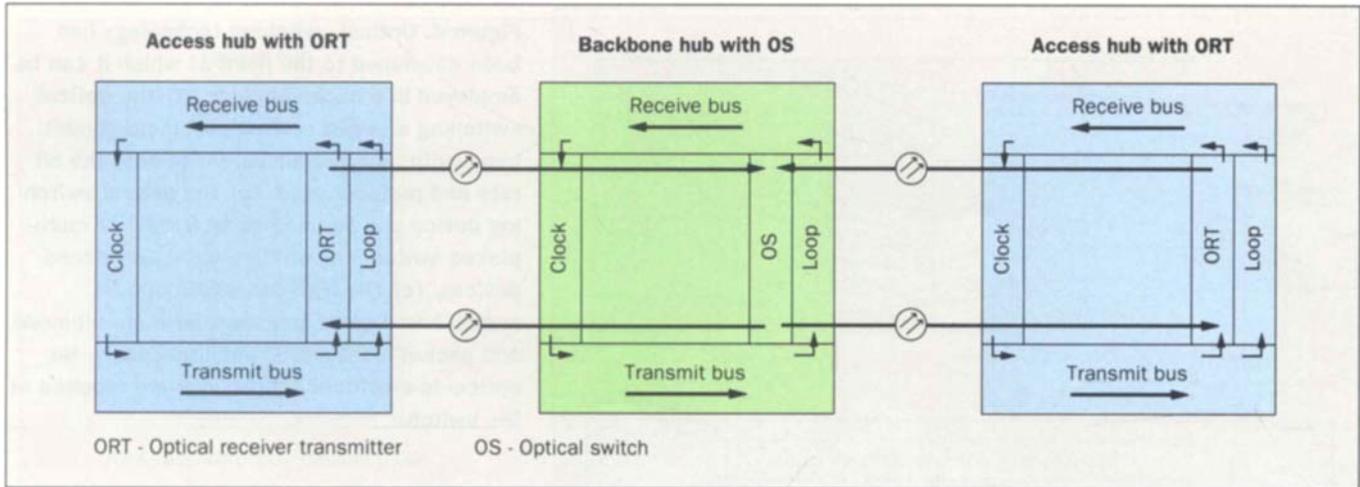


Figure 5. A network could support optical switching among hubs. An optical receiver-transmitter board in each access hub could use the switching capability of the optical switch board in the backbone hub. Thus, instead of requiring a four-octet electrical backbone hub, only a one-octet hub would be required to support inter-access hub traffic.

a network of backbone hubs that provides multiple paths to every access hub. The board will provide 16 gigabits per second of bandwidth (using ORT boards with 2.5-gigabits-per-second LASERS and receivers) and use only 25 watts of power. It should be noted that the optical switch operates independently from the speed of the LASERS attached to it. Even more bandwidth can be provided at the same power by using faster LASERS and receivers on the ORT boards.

The Optical Switch Board (OS)

One of the design issues concerning optical switching involves setting the frame rates and time-slot lengths. The time slots can transport any kind of synchronous or asynchronous signals, because the signals are simply being transported by the optical switch. These universal time slots can carry both ATM cells and STM traffic. The interconnection bandwidth between any endpoints can be dynamically allocated as the load varies. The switch itself is rearrangeably non-blocking, and has negligible cross talk. These features have already been demonstrated by existing prototype switches that use a dilated Benes network.⁶ The current architecture uses a switch that is a dilated hybrid network. The outer stages

have Benes network connectivity, and the inner stages have active-splitter, active-combiner connectivity.

The optical switch can be reconfigured in approximately 10 nanoseconds. By using 488-nanosecond-period time slots, we can provide a considerable amount of time for adjusting to the optical power level and acquiring bit synchronization between the endpoints.

Optical Receiver-Transmitter Board. The ORT board is responsible for selecting inter-hub traffic from the BASE bus, and transmitting both ATM and STM traffic over the 488-nanosecond universal time slots to the proper access hub. The universal time slots are basically carriers for bandwidth between ORT boards.

During these time slots, a start burst, or signal, is transmitted. This allows the receiving ORT board to adjust to the right optical power level, and to bit synchronize its reference clock with the data it is receiving. This bit synchronization can be acquired within a few nanoseconds and, ideally, by one transition of the data. This difficult task is accomplished using a burst-mode receiver chip,⁷ which rapidly adapts to incoming light levels, and by the development of an instantaneous phase-recovery chip⁸ that can synchronize to an incoming data stream with just a single transition of the data. This circuitry, plus the laser transmitters and receivers, can then be put on one port board, which then can be plugged into a BASE-bus hub.

An Optical Interconnection Network. The mesh networks and hierarchical routing networks discussed earlier in this paper can both be implemented using the ORT boards and the OS boards described above (see

Figure 5). In this arrangement, ORT boards in the access hubs connect through OS boards located in the backbone hubs. Thus, if an optical switch board is used in the backbone hub, inter-hub traffic between access hubs must only undergo one electrical-to-optical conversion for the complete path.

If OS boards are not used to create the backbone hubs and, instead, the inter-hub traffic is switched electronically in the backbone hub, then two electrical-to-optical conversions must be performed for the inter-access hub traffic that's routed through a backbone hub. One electrical-to-optical conversion occurs in the link from the access hub to the backbone hub, and another conversion occurs from the backbone hub out to another access hub. Thus, with an optical switch, a two-to-one savings is realized in the number of optical transmitters and receivers used in the network—and this translates into a significant cost saving.

A second advantage with optical switching is that the backbone-hub backplane is relieved of the inter-access hub traffic. This means that, instead of a four-octet backbone hub system, only a one-octet backbone hub, plus an optical switch board, is required. In this case, the backbone hub can simultaneously serve as an access hub. Initially, the high-bandwidth capabilities of the optical switch may not be necessary. When more interconnection bandwidth is needed, an ORT board can be plugged into each hub, and an optical switch board can be plugged into each backbone hub.

