An ATM strategy for IBM networking systems

by M. E. Cohen
D. S. Abensour

Advances in networking technology have created new opportunities for information presentation. In turn, new opportunities have created new demands for bandwidth and services, and the need for a new kind of networking. In this essay we explore the kind of networking that is needed, explain how the asynchronous transfer mode (ATM) technology answers that need, and discuss IBM's strategy for providing networking services built around ATM technology.

Why is there suddenly such a focus on bandwidth demand? What has allowed us to conceive of an information superhighway? In addition to a hunger for more and more widely dispersed information, advances in technology have created new opportunities for the presentation of this information. Thus, the arrival of multimedia applications, fiber-optic transmission, and high-speed electronic switching indicates the urgent need for a new kind of networking. We believe that asynchronous transfer mode (ATM) technology offers the best way to meet this need.

Today's networking environment is characterized by a diverse mix of topologies, geographic spans, carrier services, equipment, interfaces, physical media, and transmission speeds. While many business processes and applications were and are well-served by the networking technologies that support this environment, new applications and expanded communications requirements for existing applications are creating a whole new set of networking challenges.

Emerging needs for local area network (LAN) interconnection, image processing, multimedia services, and client/server relationships, to name only a few examples, stretch the capabilities of today's networks in terms of both transmission speed and gross capacity. Users are becoming increasingly sophisticated and demanding in their expectations of computing services, thereby stimulating the development of increasingly complex applications that draw from diverse forms of information, e.g., voice, video, and data, and hence are termed "multimedia." These emerging multimedia applications will not only serve to entertain their users, but will also provide the competitive edge that business constantly seeks. Taken together, these observations suggest that the networking demands of the near future cannot be met by a simple extrapolation of the current networking methods and products.

The requirements of the emerging networking environment present a multifaceted challenge:

- To provide dynamic and transparent connectivity to join wide-area and local-area domains as well as to integrate private and public subnetworks. Any given enterprise may find it desirable or even necessary to create a network infrastructure consisting of some combination of local- and wide-area, and public and private components. Key to the success of this hybrid network is the ability to manage it across all segments.

 0018-8670/95/$3.00 © 1995 IBM

- To support applications with unprecedented needs for high-transmission speeds, large transmission capacities, the flexibility of bandwidth on demand, end-to-end quality of service, and effective network-management capabilities
- To offer products that enable the creation of networking solutions characterized by lower operational costs and better price-to-performance ratios

Because of the inherent complexity resulting from mixed technologies, it becomes extremely difficult to satisfy all these requirements with existing networking tools. For example, in order for an Ethernet-attached client in a branch office to access a token-ring-attached server in a different location, at least three different networking protocols must be made to communicate. While it is clear that this scenario will be extant for some time to come, the goal is to provide a technology that will not only provide the desired simplification, but will also support those emerging applications that will demand end-to-end quality of service and bandwidth reservation capability. Asynchronous transfer mode (ATM) provides the foundation for this technology.

What is ATM?

ATM is the international standard for cell relay, which the industry is developing

- To support both public and private networks
- To use the same technology for local area, metropolitan area, and wide area networks
- To multiplex voice, video, and data (including image) traffic
- To deliver bandwidth on demand
- To offer low-cost networking and use low-cost technology

To understand the ATM philosophy, consider first one view of networking history.

In the past, networks were either circuit-oriented or packet-oriented, with circuit-oriented services traditionally preferred for delay-sensitive (isochronous) applications such as telephony, and packet-oriented services preferred for delay-tolerant applications such as data transmission.

With circuit-oriented services, the entire bandwidth of a communications link is dedicated to a specific traffic connection, e.g., a telephone call. While this dedicated connection guarantees predictable service, it is also characterized by an inefficient use of the bandwidth. For example, in the case of a telephone call, when no one is talking, the bandwidth lies unused.

Packet switching offers considerably more versatility than its circuit-oriented alternative. The bandwidth in a particular link is not dedicated to a single connection; therefore, multiple sources can utilize this link by having their traffic "packets" multiplexed over it. This minimizes the idle time on the link, but also introduces an element of randomness—the packets are of varying lengths and they arrive at random times. If, for example, a voice or video connection were to be multiplexed with a data transmission, large data packets could interrupt the flow of the voice or video traffic, resulting in breakup or "jitter." That is obviously not acceptable to this isochronous form of traffic. Nevertheless, packet switching optimizes the use of bandwidth, which makes packet switching the more economical choice for data-oriented transmission.

