# FAST PACKET TECHNOLOGY FOR FUTURE SWITCHES

**John J. Degan, Gottfried W. R. Luderer, and Avinash K. Vaidya**

New network services other than the traditional voice communications between people will be increasingly bursty in bandwidth needs and will require distributed processing and control. The new fast packet technology for transmission and switching is ideally matched to these future network service needs. In this paper, we give a high-level sketch of the evolution of fast packet switch technology starting with the first packet switches and ending with some architectures that offer promise as potential switching vehicles for Broadband ISDN. We shall be using examples of AT&T experimental switch architectures to illustrate the evolution and technical capabilities.

***John J. Degan*** *since April 1989 has been executive director of AT&T Bell Laboratories' Interexchange Switching Systems Development division. Prior to this new assignment, beginning in 1980, he was executive director of the Network Services Development and Technology division of AT&T Bell Laboratories at Columbus, Ohio and Naperville, Illinois, and was responsible for managing all aspects of research and development into AT&T's network products and services. Mr. Degan holds a B.S. in electrical engineering from Northeastern University and has completed the AT&T Bell Laboratories Graduate Study Program.* ***Gottfried W. R. Luderer*** *is head of the Advanced Switching Networks Department at AT&T Bell Laboratories' Indian Hill Park Laboratory in Naperville, Illinois, where he is responsible for ap-*
(continued on page 50)



## Introduction

The evolution of the telecommunication network will be shaped by the interaction of new customer needs and emerging new technologies. In this paper, we will outline one emerging new technology that is well suited to satisfy the expected needs of the future network customers: fast packet technology.

Today's network has been designed to optimally handle a vast amount of voice calls—a traffic load that can be considered rather homogeneous. Contrast this with the traffic mix expected for the future telecommunication network: an unpredictable mixture of voice, data, image, and even video communications. Whereas the evolving network in the near term is well served with the fundamental building block of a 64 kilobits per second (Kb/s) digital bearer channels and out-of-band signaling, a more flexible paradigm of communication will be needed in the future. There are three main characteristics that pose new requirements for the future network.

- *Bursty communications* will require dynamic bandwidth allocation, preferably with instantaneous response;
- *Human-to-machine communications* will increasingly replace human-to-human communications, as more and more information services

**Panel 1.Terms and Acronyms in This Paper**

| | |
|---|---|
| AFIPS | American Federation of Information Processing Societies |
| ANSI | American National Standards Institute |
| ATM | asynchronous transfer mode |
| CMOS | complementary metal-oxide semiconductor |
| FDDI | fiber distributed data interface (ANSI standard) |
| GaAS | gallium arsenide |
| Globecom | IEEE Global Telecommunications Conference |
| ICCC | International Conference on Computer Communication |
| IEEE | Institute of Electrical and Electronics Engineers |
| ISDN | Integrated Services Digital Network |
| ISS | International Switching Symposium |
| MARS | multistage alternate routing switch |
| NTSC | National Television Standard Code |
| VLSI | very large scale integration |
| TDM | time-division multiplexing |

appear in the network. This phenomenon will accelerate the trend towards distributed control.

- *Machine-to-machine* will play a more important role as workstations increasingly become "natural" extensions of the communicating persons. This development will create additional challenging performance needs.

### The New Service Paradigm

Given these trends, we see the emergence of a new communication paradigm that could replace the voice call: this paradigm is the movement and management of a unit of digital information ("a set of bits") under the constraints of time, space, accuracy, and variability.

*Time constraints* imply that delivery has to occur at a specified time or with a bounded delay. For interactive voice, this may be on the order of milliseconds; for bulk data, it may be minutes. As with messaging services, there will also be timed delivery. Furthermore, the tolerable time delay before the network accepts data may itself be a constraint; consider the case of connection setup time going to zero, e.g. for connectionless service or on permanent virtual circuits.

*Space constraints* should be more remarkable by their absence, as we hope to see increased flexibility in supporting unconstrained connectivity (conferences), and an increased trend toward geographic transparency, i.e., finding targets on a logical or name basis rather than by location.

*Accuracy constraints* will differ for various types of communications needing varying degrees of accuracy. There are data for which error-free delivery must be guaranteed, even if such delivery requires temporarily increased delay. There are other communications, such as interactive voice, for which timely delivery is paramount and isolated "glitches" can be tolerated. Or, we may have image transmission, where an errored picture element causes a blemish that may be tolerated and can be corrected later and out of sequence.

*Variability constraints* may restrict how smooth or "bursty" the accepted traffic load can be, that is, how much the peak rate may differ from the average rate.

These aspects of the future telecommunication paradigm suggest that the network must deal with digital information units that require somewhat customized treatment for movement and management. Packet technology has been the traditional way to deal with such requirements in the past; we have seen its successful use in data communication and the signaling network. It was only natural to try to generalize the use of packet switching across other media, and build an integrated packet transport network; one such proposal was made by Turner and Wyatt.[1] The last few years have seen significant advances in packet technology; in the following pages, we will sketch some of the results in the field of fast packet switching.

**Fast Packet Technology Definitions**

Packet-switching was invented in 1964.[2] We take the definition of packet technology broadly. It includes transmission and switching methods that rely on information units labeled with self-identifying instructions that may pertain to the units' source, destination, or intended treatment.

Advances in microelectronics such as very large scale integration (VLSI), photonics, and software technologies have all contributed toward moving packet technology from the realm of "data-only" (and signaling) to include the domains of voice and image. The term "fast packet" switching or technology can have several meanings.

- The first one implies *fast data rates.* While conventional packet switching has been operating at data rates below 1.5 megabits per second (Mb/s), typically in the 4.8 to 65 Kb/s range, new fast packet systems run at least at 1.5 Mb/s, and experimental systems are reported with rates in the 100 to 200 Mb/s range.
- The second meaning of "fast" refers to *fast response* or low delay. Although conventional packet systems incur delays in the hundreds of milliseconds, fast packet systems have been built that add only a few milliseconds to the propagation delay, excluding, of course, the packetization time.
- Finally, "fast" can refer to *fast connection.* Even though the establishment of a packet connection—the virtual circuit—is not inherently faster for the new fast packet systems, there is the opportunity and tendency to do away with the call setup entirely, by using so-called *connectionless* service or permanent virtual circuits.

A distinction has been made to categorize various fast packet switching ranges such as "wideband packet switching,"—i.e., the range from 1.5 to less than 45 Mb/s—and "broadband packet switching"—the range from 45 Mb/s to 150 Mb/s.

