Technologies for Broadband Switching

Mark A. Pashan
Marius D. Soneru
Gary D. Martin

For an asynchronous transfer mode (ATM) network to provide a wide range of services, high-capacity and low-cost network elements are needed. The capacity and cost of these elements are determined by the costs of the underlying hardware and software devices necessary to build them: ATM line cards, switching fabrics, and control structures. These devices, in turn, are built upon the advanced technologies of digital devices, optics, algorithms, and architectures. In this paper, we explore the applications of these technologies to develop ATM line cards, ATM switching fabrics, and ATM control structures.

Introduction

Advances in the fields of integrated circuits and fiber optics have led to the possibility of people communicating and being entertained as never before. Fax machines, mobile phones, laptop computers, supercomputers, wide area networks, videophones, and video distribution systems are common in the marketplace. Now, the worldwide telecommunications network must change to provide the flexible, multi-service, high-capacity, and low-cost communications these products will require. One important change in the telecommunications network is the introduction of asynchronous transfer mode (ATM) technologies and systems.

ATM is the basis of the broadband integrated services digital network (BISDN), the international standard for the next generation telecommunications network. Unlike time-slotted circuit-switched systems, such as the 4 ESS® and 5ESS® switches, which also are called synchronous transport mode (STM) systems, ATM systems are based upon a fixed-length packet or cell technique. BISDN promises flexibility, service integration, and high bandwidth. ATM technology is a flexible platform that can quickly implement new BISDN services, while providing significant operations, administration, and maintenance (OA&M) savings due to service integration. The vision of voice, data, and video services provided via a single network is possible with ATM.

To implement this vision, an ATM system should support a wide range of traffic types. The system should be capable of supporting:

- The low-jitter restriction of high-bandwidth, constant bit rate (CBR) traffic, such as clear-channel leased lines and video connections,
- The high burst rates of high-speed data connections, which could generate bursts of hundreds of cells, yet generate low cell-loss rates, and
- High use for the most important types of traffic mixes in a network, to keep the overall cost per service low.

To satisfy the needs of the telecommunications network, the ATM system should be highly reliable, easily maintainable, gracefully expandable, and cost-effectively evolvable.

AT&T has been working on high-bandwidth statistical communications for a long time, and was the first telecommunications equipment vendor to demonstrate high-bandwidth packet communications, with live customer traffic, in a wideband packet field experiment. That work, and its later laboratory extensions, were among the first steps toward ATM.[1,2] AT&T's broadband efforts are continuing to pursue a set of ATM technologies designed to meet emerging customer needs.

The ATM technologies for cell transport are the ATM line cards, the ATM switching fabric, and the ATM control structure, as shown in Figure 1.

Panel 1. Acronyms Used in This Paper

AIS — Alarm indication signal
ASN.1 — Abstract Syntax Notation language
ATM — Asynchronous transfer mode
BISDN — Broadband integrated services digital network
CBR — Constant bit rate
CNM — Customer network management
DS-3 — Data rate of 45 Mbits/s
FERF — Far-end receive failure message
GUI — Graphical user interface
HEC — Header error control
ISDN — Integrated services digital network
Mux/demux — Multiplexer/demultiplexer
OA&M — Operations, administration, and maintenance
OC-3 — Optical carrier rate of 155 Mbits/s on optical facility
RAM — Random access memory
STM — Synchronous transfer mode
STM-1 — Synchronous transfer mode signal rate of 155 Mbits/s on optical facility
STS-1 — Synchronous transfer signal rate of 53 Mbits/s on electrical facility
STS-3 — Synchronous transfer signal rate of 155 Mbits/s on electrical facility
VC — Virtual channel
VBR — Variable bit rate
VCI — Virtual channel identifier
VP — Virtual path
VPI — Virtual path identifier
VLSI — Very large scale integration

- The *line cards* are based on high-speed VLSI technology and optics to implement the ATM processing functions.
- The *switching fabric* uses a cell-expansion network, followed by a shared-memory switching fabric, that allows the AT&T system to handle a wide range of traffic types.
- The *control structure* allows for modular growth and increasing service functionality.

The first product based upon AT&T's ATM technology is the GCNS-2000, developed by the Operations Systems Data Business Unit.

System Principles

AT&T's design philosophy is to meet the long-term, broadband service traffic needs with an expandable and modular architecture that also meets the highest performance standards. The various ATM line cards (such as those interfacing DS-3, STM-1, STS-3c, OC-3, etc., data streams) are plug compatible, which means various input/output configurations can be mixed and matched in one switching-unit shelf, using standard line-card slots for a variety of line-card terminations. This flexibility will meet the wide range of services and network configurations proposed for ATM systems.