ATM combines the best of both worlds to provide the predictability of circuit switching with the flexibility of packet switching. To accomplish this, ATM uses simplified packet-switching techniques, but segments the packets into 53-byte cells—each with 48 bytes of user data and a 5-byte header—to be switched into multiple virtual channels and paths that operate at transmission speeds into the gigabit-per-second range. Since the switching can be done in hardware, as described later, overhead inefficiencies can be held to a minimum. Isochronous traffic and data cells can be interleaved in a way that allows quality-of-service guarantees. In effect, ATM has isochronous support "built in"; consequently, ATM can transport voice, data, and video, all on the same circuit, as shown in Figure 1.

Although ATM was first proposed by the public telecommunication carriers to address the projected multimedia requirements in the wide-area environment (e.g., picture phones), ATM is not just a wide-area technique. A LAN can be built from the same high-speed ATM technology. To provide local-area users with dedicated bandwidth, an ATM LAN can be configured around a central switch, thereby avoiding the multiple-access use of a shared transmission medium that typifies earlier LANs. Using a central switch in this way provides a clear-cut advantage over the shared-medium approach. The speed of the device adapter is uncoupled from the

Figure 1 **ATM combines voice, data, and video traffic on a common circuit**

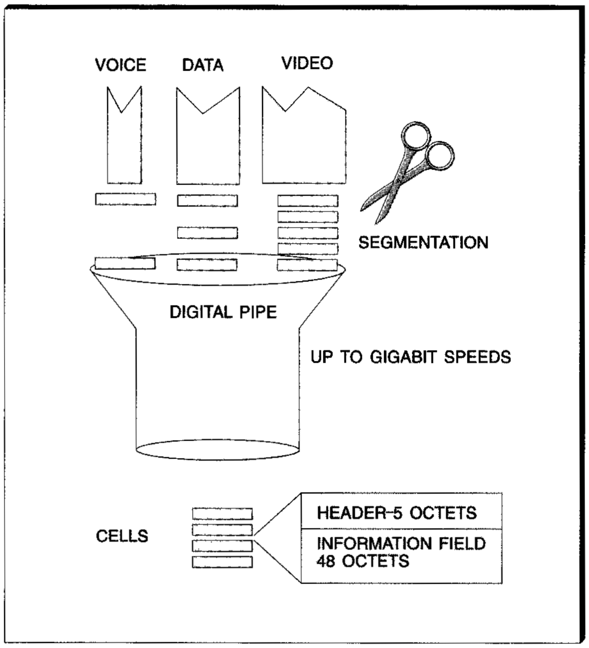

aggregate speed of the switching fabric, which means that the device adapter for a switch-based LAN can be a slower speed adapter than a shared-medium LAN of comparable performance. Moreover, an "ATM LAN" is scalable, that is, its attaching devices can employ differing adapter speeds, so a server can run at a faster speed and have greater throughput than its clients.

What is the commercial value of ATM?

State-of-the-art technology is commercially valuable only to the extent that it adds value to a customer's enterprise. Such value may be a direct financial saving or may be a better way to do things. ATM offers commercial value in both regards.

For the local-area user, ATM technology provides more bandwidth than today's LANs provide, and thereby supports new applications and improves the performance of present applications. For the wide-area user, ATM offers these same advantages along with enhanced bandwidth management and the provision of bandwidth on demand, both of which can reduce the expense of common-carrier telecommunication services.

Following are two examples of high-bandwidth contexts that can exploit the benefits of ATM:

- In the client/server model, geographically dispersed clients require access to multiple servers to obtain data and other services. Traffic often comes in bursts and sometimes spans both local and wide-area domains. Because large volumes of data are transferred, the aggregate bandwidth demands can exceed the capacity of today's LANs and WANs (wide area networks). The value of ATM in this context is twofold: (1) in providing the requisite bandwidth *on demand*, thereby reducing or eliminating bottlenecks, making possible new services and additional clients, and reducing wide-area telecommunication costs, and (2) in integrating the LAN, MAN (metropolitan area network), and WAN environments, thereby extending the full benefits of high-speed transmission throughout the geographic domains, while supporting clients attached to existing "legacy" LANs.
- In the multimedia model, the network is required to support high-function workstations that run applications involving combinations of voice, video, image, and data. Not only do high volumes of data transit the network, but delay-sensitive voice and video traffic must also be accommodated. In this case, ATM provides high-bandwidth on demand for both data and isochronous traffic *on the same physical links*, thereby providing the transmission performance to make multimedia applications technically feasible and the cost savings to make them economically viable.

How will ATM communications emerge?

ATM has proven to be unique in the sense that it enjoys the support of both the data processing community and the telecommunications carriers. The promise of ATM is great, but how will that promise be fulfilled? We believe that commercial ATM offerings will appear in four waves.

The first wave will see the appearance of high-speed ATM backbone networks. On the campus, ATM LANs will be organized around ATM switches interconnected to form high-speed backbones as shown in Figure 2. The presence of these ATM backbones, because of their isochronous capabilities, may well encourage the growth of high-speed, high-volume, multimedia applications for use within the confines of the campus.