Fast packet technology applies to both transmission and switching systems. In *transmission,* it replaces *time-division multiplexing* (TDM) with a label multiplexing concept—sometimes called *demand multiplexing.* In TDM, the position in a frame implicitly determines the association within the information stream. However, in demand multiplexing the association with an information stream is made explicit through the header or label on the data unit. The term *asynchronous time-division multiplexing* has been used for this multiplexing technique[3] because information units appear not synchronously, but asynchronously, on demand. For the upcoming Broadband ISDN standard (see below), the term *Asynchronous Transfer Mode* (ATM) has been introduced, where "asynchronous" has the same connotation.[4]

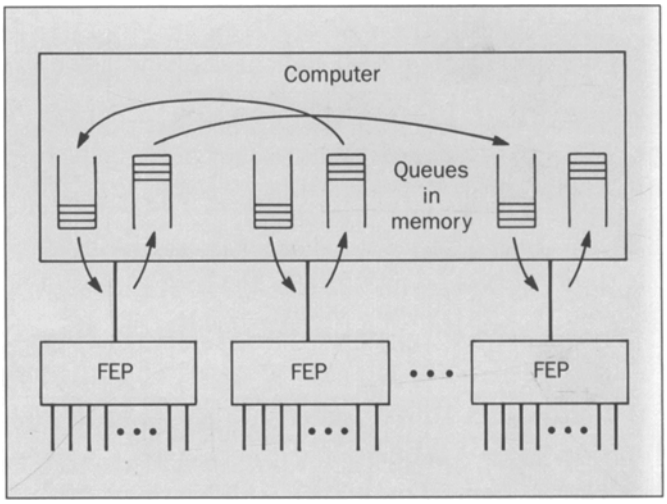


**Figure 1. First generation packet switches.**

Fast packet technology has two fundamental properties that make it an ideal match for the future network. One is its inherent ability to *allocate bandwidth* dynamically and instantaneously for both transport and distribution. The other is the provision of a handle in the form of the header that enables *distribution* of a wide range of processing and control functions.

The switching function in *circuit switches* must be performed continuously for the duration of the connec-

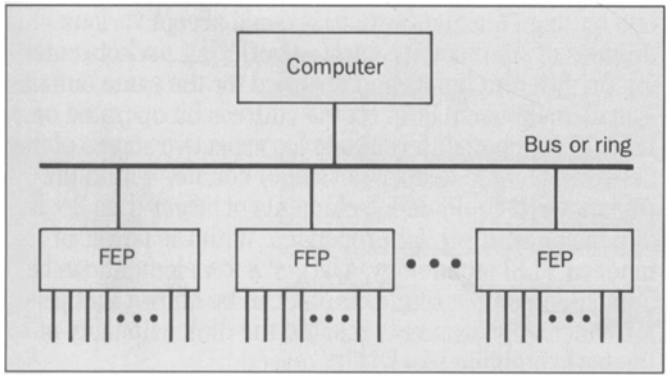


**Figure 2. Second generation packet switch.**

tion. In the packet switching *virtual circuit*, however, the switching function need be performed only when information is present. Thus, there is an inherent throughput gain due to better utilization. This aspect of packet switching gave rise to the notion of *statistical switching*.[5]

The second aspect of packet switching concerns the support of distributed processing. As we will see, modern fast packet switch architectures show a high degree of distributed control because the switches use a *self-routing* that allows significant parallelism. Rather than have a central control subsystem, modern packet switches inspect packet headers individually and perform a local switching function by acting upon the information in the header. This is illustrated in the first example of a fast packet switch, which employs a Banyan network. Before we turn to this example, let us first briefly examine the recent evolution of packet switch architectures.

### Packet-Switch Evolution

The first generation of packet switches, such as those used in the ARPANET,[6] were simply computers in which the processor performed the switching function by moving packets (or pointers to them) from an input queue to an output queue. Soon certain lower level protocol functions such as link error control were moved to front-end processors (FEP), as shown in Figure 1.

In these switches, the aggregate packet throughput was limited by processor speed, typically to a few hundred packets per second. This is a reasonable match for transmission link speeds in the 1-20 Kb/s range.

The second generation of packet switches removed the computer from the bottleneck of switching and introduced a shared medium to interconnect the front-end processors, as shown in Figure 2. In second-generation switches, the computer is used only during the virtual circuit setup by updating entries in the associated front-end processors. The transport of packets between front-end processors goes over the shared medium, typically a bus or ring. Examples of such second-generation packet switches are the AT&T 2STP signal network packet switch[7] or the AT&T 1PSS Release 4 packet switch used for Accunet® public packet service.[8] Both use a duplicated ring architecture as a shared medium. Another example of such a packet switch is the AT&T Datakit® virtual circuit switch, which uses a folded bus architecture.[9]

The second-generation packet switches actually use the principle of time-division multiplexing on the shared medium. Only one packet is in transit at a given time. Naturally, the throughput limitation of these switches is the shared medium's bandwidth. This can be increased by using *parallel* shared media, either a multiline bus or a multiple bus architecture.[10]

The next generation of packet switches removed the shared medium bottleneck by using interconnection networks, in essence introducing the principle of *space-division switching*. In circuit switching, space-division switches predated time-division switches. The crossbar switch is the preeminent example of a space-division switch. This path mostly has been avoided in packet switching. With the expense of electronic crosspoints and the need of $n^2$ crosspoints for a switch with $n$ terminations, there has been a successful search for ways to reduce the number of crosspoints. Work in this area predates packet switching, and substantial contributions

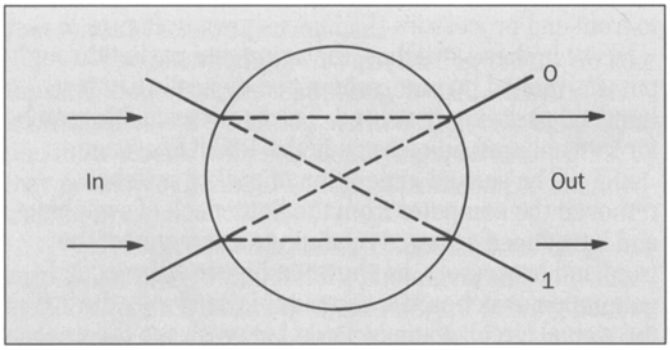


**Figure 3. Basic 2 × 2 packet switch element.**

have come from circuit switch network research and computer architecture research.[11]

In the following sections, we shall offer an overview of recent contributions in the area of third-generation packet switches.

Having seen the application of the time-division and space-division principles to packet switch architectures, it seems logical to ask for examples of the use of frequency (or wavelength) division to the field. Indeed, there is a recent paper on using frequency division as a switching principle. This so-called *multihop switch*[12] can thus be called a fourth-generation packet switch.

#### Banyan Interconnection Networks

The basic building block of third generation switches has been an $n \times n$ self-contained switching element that can be repeated and stitched together in many ways. In its simplest form, it is a two-by-two element as shown in Figure 3. The element has two inputs on the left and two outputs on the right. Its function is to switch packets entering on the left to one or both outputs on the right. The packet header is inspected to make a selection. In its simplest form, a single bit determines one output or the other. The switch element also performs a sorting function by moving the lower address packets to the top.

Many variations of this element are possible. It could support broadcasting by operating on more than one bit (e.g., one per output). It could accept various degrees of simultaneity, such as buffering packets entering on different inputs and destined for the same output. The element could strip off the address bit operated on. It could act upon different bits in respective stages of the network. It could keep a bit (stage) counter within the packet. Or, it could define elements of larger than $2 \times 2$ dimension with similar properties. With the power of modern VLSI technology, a larger $n \times n$ element can be implemented on a single chip. It can be shown that performance improvements result if the dimensionality of the basic building block is increased.