The ATM switch fabric architecture provides a way to modularly construct large ATM fabrics out of ideal performance memory modules, that is, modules in which no input or output can block any other input or output to affect overall delay/throughput performance. The architecture consists of a bufferless front-end cell expansion into an array of shared-memory ATM modules, where optimal buffering can be achieved. The overall architectural design is based on several key features:

- Increased *processing power* of the ATM line cards,
- A front-end *expansion/concentration network* that provides self-routing, near-optimal and conflict-free distribution of cells to the shared-memory switching fabric, which eliminates internal congestion,
- A *shared-memory switching fabric* that requires minimal RAM buffering to route traffic to the appropriate output ports, and
- *Buffered output port queuing* that yields the best delay/throughput performance.

These features are discussed in this paper.

The ATM control structure is based upon distributed processing for modularity and growth. The three-level hierarchy of distributed processors map into the ATM Protocol Reference Model. (For a description of the ATM Protocol Reference Model, see "Network Aspects of Broadband ISDN" in this issue). The overall architecture is duplicated for reliability, and unique cell processing algorithms allow hitless, or errorless, switching between active and inactive units with minimal overhead. Finally, the modularity of the architecture allows for the components to independently evolve.

ATM Line Cards

Line cards provide the interface between the physical media used to transport ATM cells and the ATM switch fabric. The physical media may be either electrical or optical, and there are many possible data rates and signal formats (such as DS-3, STS-1, STM-1, OC-3, etc). Many types of line cards are needed to accommodate the wide range of transport options.

Except for very high-rate signals, line cards generally contain multiple port interfaces per card, supporting multiple access or egress lines. Advances in device, optics, and packaging technology are allowing more functions and more ports to be packaged onto each line card. Line cards are usually bi-directional, but there are applications, such as one-way video on demand, in which the source of a line card input is different from the destination of its output.

Figure 2 depicts some of the signal-processing functions that could be performed by a quad, or four-port, STM-1 line card. The functions for a generic line card are described below.

Ingress Functions. For cells coming into the network switch, ATM line card hardware first synchronizes the various received signals with an associated recovered clock. Then, the line card will locate the transmission frame pattern, such as a DS-3 or STM-1, and then strip the transmission overhead from the signal. Next, the actual ATM cell boundaries are located from the information in the header error control (HEC) field. The HEC field also is used to correct errors in the header or, if the errors are too severe, the errored cells may be dropped entirely. The ATM cell payload then is descrambled, meaning the pseudo-random signals, inserted in the data stream for clock recovery (also called 1's density), are removed.

Cells are presented to the ATM switch fabric before they are synchronized to the system cell clock. This is done by passing the cells through an elastic store, or expandable buffer, and inserting or deleting idle cells as needed. Once cells from different input ports are synchronized to a common clock, they can be interleaved to form a higher-rate cell stream.

Several ATM processing functions also are performed by line cards. For each virtual path (VP), that is, a logical path through the network that can be changed in real time, and for each virtual channel (VC) in a virtual path, the received cells are counted, and the bandwidth "policed" for excessive user throughput. This policing is done by matching the customer's input to the allocated virtual path or virtual channel data rate, as determined by the customer when the service was provisioned. If the cell rate exceeds the provisioned rate, the cells either can be dropped or they can be marked for possible elimination later, if they encounter congestion in the network.

Line cards also translate the incoming VP/VC identifier (VPI/VCI) fields in each cell from the incoming value, or address, to the outgoing value. The incoming VPI/VCI identifier is used to access, or index, a routing table, which contains the destination-port code for the cell. This destination-port code is then attached, or prepended, to the cell for routing by the switching fabric. ATM operation, administration, and maintenance (OA&M)