Figure 2 ATM LANs

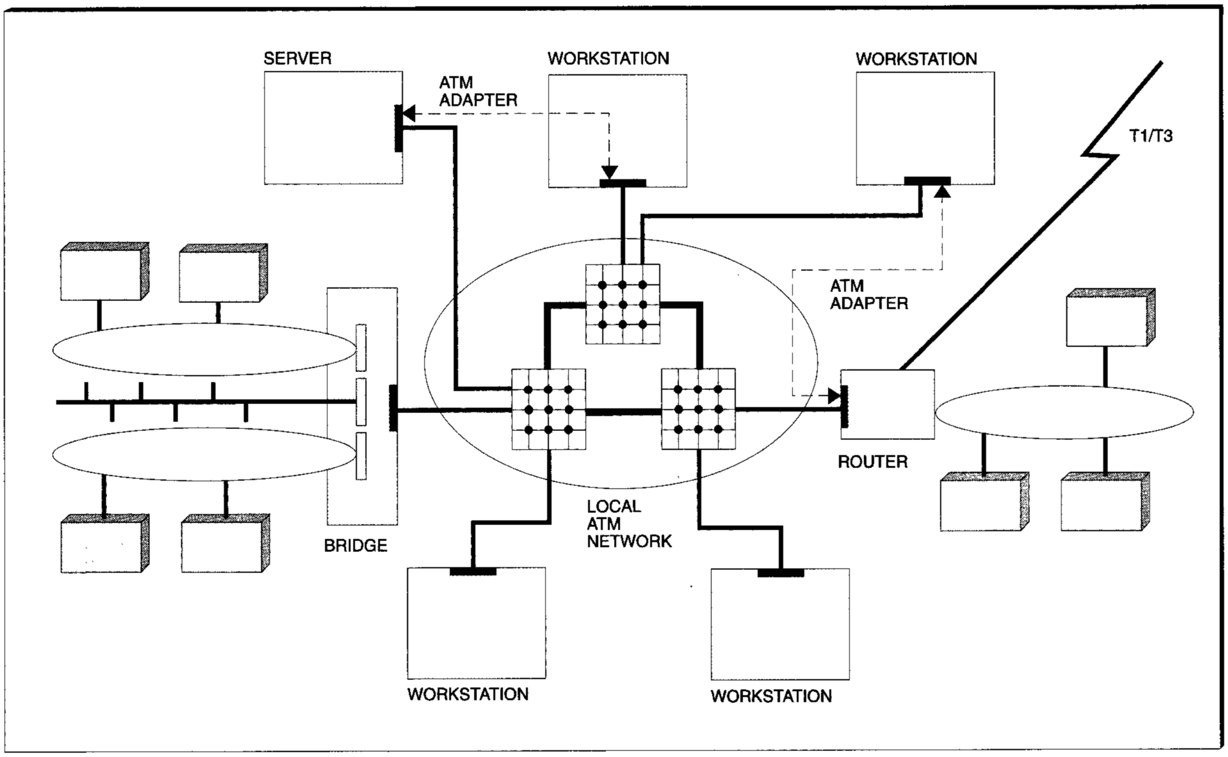

The rationale for this early campus emergence lies in the need to transcend the boundaries of the shared media backbones of today. Even the Fiber Distributed Data Interface (FDDI), with its 100 megabits per second (Mbps) capacity, is proving to be inadequate in some of the large and growing campuses. Universities, in particular, are starting to feel constrained by their existing technology—the bandwidth simply is not sufficient. Obviously, the new ATM backbone presents the opportunity to achieve greater performance and efficiency by the direct attachment of ATM-enabled servers and workstations. Existing devices will connect to this backbone via bridges and routers designed for that purpose, and mechanisms such as "LAN emulation" and "Internet Protocol (IP) over ATM" ("classic IP") will enable existing applications to operate in this new environment. Figure 2 shows this physical configuration.

The initial interconnection of these campus backbones across private wide area networks will be made using routers and leased lines, as shown in Figure 3. This pre-ATM wide area network will be generally limited in its speed and cost effectiveness. Moreover, sophisticated bandwidth-management capabilities lacking in this environment will be required to support isochronous connections over wide-area links shared by multiple users. So although it may be relatively easy to accommodate an interactive multimedia connection between two ATM workstations located on the same campus, the problem becomes more complex when the two workstations are separated by a wide area network. These complexities and limitations set the stage for the second wave.

The second wave will introduce commercial ATM technology to the private wide area networks that interconnect ATM LANs, as shown in Figure 4. These second-wave networks will be built around wide area network nodes (WANNs) that offer the full benefits of ATM: isochronous connections, fully managed bandwidth, and reduced transmission costs. The first evidence of this wave began around year-end 1994, and we expect to see a significant presence emerging in 1996.