The simplest network one can construct out of these basic elements is the Banyan network, i.e., it has the fewest number of elements.[13] An example of an $8 \times 8$ Banyan is shown in Figure 4. The self-routing property of the network is illustrated by following the progress of two sample packets with destination address 110 and 001, respectively, in their headers. Regardless of the input by which these packets enter the network on the left, the switch elements that work on subsequent bits to select the 0 or 1 outputs will always route the packets to the respective outputs—010 or 111—on the right side of the figure. If the network is built from $2 \times 2$ elements, an $n$-stage network will support $2^n$ terminations. For example, a ten stage network with 512 elements in each of the ten columns would need a total of 5,120 elements to support 1,024 terminations. AT&T used a $16 \times 16$ Banyan network in a field experiment in 1985 and 1986.[14]

The simplicity of the Banyan is attractive but has some disadvantages. To avoid collisions of packets having to traverse the same element, one can introduce buffers, operate the network at higher speeds than the inputs, or both. Or collisions can be avoided totally by sorting the inputs so their paths will not cross, as illustrated in the next example of a new packet switch.

#### The Starlite Packet Switch

The basic switching elements also can be used to perform a variety of other functions that involve reordering based on address relations. For example,

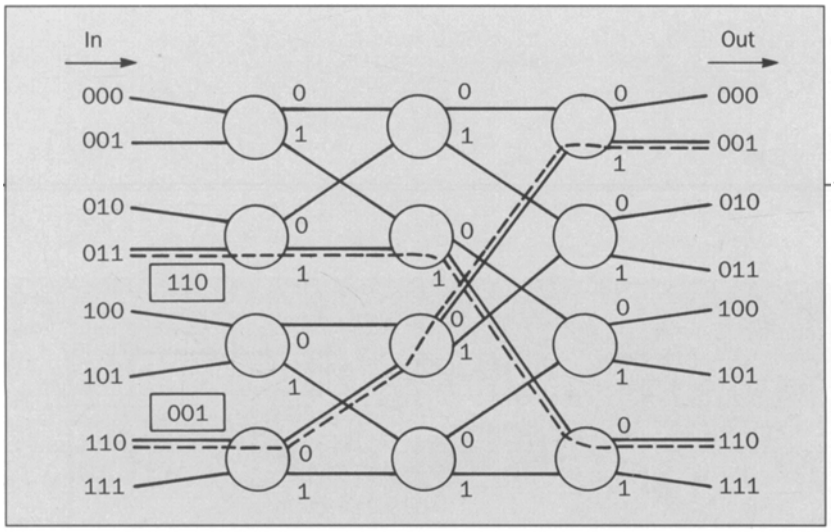


**Figure 4. Banyan interconnection network.**

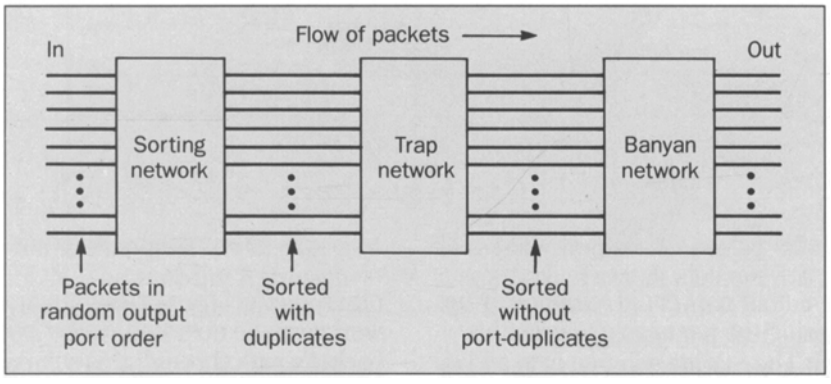


**Figure 5. Basic Starlite or Batcher Banyan network.**



Batcher[15] proposed to build networks that sort the incoming packets in ascending (or descending) address order. For example, if packets with destination addresses 6, 5, 0, 7, 3 enter an $8 \times 8$ Batcher sorting network on any set of 5 of the 8 inputs, the output would be ordered 0, 3, 5, 6, 7, $x, x, x$, where $x$ represents no outputs for the highest three ports. To complete a switching function, one must include *dummy* packets for the unaddressed outputs. Batcher proposed a scheme that merges guide packets in a bidirectional merge network with the output of the sorter. Also, if there are multiple packets for the same output port, his arrangement would let only one packet pass and would drop the others.

A simpler approach to use sorting for switching has been proposed by Huang and Knauer. For their switch known as Starlite,[16] they use the idea of combining a Batcher sorting network with a Banyan switching or "expansion" network. If the inputs to the Banyan are presorted, collisions within it are avoided. Thus, the Banyan network becomes unblocking. The exception is *output blocking*, in which multiple packets aim for the same output port. In this instance, the duplicates are intercepted in a trap network, illustrated by Figure 5. Seen differently, the sorted output from the Batcher network (0, 3, 5, 6, 7, $x, x, x,$ in our example) is expanded to skip over unaddressed slots (0, $x$, 3, $x$, $x$, 5, 6, 7). The network elements interact so unsorted packets entering on the left leave the leftmost network sorted. They then enter a trapping stage in the middle that removes duplicates. The switching function is completed when the sorted but port-unique packets finally are expanded by the Banyan network on the right. Starlite switching networks of this type also are known as *Batcher-Banyan networks*. Such networks can be considered the ancestor of several space division packet switches.

The major drawback of this architecture is the loss of trapped packets if multiple packets are headed for the same output port. If all but one of the multiple pack-

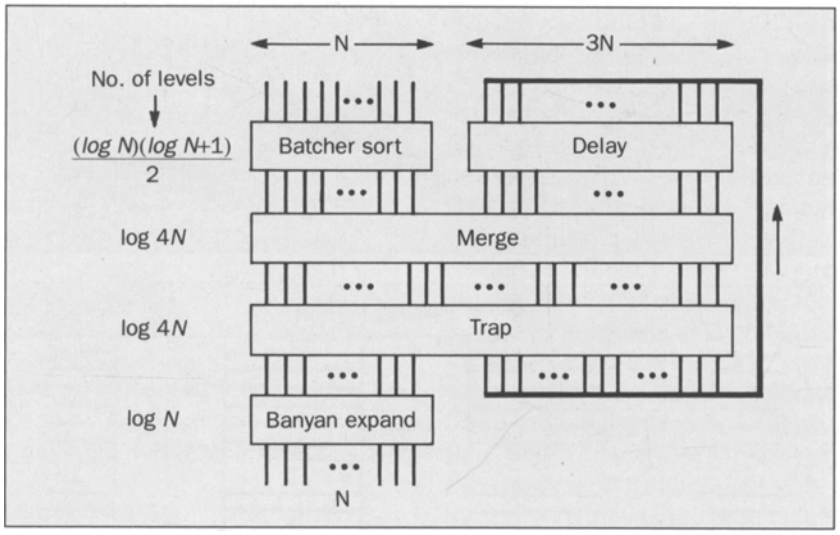


**Figure 6. Starlite switching network with recirculation.**



ets headed for the same output port could temporarily be held back for reentry during the next packet cycle, this problem could be solved. The Starlite solution is to add a trap network with recirculation to the sorting and Banyan networks. Duplicate address packets are caught and recirculated, i.e., they are reintroduced into the next sorting cycles. The number of recirculations needed in such a switch depends on the traffic mix. The complete Starlite switch with Batcher sorting, Banyan expansion, and trap stages also requires delay and merge stages with dimensionality for an $N \times N$ switch shown in Figure 6. Huang, Knauer, and O'Neill have shown that such networks of arbitrary size can be constructed from just three custom chips, two for the Batcher sorting network and one for the Banyan network.[17] Using fewer chip codes is possible if engineering considerations warrant this to be beneficial. Huang and Knauer also have proposed variations of this switch to handle broadcasting by introducing additional copy-network stages.