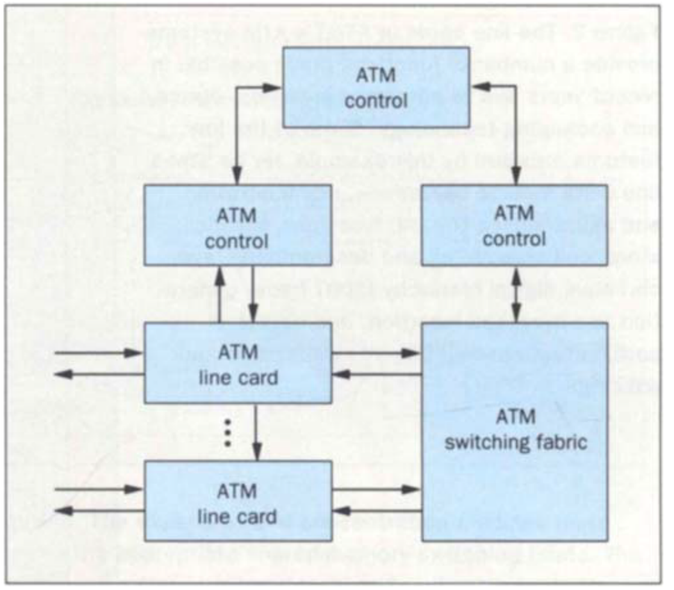

Figure 1. Three key elements in AT&T's ATM switching products are the ATM line cards, the ATM switch fabric, and the ATM control structure. *Line cards* were developed to interface a variety of electrical and optical media at a variety of rates. The cards also provide the "low-level" control functions of the three-level control structure. The *switch fabric* uses cell-expansion and concentration modules and a shared-memory switching fabric to handle a variety of traffic types with minimal buffering, providing low delay and optimal throughput. The *control structure* includes line card software, and two other levels of control, that are mapped to the control functions required in each level and plane in the ATM Protocol Reference Model.