Figure 3 Limited wide-area connectivity

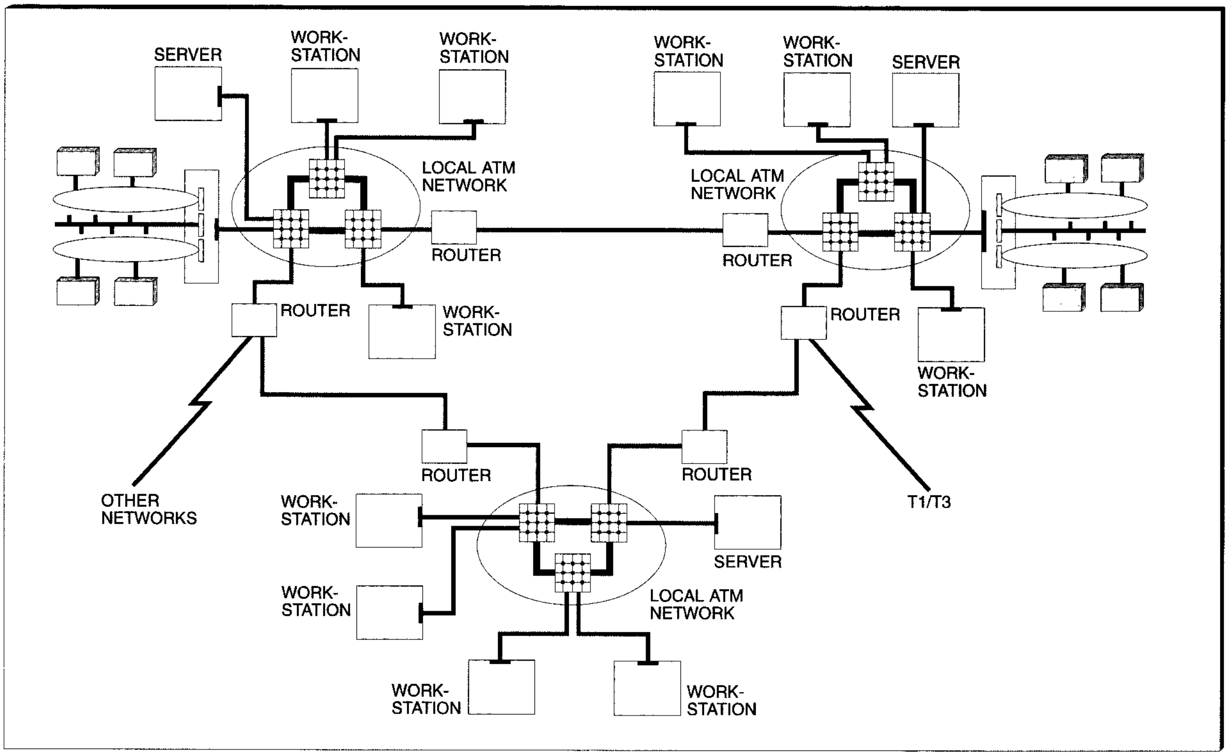

The third and fourth waves of commercial ATM technology will be marked by the coming of a pervasive public ATM service and ubiquitous ATM to the desktop. Some networking consultants have suggested that ATM ambitions will not be realized before the end of the century; however, much will depend on affordability and the changing paradigms in communications. Issues include adapter cost, ATM tariffs, and customer willingness to entrust data to the carrier. It should be noted that field trials and customer-specific contracts have already been initiated by a few of the carriers and "value-added" networks (those with services beyond basic transmission facilities). As these final waves swell, the complete convergence of the local area and wide area network seems within reach, although technical questions concerning directory services, data integrity, the consistency of addressing plans and physical interfaces, quality-of-service guarantees, and the substantial questions of finance and economics, must be resolved before the unified LAN-MAN-WAN infrastructure, optimistically shown in Figure 5, can become a reality. Nevertheless, we believe that many of the technical questions will be addressed successfully by the ATM standards bodies.

The IBM ATM strategy

In brief, IBM's strategy rests on the following objectives:

- To develop cost-effective and leading-edge technologies that will result in best-of-breed ATM products
- To provide a wide range of ATM solutions for the desktop, the campus, and the wide-area environment, along with a full network-management capability, according to a delivery schedule that satisfies customers
- To offer an ATM-based networking system that preserves and enhances customers' pre-ATM investments in wiring and applications, yet enables their networks to evolve and grow
- To provide a new high-speed communication architecture that provides functional and economic advantages to the customer by exploiting ATM capabilities