A sorting network is *packet-synchronous* if it requires all packets to enter the network at the same time. In these networks packets must be compared with each other to effect the sorting function. However, Banyan networks can be operated *packet-asynchronously* because the packet's path through this self-routing network is—at least as a first order—unaffected by the presence of other packets.

The placement and required size of storage are important design considerations for packet switch architectures, since these are important cost components. Storage is needed for many purposes, including input latches for synchronization, input and output latches for speed matching, disassembly/assembly buffers for segmentation, registers for processing (addresses, checksums, etc.), and buffers for queuing. Basically, storage can be at the input, output, or middle of the interconnection network, or can exist in several of these places. Packet-synchronous networks require latch storage at the input. The buffered Banyan approach has buffers in each element. The Starlite approach has a pool of shared buffers in the recirculation network. Queuing and buffer sharing in space-division packet switches have been extensively studied.[18 19 20]

Packet-synchronous architectures of the Starlite

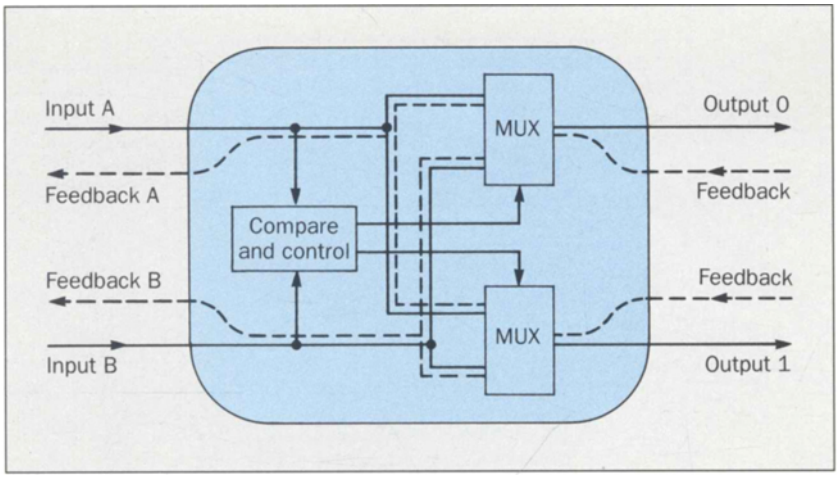


**Figure 7. Switch element with feedback links.**

type show great promise for handling the traffic expected in the future broadband ISDN. The evolving ATM standard is converging on a fixed length cell size of 64 octets payload with 5 octet header.

Recent reports show that packet switches operating at input rates of about 150 Mb/s per access link have been realized in various places in the laboratory,[21 22] and that there is increased interest in high-speed packet switching with broadcast capability.[23 24]

#### Batcher Banyan Network With Feedback

Higher circuit speeds usually accompany reduced gate counts per chip as evidenced in semiconductor technologies such as Gallium Arsenide (GaAs), or in bipolar technology. Thus, if one wants to increase the data rate, it is also desirable to reduce the storage needs of the switch architecture. If the input-side latch is needed for packet synchronization—as it presumably is for fast multiswitch networks that cannot be operated in total synchronism—it is possible to use various kinds of retransmission schemes instead of the recirculation buffers. One such scheme uses probing packets to find out about potential contention and then send one packet per output port only.[25] Vaidya has proposed another scheme: to modify the Batcher sorting network by building it from a new type of switch element that includes feedback lines, as shown in Figure 7.

An example of a sorting network constructed from such elements is shown in Figure 8. In this packet-synchronous architecture, packets entering from the left side may collide somewhere within the network. A contention resolution circuit in each $2 \times 2$ switch element reports back which input has lost the contention. This signal is fed back to the input buffer stage, which stops the transmission.

For such an architecture, M. A. Pashan has implemented an experimental $2 \times 2$ switch element using GaAs technology. Laboratory measurements have shown switching of packets at a rate of 250 Mb/s. The actual transfer rate through the chip has been measured at 1 gigabits per second (Gb/s), though no packet load could be generated at this rate. These results hold promise that fixed-length packet switching could be practical in the 500 Mb/s range within a few years.

#### Some Challenges

Fast Packet Switches that operate at such high data rates as hundreds of megabits per second and that