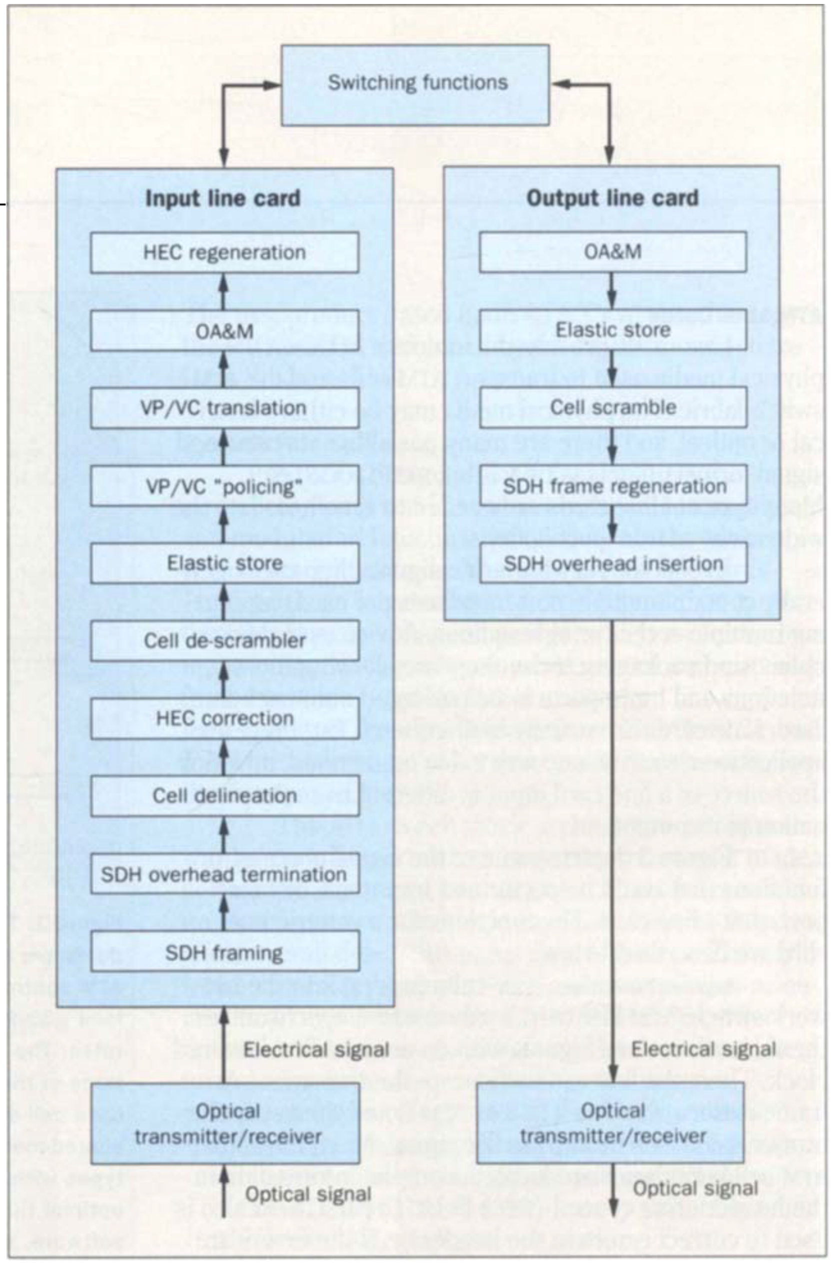

Figure 2. The line cards in AT&T's ATM systems provide a number of functions made possible in recent years due to advances in device, optics, and packaging technology. Some of the key features provided by this example, for an STM-1 line card, include operations, administration, and maintenance (OA&M) functions, elastic store, cell scrambling and descrambling, synchronous digital hierarchy (SDH) frame generation and overhead insertion, and virtual path/virtual channel (VP/VC) translation and policing.

functions, such as generating virtual path and virtual channel alarm indication signals (AISs), also are done by line cards. Finally, the head error-control field is recalculated for all cells whose headers are modified by the line cards.

Egress Functions. After cells are routed by the switch fabric, which will be discussed below, they are returned to output-line cards for further processing. If an output-line card supports multiple output ports, cells arriving from the switch fabric are demultiplexed to their appropriate ports. OA&M functions also are performed, such as generating a VP/VC far-end receive failure (FERF) message, which informs the switch to discontinue sending data when a line failure is detected. Any non-standard ATM cell bits (such as prepended destination-code bits) are removed, and the cell rate is adapted to the available line rate by passing the cells through an elastic store and adding or deleting idle cells when necessary.

The cells are then scrambled, that is, pseudorandom signals are inserted in the data stream for clock recovery, and are packed into the payload signal of the transmission format generated by the line card. The resulting data stream is then converted into the line electrical, or optical, signal for transmission to the next ATM switch in the network, or to the end user.

ATM Switching Fabric

Broadband ATM cell switching has been extensively studied at AT&T Bell Laboratories, where significant contributions have been made in fundamental theory, as well as innovative fabric designs.[2-6] Throughout these research efforts, the critical issues of self-routing interconnect, destination conflict resolution, queuing arrangement, delay and throughput performance, and switch complexity were studied carefully and now are well understood.

Although various switches can, on paper, be designed to large dimensions, technological and physical constraints often impose a practical limit on their maximum size, say $M \times M$. If we want a larger switch system, say $2M \times 2M$, then multiple $M \times M$ switches have to be interconnected. A conventional way to do so is to employ three columns of $M \times M$ modules. This is an effective and proven technique in circuit-switched systems, but a hardware-expensive solution, nonetheless. Although each individual switching module can be constructed as the *ideal* cell switch, with optimum throughput, the complete multiple-column switch fabric no longer reflects the optimal design conceived for its individual switch modules. A principal disadvantage in such multistage-buffered approaches is the performance degradation due to internal congestion, usually arising from non-uniform traffic loads. Therefore, it was our architectural goal to achieve the same high performance expected of an ideal module for the *entire* switch fabric, regardless of traffic conditions.

The switch fabric architecture is depicted in Figure 3. The ATM switch fabric contains two stages, the cell expansion and concentration network, and the shared-memory switching fabrics. The expansion and concentration network routes cells to the appropriate shared-memory fabrics. The shared-memory fabrics provide the last stage of switching, and most of the queuing, required for the ATM traffic.

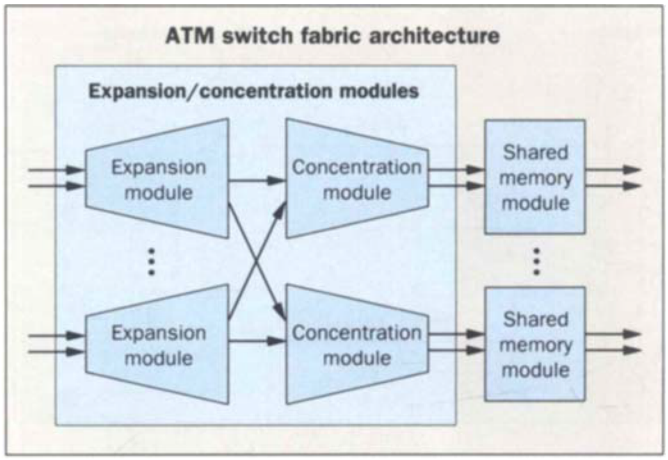

Figure 3. The expansion and concentration modules route cells to the appropriate shared-memory switching fabric. The expansion module is memoryless, and routes the cells to the concentration module. This module, in turn, requires a small amount of buffering, since it provides 32 inputs and eight outputs to the shared-memory switch fabric.