Figure 4 Private wide-area ATM

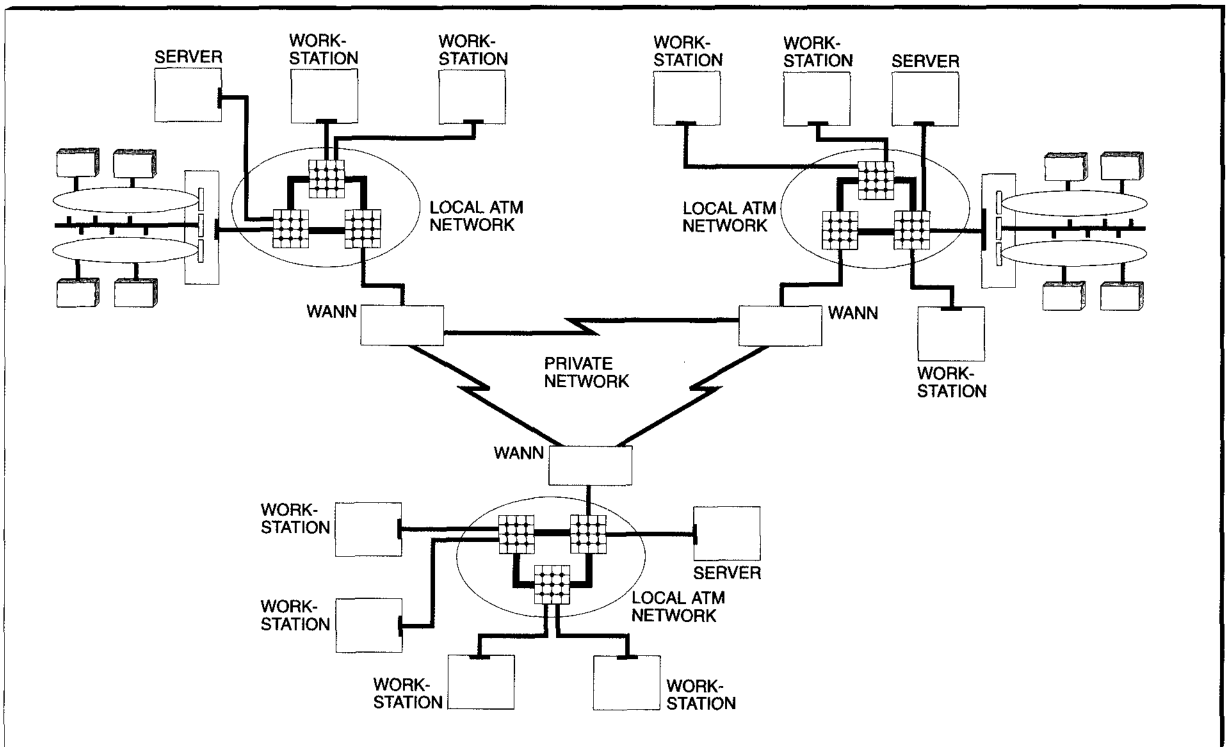

IBM is a principal member of the ATM Forum and, as such, has been participating in the development of the emerging ATM standards. IBM is committed to compliance with these standards while also building on its leading-edge technology and store of networking expertise to offer products that differentiate themselves in areas that lie beyond the scope of the standards.

IBM's ATM system and its family of products includes:

- A family of high-speed workstation adapters for clients and servers that, over time, will provide ATM communications over a range of transmission speeds, with a choice of copper or fiber transmission media
- An affordable 25 Mbps ATM adapter for the client workstation that can carry ATM traffic over voice-grade wiring that is prevalent on the customer's premises
- A 25 Mbps concentrator that aggregates multiple 25 Mbps ATM links into a higher speed (100–155 Mbps) ATM link for more efficient switch port utilization. This product acts as a multiplexor, as opposed to a switch.
- A switch-based ATM LAN built on the IBM 8260 Intelligent Hub. The 8260 provides LAN (token ring, Ethernet, and FDDI) concentration and management function. This capability is enhanced by the addition of an optional ATM backplane and ATM switch, and ATM concentration modules to enable ATM communication between hubs and among ATM-attached devices at a variety of standard speeds.
- An ATM to LAN (token ring or Ethernet) bridge that will enable devices on existing LANs to communicate across an ATM backbone with either ATM devices or other LAN devices
- An ATM workgroup switch that will support switching of multiple 25 Mbps ATM devices, and provide a high-speed (155 Mbps) ATM link for connection to an ATM backbone
- A LAN (token ring and Ethernet) switch that will have an ATM link into an ATM campus backbone network