Conclusion

Several key components of a broadband technology platform have been described. The multi-gigabits-per-second technology of the BASE bus, and the even higher-bandwidth optical connectivity offered by the optical switch board, could be used together to satisfy a variety of broadband networking needs. Both of these technologies offer the potential for producing a dramatic increase in corporate network data bandwidths. With this increase in bandwidth comes a corresponding simplification in the installation, operation, and maintenance of multimedia networks. High bandwidth can be substituted for network complexity. The idea of providing universal port slots for broadband systems is a powerful concept, supporting the flexibility that users want and need. The variation in future multimedia networking needs will dwarf the variations that have previously occurred in

dealing with voice-only or data-only networks. Flexible components, which can be easily interconnected, are key requirements. We know that customer demands for higher bandwidths will continue to grow. With this flexible platform approach, we can follow the demand instead of restricting it.

Acknowledgments

The broadband technology platform work has come this far due to the tireless efforts and keen insights of many people. Special thanks are due to the rest of the current broadband hardware team members, including Ruel Brown, William L. Aranguren, John Camlet, Steve Durney, Mark Katz, Scott Magnuson, Tim Martin, Bob Pritchard, and George Shevchuk. Special thanks for technical contributions also are due to Rod Alferness, Al Dunlop, Saul Einbinder, Thaddeus Gabara, Ed Murphy, Tim Murphy, Arun Netravali, Yusuke Ota, Kevin Oye, Krishnan Padmanabhan, A. J. Rainal, Bob Swartz, and Rich Thompson.

References

1. "North Carolinians Prepare for Life in the Fast Lane," *Broadband Networking News*, Phillips Business Information, Inc., May 19, 1993.
2. G. D. Bergland, "A Photonic Time-Division Switching Experiment," *OSA Topical Meeting on Photonic Switching Proceedings*, March 1993.
3. R. A. Thompson, J. J. Horenkamp, and G. D. Bergland; Photonic Switching of Universal Time-Slots, *XIII International Switching Symposium Proceedings*, Session C2 paper 4, Stockholm, May 1990.
4. P. Granestrand, B. Lagerstrom, P. Svensson, L. Thylen, B. Stoltz, K. Bergvall, J-E Falk, H. Olofsson, "Integrated Optics 4 X 4 Switch Matrix With Digital Optical Switches," *Electronics Letters*, Vol. 26 1990, pp. 4-5. (ELLEMTTEL and Ericsson Telecom)
5. H. Okayama, T. Ushikubo and M. Kawahara, "Reduction of the voltage-length product for a Y-branch digital optical switch," *OSA Topical Meeting on Photonic Switching Proceedings*, February 1991. (Oki Electric Industry Co., Ltd., Tokyo.)
6. K. Padmanabhan and A. N. Netravali, "Dilated Networks for Photonic Switching," *OSA Topical Meeting on Photonic Switching Proceedings*, March 1987 and *IEEE Transactions on Communications*, December 1987.
7. Y. Ota, R. G. Swartz, V. D. Archer III, "DC - One Gigabit per second Burst Mode Compatible Receiver for Optical Bus Applications," *Journal of Lightwave Technology*, Vol. 10, No. 2, pp. 244-249 (1992).
8. M. Banu and A. Dunlop, "A 660 Mb/s CMOS Clock Recovery Circuit with Instantaneous Locking for NRZ Data and Burst-Mode Transmission," 1993 IEEE International Solid-State Circuits Conference, June 1993.

(Manuscript approved September 1993)

Glenn David Bergland is department head of the Switching Systems Research Department at AT&T Bell Laboratories in

Holmdel, New Jersey. He joined AT&T Bell Laboratories in 1966 and works on broadband optical and electronic switching systems, multimedia services, and software project management. He has a B.S.E.E., M.S.E.E., and Ph.D. in electrical engineering from Iowa State University, Ames, Iowa.

James J. Ferenc is a supervisor in the Broadband Solutions Development Department at AT&T Bell Laboratories in Denver, Colorado. Mr. Ferenc joined AT&T in 1979, and is currently responsible for hardware development of broadband multimedia business communication systems. He received a B.S.E.E. degree from the University of Michigan in Ann Arbor, and an M.S.E.E. degree from Stanford University, in Palo Alto, California.

David A. Morano is a member of technical staff in the Switching Systems Research Department at AT&T Bell Laboratories, Holmdel, New Jersey. Mr. Morano joined AT&T in 1982 and works on high-speed switching systems. He has a B.S.E.E. and an M.S.E.E. from Northeastern University, Boston, Massachusetts.

Walter M. Pitlo is a member of technical staff in the Broadband Switching Technology Group at AT&T Bell Laboratories in Holmdel, New Jersey. He is working on broadband guided-

wave optical switching and high-speed electronic bus technology. Mr. Pitlo has a B.S.E.E. degree and an M.S. degree in computer science from Stevens Institute of Technology in Hoboken, New Jersey. He joined American Bell in 1983.

Donald D. Shugard is a member of technical staff in the Switching System Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. He joined the company in 1984. Mr. Shugard has a B.S.E.E. from Rutgers College of Engineering, New Jersey, and an M.S. in computer science from Columbia University, New York, New York.

Thomas L. Smith is a technical manager in the Wireless Communications Development Department at AT&T Bell Laboratories in Middletown, New Jersey. He joined AT&T in 1968 and was involved in research and development of high performance switching fabrics for Global Business Communications Products prior to his current assignment in wireless communications. Mr. Smith has a B.S.E.E. and an M.S.E.E. from the Massachusetts Institute of Technology in Cambridge.
