

Enabling ATM networks

by R. O. Onvural
H. J. Sandick
G. M. Kump
E. A. Hervatic
P. F. Chimento

Asynchronous transfer mode (ATM) is the primary candidate technology for emerging high-speed networks. ATM standardization provides the basic framework for building ATM networks. Although the development of ATM standards and the availability of high-speed links in the network is essential, what really matters is how ATM networks are enabled so as to integrate different services with different traffic characteristics and service requirements. Two important contributions to this integration are the development of network control services complementing ATM standards while achieving high utilization of network resources and the efficient implementation of the ATM standards in the network. In this paper, we present a short overview of the current status of ATM standards, discuss how various control services that were developed in IBM's Networking BroadBand Services architecture complement the ATM standards, and provide an overview of various functions included in the architecture for their efficient implementation.

In the early 1980s, the telecommunications industry started to develop the concept of the Broadband Integrated Services Digital Network (B-ISDN). B-ISDN was conceived as an all-purpose digital network that integrates different types of services with different quality-of-service requirements. The technology chosen to deliver this promise is referred to as *asynchronous transfer mode* (ATM). The International Telecommunication Union-Telecommunications sector (ITU-T) published its initial set of B-ISDN recommendations in 1988, including the specification of ATM as the transfer mode of choice for B-ISDN. Strictly speaking, B-ISDN is a public service, and the B-ISDN architecture has been developed for public networks. However, ATM is a technology that can be used in

local, campus, and private wide area networks as well. To a very great extent, ATM and B-ISDN have been used interchangeably in the literature. In this paper, we use the term ATM exclusively.

Simply defined, ATM is a connection-oriented packet-switching technology that uses fixed-size packets, referred to as *cells*, to carry the traffic in the network. ATM embodies various design objectives that include:^{1,2}

- Integration of voice, video, image, and data services in a single framework
- Minimization of the switching complexity and the processing burden at the switches
- Minimization of the buffer management complexity and the buffering needed at the switches

ITU-T decided that these design objectives can be met at high transmission speeds by keeping the fundamental unit of transmission, the ATM cell, short and fixed in length.¹ This constraint in turn provides the flexibility to support a wide range of services that emerging high-bandwidth, multimedia applications require. It also fosters statistical multiplexing. That is, the total bandwidth needed by an aggregate traffic mix may be less than the sum of the maximum bit rates generated by the indi-

©Copyright 1995 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *re-publish* any other portion of this paper must be obtained from the Editor.

vidual traffic streams, provided that the traffic generated by each source fluctuates over time.

Characterizing ATM as asynchronous indicates that cells may occur at irregular times determined by the nature of the application rather than the framing structure of the transmission system. The connection-oriented nature of ATM arose out of the need to reserve resources in the network to guarantee that the service quality provided to either real-time or loss-sensitive applications or to both does not degrade because of traffic integration in the network.

ATM is envisioned as a technology that will facilitate universal networking: It integrates voice, video, and data, and offers scalability both in distance (a single technology in the local area, campus, and wide area) and in speed (currently defined physical layer interfaces vary from 1.5 Mbps up to 622 Mbps). The types of services available in ATM networks include:³

- ◆ Bandwidth on demand
- ◆ Guaranteed service levels
- ◆ Point-to-point and point-to-multipoint connections
- ◆ Constant as well as variable bit rate services
- ◆ Connection-oriented or connectionless application services

ATM supplies the basic framework to assign usable capacity dynamically on demand. The basic operation of an ATM switch is independent of what is carried in the ATM cell payload, thus having the capability of switching different types of traffic with different requirements.

Figure 1 illustrates an ATM network and the protocol stack at the ATM end stations for user traffic. An ATM network provides ATM layer connectivity among the end stations. Users access the network through standard interfaces.

The ATM layer is the source of the transmission, multiplexing, and switching functions for the end-to-end, in-sequence transmission of ATM cells. There is no awareness at the ATM layer of the contents of the payload. It is also independent of the physical layer. The ATM layer functions are kept very simple mainly to keep up with the high-speed transmission links in the network. For example, a cell is required to be processed in less than three microseconds to keep up with a 155 Mbps link. This

time reduces to less than one microsecond with a 622 Mbps link. Hence, there is not much time to process much of anything for a cell in transit at the intermediate switches.