Cell-Expansion Module. When the input cell arrives at the shared-memory module, it is first processed by the cell-expansion module. The cell-expansion network is memoryless, that is, buffering is not used in this module. Instead, all the input cells are cleared out in each time slot in the next clock cycle, and there is no internal congestion. Because the input cell is self-routing, due to the prepended destination-port code, path conflicts are automatically resolved without the need for a centralized controller. The function of the cell-expansion network, therefore, is merely to transport cells to each concentration module in every slot.

It should be noted that address discrimination can be performed at both the cell-expansion module and the concentration module. Studies indicate that for small switching systems, address discrimination is more efficient at the concentration module; for large systems, it is more efficient at the cell-expansion module. We describe the former situation here.

The cell-expansion module copies an incoming cell to all of the output ports leading to the concentration modules. For example, if the module is 4×4 (four inputs and four outputs), a single input would be written to all four outputs, that is, transported to all four concentration modules connected to the cell-expansion module.

Concentration Module. The concentration module, in turn, funnels the cells into an associated shared-memory switch fabric. Unlike the cell-expansion module, which is connected to multiple concentration modules, each concentration module is connected to only one shared-memory module (excluding redundancy for reliability). Since the expansion module broadcasts the cell to multiple concentration modules, each concentration module first must determine the address of the cell, and drop any cell not destined to its shared-memory module. The appropriately addressed cell is then routed to the shared-memory fabric. A small amount of buffering is required in the concentration modules, since the module accepts 32 inputs, but delivers only eight outputs in one cell time, or clock cycle. The 32×8 concentration is possible due to the statistical nature of data transmission, wherein not all of the inputs will be active at the same

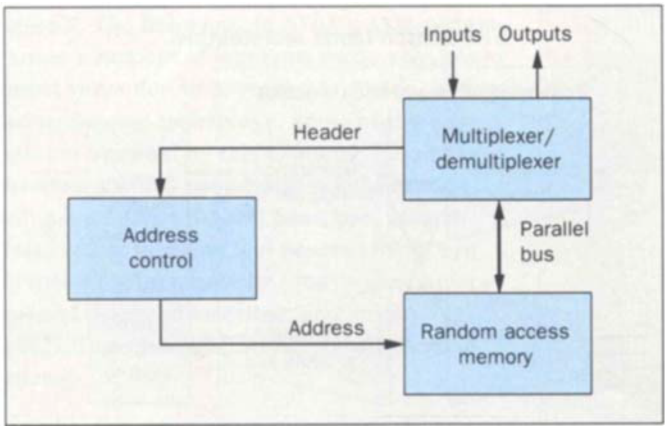

Figure 4. AT&T's ATM shared-memory module architecture is shown in this illustration. Input cells are routed to the multiplexer/demultiplexer, and header information is routed to the address control, while the user data is stored in RAM until processing is completed. Then the user data is read from RAM to the multiplexer/demultiplexer for transmission.

time. Therefore, since the chance of contention is statistically small, only a small amount of buffering is required. The control consists of a single queue with up to 32 'writes' and eight 'reads' to the queue in one cell time.

There are three basic approaches for providing the queuing in an ATM switch: at the input port, throughout the fabric, or at the output port. It has been shown that the best possible delay/throughput performance for an ATM switch is achieved by output queuing,[7] that is, by providing any queuing *after* the switching functions are performed. This approach has been taken by AT&T. Most cell buffering in the ATM switch is done at the shared-memory modules, and the optimal output-queuing design of the shared-memory modules ensures the best delay/throughput performance for the entire switch.

Therefore, the input cell arrival rate to the switching fabric has no effect on the design of either the cell-expansion network or the cell-concentration network. This is because the input cells are cleared into the concentration modules on a slot-by-slot basis. The only effect is on the sizing of the buffer space, inside the concentration module, to handle the statistically small amount of contention expected at this 32×8 stage, and the buffer space required for the shared-memory module.

Shared Memory Module. The general architecture of the 20-Gbits/s shared-memory module is shown in Figure 4. The ATM fabric cells (53-byte ATM link cell enveloped in a larger fabric cell) arrive at and leave from each module synchronously. The shared-memory module writes arriving cells into, and reads leaving cells out of, a single, large random-access memory (RAM). It should be emphasized that there is no buffer-specific RAM—thus the name, shared memory. Indeed, the entire RAM in the shared-memory module could be used for buffering, if necessary. The high-bandwidth RAM is accessed via a wide parallel bus. By controlling the addresses for the RAM read and write operations, logical output buffers for each of the module outputs can be created, as required, within the single physical RAM.