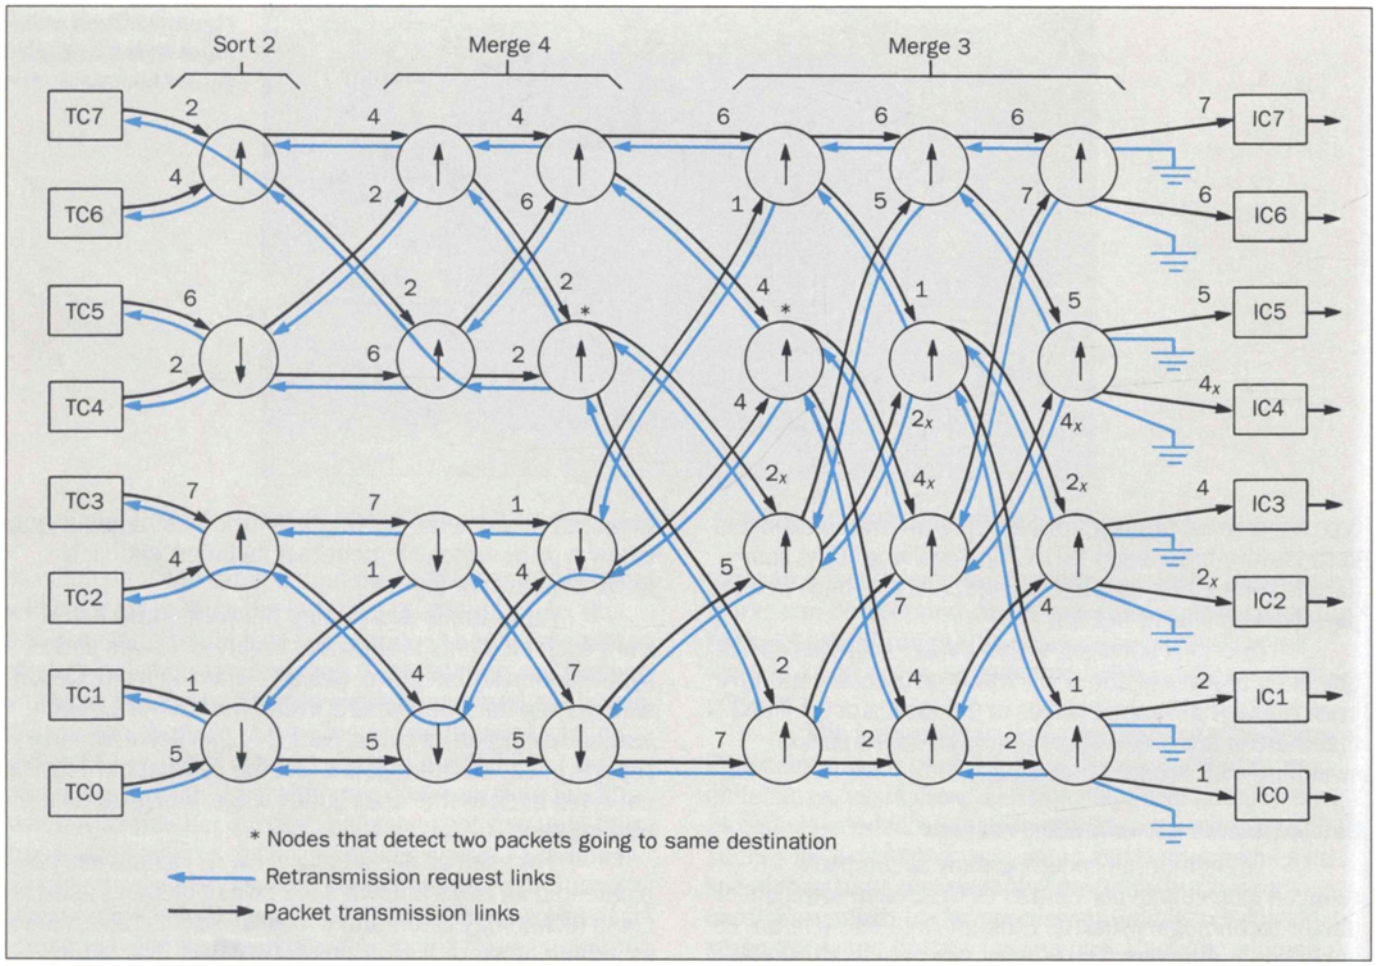


**Figure 8. Modified sorting network.**

can handle a large number of users pose a number of engineering challenges.

**Operating Rate.** One can operate the interconnect network at a lower clock rate by using parallelism. Thus, only the access links have to be operated at the high access speed. For variable packet length approaches, one must recognize the delimiting flag. This need not be the traditional flag pattern, but use of an easier scheme such as 4 out of 5 bit coding has been proposed.[26]

**Header Processing.** Many architectures require the destination address in the header be translated into some other address, e.g., for routing through the switch. This address substitution can become a limiting factor on how many packets per second can be processed on an access

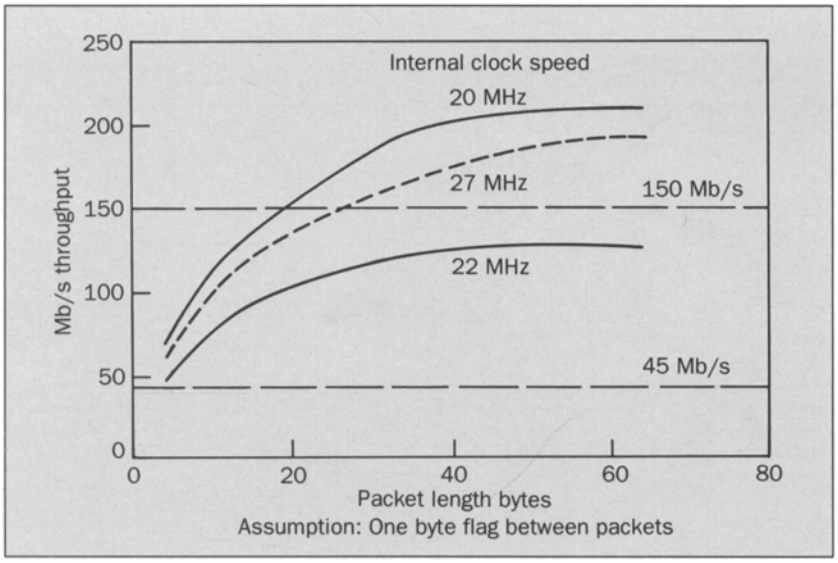


**Figure 9. Throughput of a packet translation unit.**



link. Figure 9 shows the simulation results from Pashan's model design. Throughput increases with packet length, because the per-packet-processing is a constant requirement. The three curves in the figure illustrate three different clock rates representative of different generations of VLSI technology.

**Packaging.** Short and constant length interconnect paths are required, especially for packet-synchronous architectures. Wilson of Bell Communications Research has demonstrated an interesting packaging concept where subsequent stages of the switch boards are rotated by 90 degrees in space to create a uniform "pipeline" structure.[27]

### Multistage Alternate Routing Switch

Another approach towards building large switches is the use of a *packet-asynchronous* design. The experimental multistage alternate routing switch (MARS) invented and designed by Kari Teraslinna of AT&T Bell Laboratories is shown in Figure 10. The dotted lines show four different paths between inputs 5 through 8 on the left and outputs 1 through 4 on the right. Each board implements a non-blocking switch element of $4 \times 4$ that contains its own buffers. It accepts and switches packets in an asynchronous fashion. That is, variable length packets entering the board are immediately put on their way, since no sorting is required. These boards are interconnected as a Beneš network with 4 alternate routes between each input and output.[28] The switch operates internally in parallel at an aggregate rate of 160 Mb/s, and can accept input at 150 Mb/s on each port. A multicasting capability is provided. Donald W. Peterson, in as-yet unpublished research, has demonstrated packetized broadcast video switching on this switch by encoding each National Television Standard Code (NTSC) television image line into a 668 byte packet by sampling three times the color carrier frequencies and omitting the synch interval. This yields an aggregate rate of 84.168 Mb/s in uncompressed form. A $16 \times 16$ prototype of this switch was completed in 1986 and has been used for experimentation with various aspects of fast packet technology.