Figure 5 Public wide-area ATM connectivity

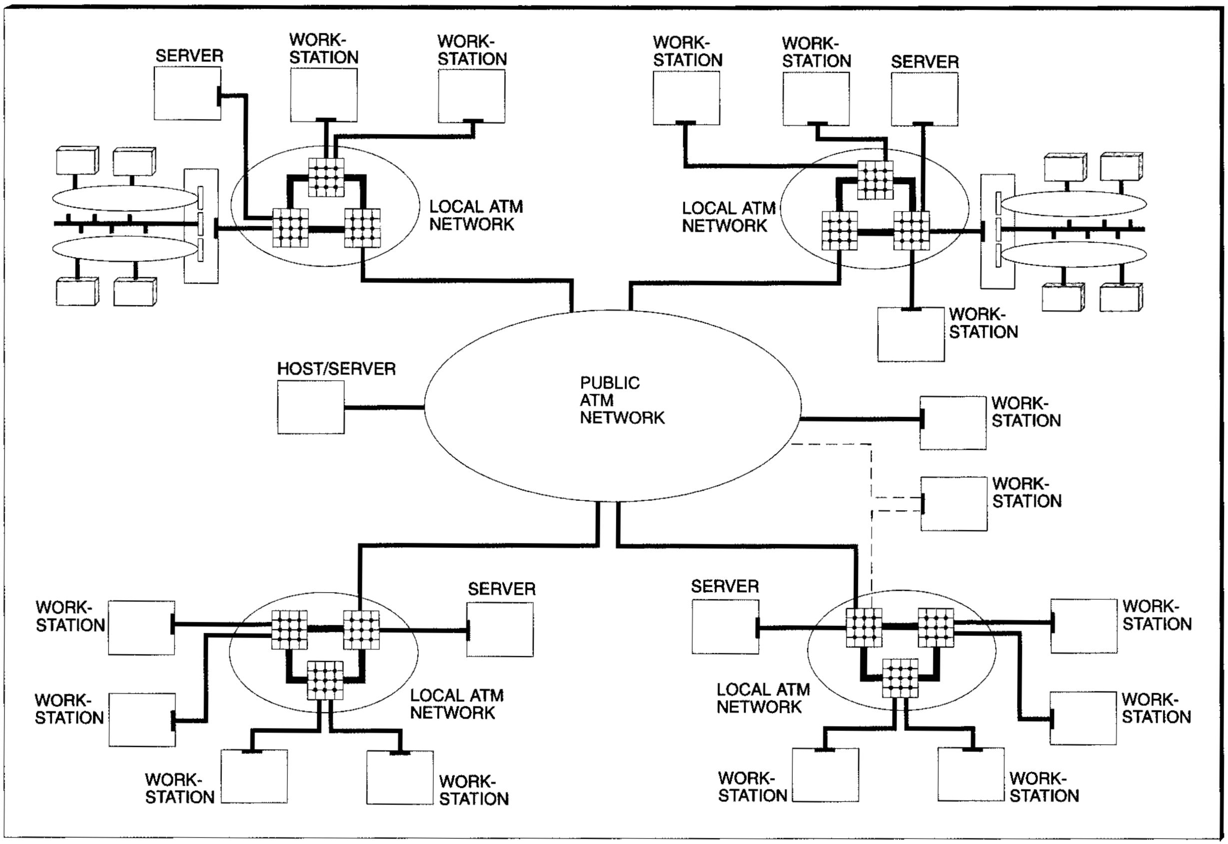

- LAN emulation software to enable existing LANs and applications to take advantage of the enhanced performance provided by an ATM campus backbone network
- A family of wide-area ATM switches that brings the benefits of IBM's new high-speed networking architecture to both private and public ATM networks. Models of these switches will range in size and capacity from a small campus access node to large gigabit switches appropriate to public carrier networks.

ATM, IBM, and worldwide telecommunication standards

Because the ultimate objective of the ATM effort in the broader community is to enable the convergence of local area, metropolitan area, and wide area subnetworks while maintaining openness to equipment made by a wide selection of vendors, ATM has become the focus of cooperative organizations devoted to telecommunication standards. IBM continues its commitment to openness and convergence by virtue of its active participation in these efforts.

Committees that develop standards for the various aspects of ATM function mainly under the auspices of two groups: the International Telecommunications Union-Telecommunication Sector (ITU-T), formerly the International Telegraph and Telephone Consultative Committee (CCITT), and the ATM Forum. The ITU-T represents the interests of the public telecommunication carriers, whereas the ATM Forum represents the interests of the community of vendors that also sell private-network equipment.

Some of IBM's contributions to the ATM Forum target the carryover of pre-ATM technology and components to provide a cost-effective base for ATM networking. For example, some IBM ATM Forum efforts include:

- Leadership of the effort to base ATM Forum traffic management on rate-based rather than credit-based flow control
- Initiation of the work effort for a low-cost ATM user-to-network interface (UNI) physical layer on existing copper media (untwisted shielded pair, category 3 [UTP-3])
- Initiation and leadership of the work effort for end-to-end network management of private and public ATM networks and services
- Initiation and leadership of the work effort to focus on LAN and WAN migration over ATM networks, e.g., LAN emulation and frame relay emulation
- Initiation and leadership of the work effort to focus on definition of native ATM application programming interfaces (APIs)
- Initiation of the work effort on end-to-end quality-of-service guarantee
- Leadership of the work effort on private network-to-network interface (PNNI) source routing and establishing link state metrics
- Initiation of cell delay variation tolerance studies

In fact, IBM has contributed over 10 percent of the submissions to the forum since its inception.