The data path is composed of a multiplexer/demultiplexer (mux/dmux), a wide data bus, and the large RAM. The mux/demux multiplexes the incoming and outgoing cell streams onto the wide data bus. The data bus also is connected to the RAM data inputs—which could, as noted, be as wide as an entire fabric cell—to achieve the high RAM access bandwidth. The address control block uses copies of the switch fabric headers as input to generate the RAM address information. The operations of each module are pipelined, that is, an operation is divided into subprocesses that are processed simultaneously.

In each cell time, the 8×8 (for eight inputs and eight outputs) shared-memory module performs eight write operations (one for each arriving cell) and eight read operations (one for each leaving cell). In general, the address control logically contains an output control queue associated with each module output. Each output control queue holds the RAM addresses of all the cells stored in the RAM destined for a particular module output. The process of possibly dropping cells that exceed the customer's contract is controlled by the terms of the provisioning process, the size of the RAM, and, of course, by ever-changing network conditions. If conditions are favorable, the customer's cells will not be dropped.

Output Queuing. When multiple cells arrive simultaneously, for the same output port in the 8×8 switch, they have to be buffered. The output can only serve one cell at a time. As previously noted, it has been shown that the best possible delay/throughput performance is achieved by buffering, or queuing, at the output port. That is, all cells in each time slot are swept through the switch fabric to their destinations such that, if several cells arrive at the same time at the same output, they can be buffered in respective output queues. They are not delayed by, nor do they delay, cells destined for other outputs. The comparison between delay-throughput performance for input queuing, versus output queuing, is shown in Figure 5.

Consequently, the shared-memory fabric is a 20-Gb/s ATM output-buffered switch module. If the switch is expanded later, the concentration modules and the shared-memory modules can be reused.

ATM Control Structures

The control structure is based on the mapping of the control functions—required by each layer and plane in the ATM Protocol Reference Model—into a three-level hierarchy of distributed control processors. The mapping takes into account the characteristics of the control functions, such as the required response time, processing needs, and expected reliability. It also identifies corresponding control processors with the associated execution environments.

First, the *low-level* control functions in the three-level control hierarchy are associated with the ATM layer, and are performed by processors residing on the ATM line cards. As previously described, the control functions performed here include source policing, VPI/VCI routing, traffic and performance measurements, and error detection and reporting. A powerful microprocessor with a real-time execution environment and a fast scheduling and dispatching algorithm implements these functions on each line card. The line card-based microprocessors communicate with the higher-level processors using a functional, message-based interface that hides many of the hardware-specific characteristics of the line cards from the higher levels.

The *middle level* of control in the hierarchy provides operational and administrative control for the ATM fabric and the line cards, including basic connection management, measurements collection, and event reporting. It also is responsible for the reliable operation of the ATM fabric and line cards, providing fault recovery and diagnostics.

The reliability requirements on public switching systems have led to duplexed (side 0/side 1) architectures. This is also true for ATM systems. In the event of an error in one side, the other side can take over instantly, since both sides normally operate in lock step.

For the duplex hardware, when fault recovery detects a fault in the active side, it switches to the standby side, initializes the newly active side, and returns the unit to service. Then fault recovery initiates a diagnosis of the faulty unit. This procedure leads to a high level of availability for the transport hardware. These functions are implemented by a highly reliable, fully duplicated processor, with a real-time execution environment that is capable of continuous operation and fast response time.