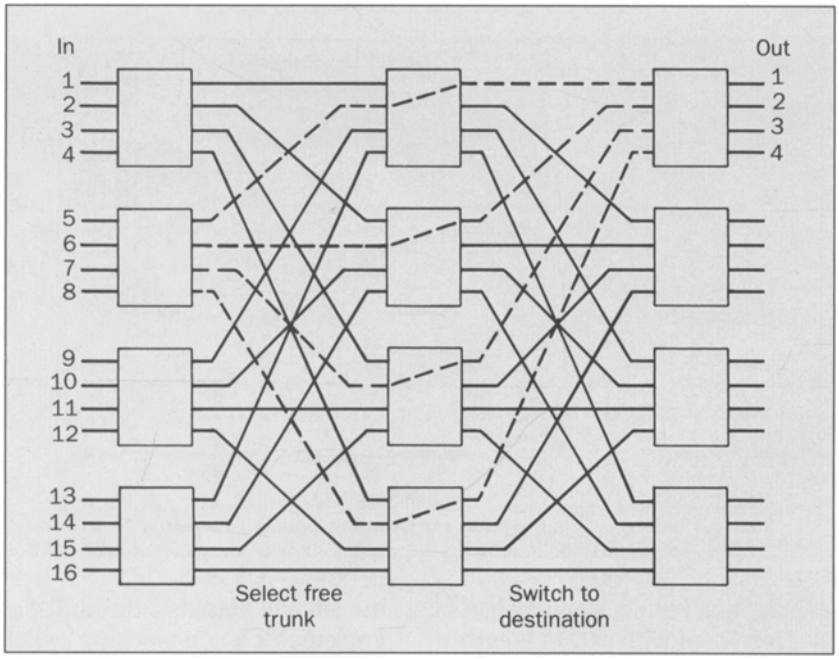


**Figure 10. Multistage alternate routing switch (MARS).**



The MARS architecture uses multiple paths to reduce the likelihood of collisions and increase the throughput of the fabric. A similar approach is used in "load-sharing Banyan" networks proposed by Chin-Tau Lea,[29] which uses extra intrastage paths rather than multiple interstage paths in MARS.

#### The Knock-Out Switch

As we have seen, several approaches exist to avoid collisions—and resulting delays or packet loss—in switch architectures. However, *output blocking* is a situation that *no* architecture can really prohibit, since it is intrinsic to packet switching. While the called party can be busy in a circuit switch, in a packet switch it is usually considered legitimate to have multiple simultaneous packets for the same destination. The network is assumed to provide some storage capability. To our knowledge, this problem has been first addressed specifically in an unpublished switch architecture by Nian-Chyi Huang, called a Contention-Free Banyan Network. To explain the motivation for designing this switch, we start by stating a basic network engineering question: *How many simultaneous inputs from the switch is a switch termination willing to accept?* Huang has found that the number eight (8) goes a long way for many different realistic network and traffic assumptions they have modeled. With 8 outputs per termination, packet loss probability can be kept below $10^{-6}$. Huang proposed to interleave the appropriate number of Banyan networks. An example of an expanded Banyan that assumes each output can receive up to four simultaneous packets is shown in Figure 11.

One must make provisions to resolve the contention for the rare case when more than the maximum allowed number of packets per destination occur. This problem has been explicitly addressed in a new switch architecture by Yeh, Hluchyj, and Acampora, called the *knock-out switch*.[30] The premise is that an individual subscriber can receive at most eight packets. Should there