IBM's communication services for ATM

Although two independent organizations are concerned with ATM standards—the ITU-T and the ATM Forum—they share a common goal: to provide seamless local and wide area networking, i.e., LAN = MAN = WAN. In line with this common goal, ATM standards have focused on the interfaces that enable end system attachment, as well as those that describe communications between switches, while not dictating individual switch architecture. This approach of conscious demarcation empowers the community of ATM vendors to design products that interoperate, yet maintain distinct product identity. It does not dictate how these products are implemented.

To distinguish its products in a way that provides the best value to customers, IBM has developed new services for high-speed networking called Networking BroadBand Services, or NBBS. These services provide:

- Distributed network control that offers path selection to eliminate bottlenecks for improved availability
- Integral topology and directory services that maintain and retrieve information on network resources and users
- Hardware-based spanning-tree routing that rapidly distributes connection-control information. This routing technique guarantees that control information will be sent to each node over a single link, thus eliminating redundant broadcasts and potential broadcast loops.
- Dynamic bandwidth allocation that optimizes the use of bandwidth, provides bandwidth on demand, readily accommodates isochronous traffic, and provides multicast capability
- Guaranteed quality-of-service (QOS) measures that ensure appropriate end-to-end bandwidth, delay, and packet-loss characteristics
- Automatic routing that provides nondisruptive rerouting of connections around network-link failures
- Congestion control that focuses on avoidance of, rather than reaction to, network congestion

The technical capabilities provide mechanisms that deliver a substantial economic advantage by saving customers up to 50 percent or more of the recurring cost of wide-area bandwidth, depending on the traffic, or conversely, by supporting more connections across a finite-capacity network link.

Prototype networks based on NBBS have been implemented as testbeds for high-speed applications. These have included both research-oriented and commercial testbeds.

It is important to note that NBBS is not "IBM's version of ATM." ATM provides a set of standards on which a working network can be constructed. NBBS defines the working network by exploiting the capabilities of ATM and adding additional control and management function. IBM products that implement NBBS will fully support the ATM standards as they emerge.

Other papers in this issue provide an in-depth description of NBBS.

IBM's switch fabric—"switch-on-a-chip"

The advances in very large scale integrated (VLSI) circuit technology will allow the development of high-speed packet switches able to accommodate the future transmission rates. As opposed to conventional packet switches, which are mainly based on software processing, these new advanced packet switches are hardware-based. Their topologies and queuing arrangements for contention resolution employ a high degree of parallelism. Their routing function, referred to as self-routing, is typically performed in a distributed way at the hardware level. In fact, many different designs for advanced packet-switching fabrics have been proposed and developed over the last few years, but we feel that the "switch-on-a-chip," developed by the IBM Research Center in Zurich, represents the state of the art.

Utilizing NBBS and cost-effective technology, IBM is providing a broad range of ATM products.

This single-chip switch element exploits the performance advantage of output queuing and can be used to construct larger, self-routing single-stage or multistage switch fabrics in a modular way. High performance is achieved by output queues that are configured as a dynamically shared memory. This shared memory can be expanded by linking multiple switch elements. This design allows for transmission rates on the order of a gigabit per second per port.

Switch-on-a-chip is a self-routing, nonblocking, scalable switch fabric particularly well-suited for transmission of multimedia (voice, video, and data) traffic. It is employed in both local- and wide-area IBM ATM switches.

IBM's comprehensive ATM product line

Utilizing NBBS, switch-on-a-chip, and cost-effective ATM technology, IBM is providing a broad range of ATM products that span both the local- and wide-area environments and both the public and private networking domains. Individually, they will provide standards-based network components capable of interoperating with any other vendor products that support the ATM standards. Together, these products will provide a complete *ATM system*, rather than just a set of isolated ATM products.

In June 1994, IBM announced a family of hardware and software products aimed at the distributed, multivendor networking environment. This family, called the *Nways** product family, includes both existing products (e.g., adapters, routers, wireless products, TCP/IP) and the more recently announced and previewed ATM products. Within the context of the IBM Open Blueprint*, Nways products fall into the transport network and subnetworking layers.

The Nways/ATM announcement illustrated IBM's intent to deliver the elements of its ATM system according to a schedule that is consistent with anticipated customer needs and in a way that affords a smooth migration from old technology to new. The system will enable the coexistence and interoperability of the new ATM products with today's base of LANs, workstations, and other networking products, thereby preserving the customer's prior investments. Preserving investments in older technology, however, is not the same as maintaining the status quo; instead, it implies an obligation to provide alternatives for the evolution of networks.