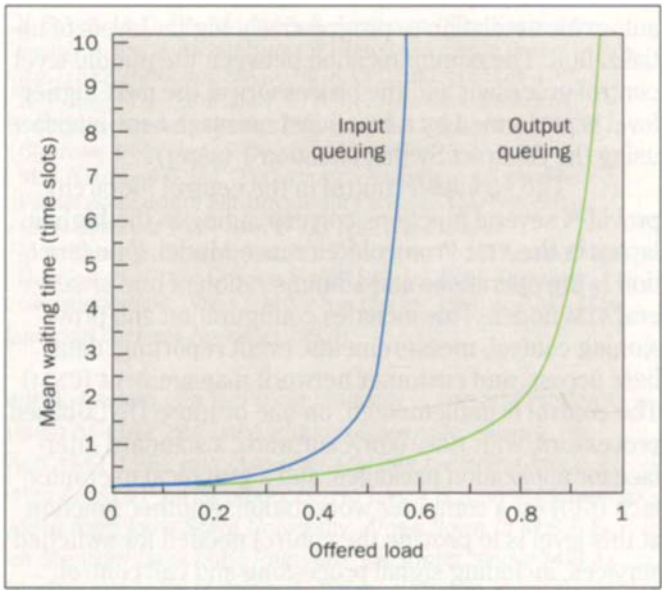

Figure 5. This figure shows the delay/throughput for input versus output queuing. If input queuing is employed, the waiting time increases significantly with an offered load of 0.5. With output queuing, a similar waiting time is not reached until an offered load of 0.8 is experienced. Thus, AT&T's ATM technology uses output queuing based on statistical probability.

Operating practices often result in occasional side switches, so that each side can be routinely tested. This is known as routine exercises. In the event that an active side detects a fault while the inactive side is completing its testing procedures, the unloaded inactive side takes over the switching operation from the loaded active side in a short period of time. The GCNS-2000, for example, handles this feature with a fast synchronization procedure.

The control software is designed for reliable operation by using a number of fault-tolerant techniques: data is continuously validated by background audit programs and software execution errors are detected, leading to corresponding recovery actions, including the automatic escalation to progressively higher levels of initialization. The communication between the middle-level control processor and the processors at the next higher level is performed by a functional message-base interface using the Abstract Syntax Notation 1 (ASN.1).

The *high-level* control in the control hierarchy provides several functions corresponding to the higher layers in the ATM Protocol Reference Model. One function is the operations and administration of one or several ATM nodes. This includes configuration and provisioning control, measurements, event reporting, database access, and customer network management (CNM). The control is implemented, on one or more UNIX-based processors, with BaseWorX software, a standard interface for application packages, and a graphical user interface (GUI) on a computer workstation. Another function at this level is to provide the control needed for switched services, including signal processing and call control. This is implemented by several highly reliable processors with real-time execution environments.

The different processors in the control hierarchy have different characteristics that match the needs of the functions they perform. They provide real-time or time-sharing environments, depending on the response time desired. They can provide open interfaces for programming customer-specific applications; reliable operation with high-availability, for example, call processing; and reliable transaction operation, such as operations and administration. The distribution of control among processors, and the well-defined interfaces between the three levels of processors in the hierarchy, lead to a modular and flexible control architecture that can evolve with the ATM network.

Cell Transport Performance

ATM cells are of fixed length. Cell traffic will have certain performance characteristics, due to the switch's incoming and outgoing link bandwidth restrictions (such as whether they are DS-3, STS-1 etc.) and the resulting traffic queuing. Unlike input-buffered fabrics that can suffer from head-of-line, or input-side, blocking, and Banyan fabrics that have internal link contention, output-buffered fabrics allow the incoming traffic to flow to the outgoing buffer/link without contention. In such a system, buffers are only required for those cases when traffic from multiple inputs are simultaneously destined for a single output, such as if there is output contention.

Table I. Utilization Comparisons

Fabric Type	Dedicated Buffering	Shared Buffering
Allowed Link Utilization	45%	88%

The size of the output buffer needed is dependent on the allowed cell-loss rate, incoming link utilization, the characteristics of incoming traffic, and the buffering technique. A 10^{-9} cell-loss rate and 80 percent incoming link load, for example, requires 45 cell buffers at each output port of a dedicated output-buffer system for random traffic. If the incident traffic is not random, but is instead bursty, such as when long packets are broken into a series of cells, then holding the output-buffer size fixed, and holding the cell-loss probability constant, results in a decreased system throughput. In fact, burst size, link utilization, and cell-loss probability are interrelated. For a given cell buffer size and technique, increasing the incident-burst size, or increasing the link utilization, or decreasing the cell-loss probability, will degrade one or both of the remaining parameters.

The only way to avoid this situation in dedicated output buffering systems is to increase the output buffer size. In other words, larger output buffers are required to carry increasingly bursty traffic at high-utilization and low cell-loss rates. In fact, the output queue size scales almost linearly with burst size, given the same cell-loss rates and link utilizations.

For example, if a switch routinely experiences burst rates of 10 cells in a row, and there is a high probability of multiple bursts, the switch might require, say, 50 output buffers. In a linear fashion, then, a switch experiencing bursts of 100 cells would require a linear increase to 500 output buffers, which greatly increases the hardware cost of the switch.