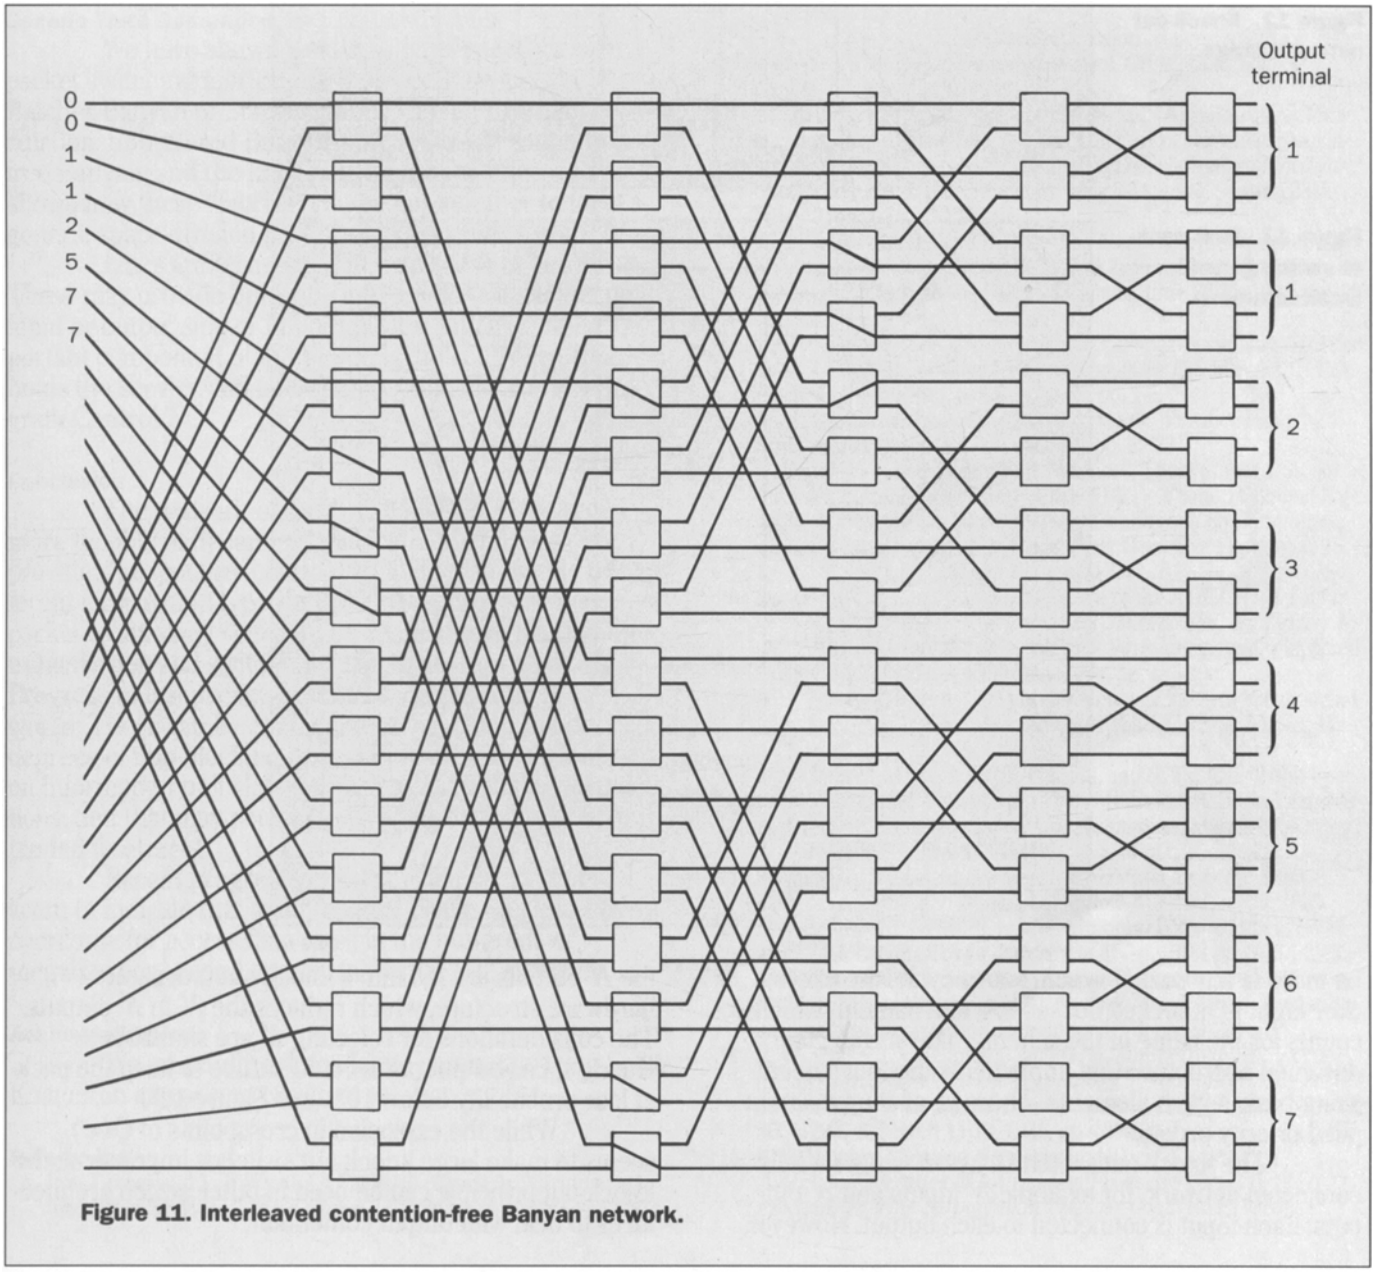


**Figure 11. Interleaved contention-free Banyan network.**

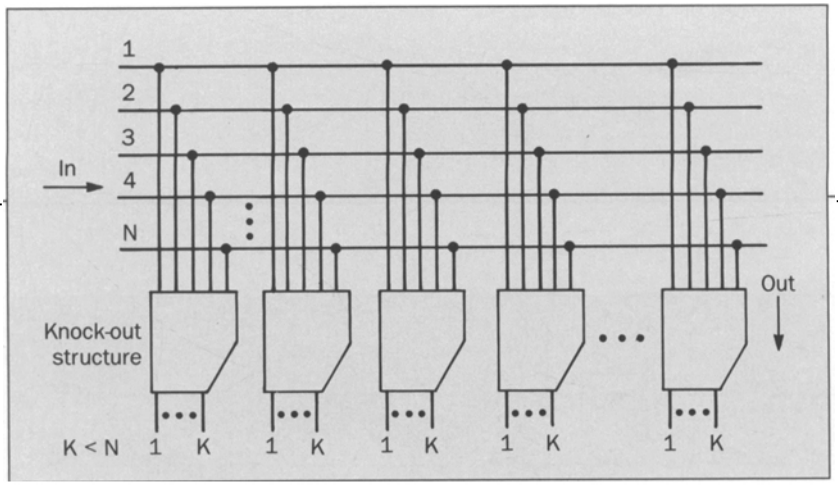


**Figure 12. Knock-out network stage.**

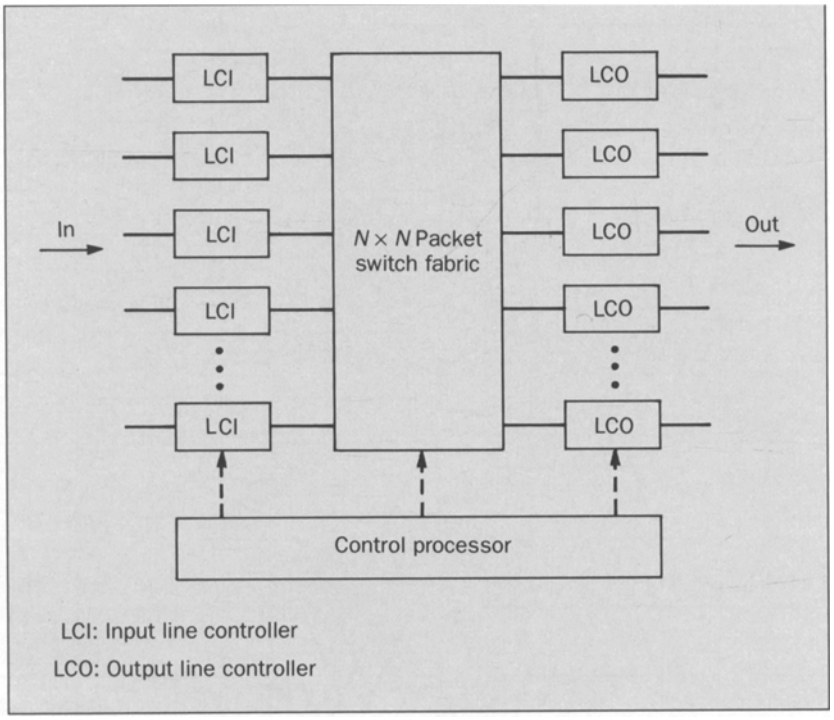


**Figure 13. Fast packet switch generic architecture.**



be more in any packet-synchronous cycle, the excess over eight is "knocked out" as in a tournament, which accounts for the name of the scheme. The architects designed a structure that implements the tournament using basic switch elements. There is, of course, an implied priority order.

The knock-out switch structure shows a fully connected network, for example $N$ inputs and $N$ outputs. Each input is connected to each output. However, the $N$ outputs are fed into a knock-out tournament hardware structure, which reduces the $N$ to $K$ outputs. The considerations for selecting $K$ are similar to Huang's, i.e., output ports could suffice to keep the packet loss probability below $10^{-6}$ (see Figure 12).

While the explosion in crosspoints of $O(n^2)$ seems to make large knock-out switches impractical, the knock-out principle can be used in other switch architectures to deal with output contention.

### Generic Third Generation Switch Architecture

We have shown a succession of space-division packet switching fabrics: the buffered Banyan; the Batcher-Banyan or Starlite fabrics with and without recirculation; unbuffered Banyan with feedback; contention-free Banyan; and the knock-out structure. Figure 13 shows how these "fabrics" can be put together to form a generic space-division packet switch architecture.

Lines are terminated in trunk (link) controllers. These may provide buffering and header translation on input or output side or in both places. Finally, a most important component of the switch is the controller, that holds the service and operational logic, the "Stored Program Control."

### Conclusion

The network of the future must become much more flexible in dynamically allocating bandwidth and providing responsive control capabilities to handle different information types in the most appropriate way. Fast packet technology seems to be an ideal match to handle transmission and switching tasks in the new network. Progress in fast packet switches uses self-routing networks. These can be engineered to deal with the desired degrees of non-blocking or packet loss methods that rely on information units labeled with self-identifying instructions, and that may pertain to its source, destination, or intended treatment.

Recent progress reports of lab experiments seem to indicate that technological challenges can be overcome for access data rates in the hundreds of megabits per second and for moderate size switches.

### Acknowledgments

We gratefully acknowledge comments from Scott Knauer on an earlier version of this paper.