Accordingly, IBM's ATM product rollout will address the difficult questions that surround the issue of migration while at the same time acknowledging commonly held industry views. For example, IBM subscribes to the commonly held view that frame-relay communication has some limitations that relegate its long-term role to that of a stepping stone to ATM. Nevertheless, frame relay shares several attributes with ATM: both employ fast-packet techniques to achieve efficient bandwidth utilization, and both target LAN interconnection. Consequently, frame relay can be seen as addressing some of the same problems as ATM. Because of this commonality between frame relay and ATM, IBM's total system solution will present the interfaces needed to support frame-relay communications in an ATM environment, thus offering customers an evolutionary alternative to ATM-only or frame-relay-only approaches. Similarly, other interfaces will be provided that continue support for other protocols—high-level data link control (HDLC), circuit emulation, etc.

Summary—IBM's total system solution for ATM

Since the June 1994 announcement, IBM has received considerable acclaim from both the trade press and consultants for its ATM strategy and products. The scope of its product offerings and the value of its technology have positioned IBM to assume a leadership role in the ATM marketplace.

IBM's experience in high-speed networking provides a solid foundation for developing leading-edge network-control and optimization features that both comply with industry standards and go beyond them in providing functional value to the customer. The focus on cost effectiveness in designing ATM components, on preserving the customer's pre-ATM wiring and applications, and on economically controlling bandwidth provide good financial reasons for customers to welcome IBM's offerings. Furthermore, the product delivery schedule is designed to allow full compliance with emerging ATM standards, a smooth migration from old technology to new, and relevance to customers' evolving needs. As noted earlier, IBM has already made some initial announcements and previewed additional products. Availability for these products started in the second half of 1994 and will continue into the future based on customer requirements. Important to note is the fact that IBM is capitalizing on its technology (switch-on-a-chip, NBBS) and its experience in data networking to address the ATM service provider market (e.g., carriers, telephone companies).

IBM will deliver standards-compliant ATM products that range from workstation adapters to full-function WAN switches in order to provide customers with the flexibility to implement any combination of private and public networking services and to manage and control their networks effectively. By providing ATM switching in its key LAN and WAN products, IBM will ensure synchronized capabilities, and ensure the coordination of new products with established products. As a result, disruption to customers' networks will be minimized as they migrate to ATM technology; moreover, the value of the customers' prior investments in networking equipment will be preserved to the greatest possible extent.

This is not to say that IBM is overlooking leading-edge implementations of ATM. In particular, several industries (e.g., education, medical, and entertainment) are already putting in place plans to design and build applications that will depend on ATM networks. Distance learning, remote diagnosis, and video on demand all require high bandwidth and, in some cases, multimedia implementations. At this time, the ATM Forum is working to define the standards for native APIs that will enable applications to take full advantage of ATM. IBM is helping to define these standards and will provide this capability in its products.

In summary, IBM is committed to ATM. Initiatives are under way to penetrate new markets (e.g., telecommunications carriers) and to develop industry-oriented solutions. Customer and vendor partnerships are being pursued as needed in order to achieve these ends and realize the goal of industry leadership.

*Trademark or registered trademark of International Business Machines Corporation.

Accepted for publication August 15, 1995.

M. Elizabeth Cohen *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: bcohen@vnet.ibm.com).* Ms. Cohen joined IBM in 1967 in the Thomas J. Watson Research Center in Yorktown Heights, New York. She held several management positions in the Corporate Information Systems organization before joining the Networking Hardware Division in 1986. She joined the ATM System Management function in 1993 and helped to define and shape the ATM strategy for the division. Currently she is a program manager in the ATM Market Development organization. Ms. Cohen holds a B.S. degree in mathematics and economics from Southern Colorado State College and an M.S. degree in operations research from New York University.

Daniel S. Abensour *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: abensour@ralvm29.vnet.ibm.com).* Dr. Abensour is the program director of the ATM Market Development organization and a senior technical staff member in the Advanced Transport Business segment. Since joining IBM in Paris in 1966, he has held several senior management positions in the Thomas J. Watson Research Center in Yorktown Heights, New York, in La Gaude, France, and in Boca Raton, Florida. His current responsibilities include the definition and development of industry solutions based on ATM, the development of the ATM technical marketing strategy, and acting as spokesman for IBM's ATM strategy. Dr. Abensour is a member of the IBM Academy of Technology and is a senior member of the Institute of Electrical and Electronics Engineers (IEEE). He graduated from the École Nationale Supérieure des Télécommunications in Paris in 1964 and holds a *doctorate es Sciences* from the University of Toulouse, France.

Reprint Order No. G321-5581.