Another method to handle the problem of bursty traffic, and the one used by AT&T to reduce the linearity problem, is that of sharing the output buffers between the outputs. In other words, use one shared buffer for all of the outputs, such that space unused by one output can be reused by another output. The impact of such a technique is shown in Table I for a traffic type consisting of a mean-burst length of five cells, a 10^{-10} cell-loss rate, a fabric size of 8×8, and total buffer of 8000 cells. The shared-buffering technique allows almost twice the link utilization as the dedicated buffering technique for this traffic type. Consequently, systems designed with a shared-buffering technique are better able to support a wide range of services with a minimal amount of buffering and, consequently, with minimal switch cost.

Due to these theoretical findings, the AT&T system uses memory-based modules that implement shared-output buffering. These architectures provide ideal delay/throughput performance characteristics and use a minimum of buffering for any traffic type or mix, such as constant bit rate and variable bit rate (VBR). Additionally, the AT&T system will incorporate delay and loss priorities to segregate traffic types and provide another level of congestion control.

Summary

AT&T's ATM technology has been designed to provide a wide range of services in an efficient and cost effective manner. The ATM line cards can be interchanged on a single shelf, providing flexible configurations. The ATM switch fabric has ideal output-buffering performance for up to 20 Gbits/s, and near ideal performance for systems larger than 20 Gbits/s. The fabric has the ability to carry random and bursty traffic at high-load and low cell-loss rates as appropriate. No fabric rearrangement or central processor routing is required. The ATM control structure is modular and extensible. The system is duplexed for reliability and fault tolerance. These technologies provide the basis of AT&T's ATM-based products, such as the GCNS-2000, the ATM Cross-connect, and the Broadband Switching System-2000.

References

1. R. W. Muise, T. J. Schonfeld, G. H. Zimmerman, "Experiments in Wideband Packet Technology," 1986 International Zurich Seminar on Digital Communications, March 11–13, 1986.
2. A. K. Vaidya and M. A. Pashan, "Technology Advances in Wideband Packet Switching," Globecom 1988, Hollywood Florida, Nov. 28–Dec. 1, 1988.
3. A. Huang and S. Knauer, "STARLITE: A Wideband Digital Switch," *Globecom '84 Conference Record*, Nov. 1984, pp. 5.3.1–5.3.5.
4. Y. S. Yeh, M. G. Hluchyj and A. S. Acompora, "The Knockout Switch: A simple Architecture for High Performance Packet Switching," *IEEE JSAC*, Vol. SAC-5, No. 8, Oct. 1987, pp. 1274–1283.
5. K. Y. Eng, M. J. Karol and Y. S. Yeh, "A Growable Packet (ATM) Switch Architecture: Design Principles and Applications," *IEEE Globecom 1989*, Nov. 27–30, pp. 1159–1165.
6. M. J. Karol and C. L. I, "Performance Analysis of a Growable Packet Architecture For Broadband Packet (ATM) Switching," *IEEE Globecom 1989*, Nov. 27–30, pp. 1173–1180.
7. M. J. Karol, M. G. Hluchyj and S. P. Morgan, "Input Versus Output Queuing on a Space-Division Packet Switch," *IEEE Transactions on Communications.*, Vol. COM-35, No. 12, Dec. 1987, pp. 1347–1356.

(Manuscript approved November 1993)

Mark A. Pashan *is a technical manager in the ATM Platform Division at AT&T Bell Laboratories, Naperville, Illinois. He joined AT&T 1982 and manages ATM cell switching fabric design for AT&T's ATM products. He has a B.S. in engineering science from Iowa State University, Ames, Iowa, a M.S.E.E from Stanford University, Palo Alto, California, and a M.B.A. from the University of Chicago, Chicago, Illinois.*

Marius D. Soneru *is a technical manager in the Switching ATM Platform Department at AT&T Bell Laboratories in Naperville, Illinois. His group designs and develops software for the ATM platform infrastructure, inter-processor communications, initialization, and measurement. He joined AT&T in 1973 and holds a M.S.E.E. degree from Polytechnic Institute of Bucharest, Hungary, an M.S. degree in computer science from the University of California at Los Angeles, and a Ph.D. in computer science from the Illinois Institute of Technology in Chicago.*

Gary D. Martin *is a technical manager in the ATM Core Hardware Department at AT&T Bell Laboratories, North Andover, Massachusetts. He joined AT&T in 1978 and currently supervises a group that designs VLSI devices. He has a B.S.M.E. and M.S.M.E. from Oklahoma State University in Stillwater, and an M.S. and Ph.D. in electrical engineering from Stanford University, Palo Alto, California.*