### References

1. J. S. Turner and L. F. Wyatt, "A Packet Network Architecture for Integrated Services," *IEEE Global Telecommunications Conference (Globecom)*, 1983, pp. 45-50.
2. P. Baran, "On Distributed Communication Networks," *IEEE Transactions on Communications Systems*, vol. CS-12, March 1964, pp. 1-9.
3. A. Thomas, J. P. Coudreuse, and M. Servel, "Asynchronous Time-Division Techniques: An Experimental Packet Network Integrating Video Communication," *Proceedings of International Switching Symposium (ISS)*, Florence, Italy, May 7-11, 1984, Session 32C, Paper 2, pp. 1-7.
4. P. E. White, "Broadband ISDN—The Next Generation Broadband Telecommunications Network," *IEEE International Conference on Communications (ICC)*, Toronto, Canada, June 22-25, 1986, pp. 385-387.
5. J. J. Kulzer and W. A. Montgomery, "Statistical Switching Architectures for Future Services," *Proceedings of the ISS*, May 7-11, 1984, Florence, Italy, Session 43A, paper 1, pp. 1-6.
6. Andrew S. Tanenbaum, *Computer Networks*, Prentice-Hall, Englewood Cliffs, New Jersey, 1981, pp. 22-23.
7. D. Rouse, R. J. Spires, and R. E. Wallace, "The Number 2 Signal Transfer Point: An Overview of the AT&T Common Channel Signaling Packet Switch," International Conference on Computer Communication (ICCC), Munich, West Germany, September 15-19, 1986, Session B7, Paper 6, pp. 370-374.
8. Gottfried W. R. Luderer, "Architecture of the AT&T No. 1 Packet Switching System," *Proceedings of the 7th International Conference on Computer Communication (ICCC) '84*, Sydney, Australia, North Holland Publishing, Amsterdam, 1984, pp. 485-490.
9. A. G. Fraser, "Datakit—A Modular Network for Synchronous and Asynchronous Traffic," *Proceedings of the ICC*, 1979, Boston, Massachusetts, pp. 20.1.1-20.1.3.
10. S. Nojima, M. Hashimoto, H. Fukuda, and E. Tsutsui, "High Speed Packet Switching Network for Multi-Media Information," *Proceedings of IEEE Computer Networking Symposium*, March 1986, Washington, D.C., pp. 141-150.
11. Howard Jay Siegel, *Interconnection Networks for Large Scale Processing*, Lexington Books, Lexington, Massachusetts, 1985.
12. A. S. Acampora, "A Multichannel Multihop Local Lightwave Network," Globecom, Tokyo, November 15-18, 1987, pp. 37.5.1-37.5.9.
13. L. R. Goke and G. J. Lipovski, "Banyan Networks for Partitioning Multiprocessor Networks," *First Annual International Symposium on Computer Architecture*, December 1973, pp. 21-28.
14. R. W. Muise, T. J. Schonfeld, and G. H. Zimmerman III, "Experiments in Wideband Packet Technology," *Proceedings of the 1986 Zurich Seminar on Digital Communication*, IEEE Catalog Number 86CH2277-2, March 11-13, 1986, pp. 135-139.
15. K. E. Batcher, "Sorting Networks and Their Applications," *American Federation of Information Processing Societies (AFIPS) Proceedings of the Spring Joint Computer Conference*, 1968,

pp. 307-314.
16. Allan Huang and Scott Knauer, "STARLITE: A Wideband Digital Switch," *Proceedings Globecom*, Atlanta, Georgia, November 1984, pp. 121-124.
17. S. C. Knauer, J. H. O'Neill, and A. Huang, "Self-Routing Switching Network," *Principles of CMOS VLSI Design*, Neil Weste and Kanran Eshraghian, eds., Addison-Wesley, Reading, Massachusetts, 1985, pp. 424-448.
18. M. J. Karol, M. G. Hluchyj, and S. P. Morgan, "Input vs. Output Queuing on a Space-Division Packet Switch," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-5, October 1987, pp. 1274-83.
19. A. E. Eckberg and T.-C. Hou, "Effect of Output Buffer-Sharing on Buffer Requirements in an ATDM Packet Switch" *Infocomm 88*, March 29-31, 1988, Session 5A, pp. 459-466.
20. M. G. Hluchyj and M. J. Karol, "Queuing in Space-Division Packet-Switching," *Infocomm 88*, March 29-31, 1988, Session 4A.3.1, pp. 334-343.
21. R. E. Staehler, J. J. Mansell, E. Messerli, G. W. R. Luderer, and A. K. Vaidya, "Wideband Packet Technology for Switching Systems," *ISS*, Phoenix, Arizona, March 16-20, 1987, CH2431-5/87/0000-0448 IEEE, pp. B6.1.1-7.
22. Hitoshi Imagawa, Shigeo Urushidani, and Koichi Hagishima, "A New Self-Routing Switch Driven with Input-Output Address Differences," *Globecom*, Hollywood, Florida, November 28-December 1, 1988, Session 49, Paper 1, pp. 1607-1611.
23. Jonathan S. Turner, "New Directions in Communications," *Proceedings of the 1986 Zurich Seminar on Digital Communication*, IEEE Catalog Number 86CH2277-2, March 11-13, 1986, pp. 25-32.
24. E. Arthurs, R. Boorstyn, and T. T. Lee, "The Architecture of a Multicast Broadband Packet Switch," *Infocomm 88*, New Orleans, Louisiana, March 29-31, 1988, pp. 1-8.
25. J. Y. Hui and E. Arthurs, "A Broadband Packet Switch for Integrated Access," *Journal for Selected Areas in Communications*, Vol. SAC-5, No. 8, October 1987, pp. 1764-1273.
26. Floyd E. Ross, "FDDI—A Tutorial," *IEEE Communications Magazine*, Vol. 24, No. 5, May 1986, pp. 10-17.
27. D. K. Wilson, "A New Architecture for Packaging Wideband Communication Equipment Using a 3-D, Orthogonal Edge-To-Edge Topology," *Globecom 88*, Hollywood, Florida, November 28-December 1, 1988, 13.5.1-5, 430-434b34.
28. V. E. Beneš, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, New York, 1965.
29. C. T. Lea, "The Load-Sharing Banyan Network," *IEEE Transactions on Computers*, Vol. C-34, pp. 1025-1034, December 1986.
30. Y. S. Yeh, M. G. Hluchyj, and A. S. Acampora, "The Knockout Switch: A Simple Modular Architecture for High-Performance Packet Switching," *IEEE Journal of Select Areas of Communication*, Vol. SAC-5, pp. 1274-83.

Biographies (continued)

*plied research in next-generation networking products. He holds Dipl. Ing. and Dr.-Ing. degrees in electrical engineering from the Technical University of Braunschweig, FRG.*

*Mr. Luderer joined Bell Laboratories in 1965.* ***Avinash K. Vaidya*** *is a group supervisor in the International Switching Systems European Performance Verification Department at AT&T Bell Laboratories Indian Hill facility in Naperville, Illinois. He is responsible for verification of switching systems for customers in the UK and Spain, and for supervision of system verification planning in other European countries. Mr. Vaidya holds a B.E. in electronics from Devi Ahilya University, Indore, India; an M.E. in electronics from the Netherlands Universities Foundation; and an M.S. in computer science and Ph.D. in electrical engineering, both from the University of Wisconsin, Madison. Mr. Vaidya joined AT&T in 1982.*

*(Manuscript received March 21, 1989)*