Because the ATM layer mainly serves to route the cells in the network, more features than those defined at this layer are required to support network applications in ATM networks. Considering various services that are envisioned in ATM networks, ITU-T grouped functions that are common to most applications at the ATM adaptation layer (AAL). The main function of the AAL makes available a generalized interworking across ATM networks by providing a service interface to the ATM layer. Therefore, AALs are application-specific.

Enabling ATM networks requires a network infrastructure that implements the ATM architecture as defined in the standards while providing a comprehensive set of network services that complements the ATM standards. Networking BroadBand Services (NBBS) is IBM's comprehensive architecture that meets this challenge while achieving high utilization of network resources.⁴ Various control services developed in the NBBS architecture are reviewed in other papers in this issue.⁵⁻⁹ In this paper, we briefly give an overview of the current status of various ATM standards, discuss the need for network control services that complement ATM networks, and discuss how NBBS provides these services in a standards-compliant manner. In particular, the current status of ATM standards and specifications is presented next. We then describe what are some of the network control services needed to deploy ATM networks and how NBBS meets this challenge by using an ATM call setup as an example. The last section is a brief overview of various features developed in NBBS to implement ATM standards in the network efficiently.

Standards on ATM interfaces

Standardization is necessary to ensure interoperability. Without standards, interfaces become proprietary and limit connectivity choices. Today, customers view interoperability among networking equipment from different vendors as essential. In this environment, network equipment vendors develop standards-compliant products and services while differentiating their products based on cost, reliability, and the value-added features built into them.

Figure 1 Layering in an ATM network and ATM end stations




To provide a common framework for emerging ATM networks and ATM end stations, various standards organizations are working to address the challenges of emerging high-speed, multimedia networks. The ATM standards have been developed by the ITU-T (formerly known as the CCITT) with contributions from various national standardization organizations such as the American National Standards Institute (ANSI) and the European Telecommunications Standardization Institute (ETSI).

Another major contributor to solving the interoperability problems in ATM is the ATM Forum, a consortium of more than 700 companies. The ATM Forum members share a common desire to get ATM-based products to the marketplace as expeditiously as possible. The principal mission of the ATM Forum is to provide interoperability specifications in order to speed up the development and

use of ATM-based products. Because the early deployment of ATM products requires that specifications be available much earlier than the schedules of the formal standards bodies permit, the goal of the Forum is to anticipate and fill gaps in the international telecommunication standards.

The ATM standards and specifications are based on ATM interfaces. Figure 2 illustrates different interfaces that have been defined or are currently being defined by the aforementioned organizations. The ATM user-to-network interface (UNI)¹⁰ defines the demarcation between ATM users and the ATM network. In this context, an "ATM user" is a device that transmits ATM cells into the network. Accordingly, an ATM user may be an interworking unit that provides interoperability between a non-ATM network and an ATM network by encapsulating non-ATM data into ATM cells, or it may be a pri-

Figure 2 ATM interfaces



vate ATM switch or an ATM workstation. The UNI is referred to as *private UNI* or *public UNI*, depending on whether the ATM network is private or public. A private network can request service from a public network across a public UNI.

The UNI specifications define a set of physical-layer interfaces and the ATM layer. For connections that are established dynamically on demand, the UNI specification defines the signaling procedures used across a UNI for the establishment, termination, and management of ATM connections.

The ATM data exchange interface (DXI)¹¹ enables routers to interwork with ATM networks without requiring special hardware by connecting routers to an ATM data service unit (DSU). The DSU provides a mapping from a data link layer protocol defined as part of the DXI to the ATM protocol layers.

In particular, the ATM Forum's DXI specification, as the demarcation point between the data terminal equipment (DTE) and data communications equipment (DCE), defines a data-link protocol and the physical layer interfaces needed to transfer data between the router and the DSU. The specification also includes the definition of a network management interface.

The ATM broadband intercarrier interface (B-ICI) supports the joining of two public ATM networks.¹² B-ICI provides end-to-end connections across two or more public networks for a variety of telecommunication services. For this purpose, the Forum's B-ICI specification covers the physical layer, ATM layer, and interworking for non-ATM services such as frame relay, cell relay, circuit emulation, and switched multimegabit data services (SMDSs).

The ATM network-to-network interface (NNI) defines the demarcation between two ATM *switching systems*. A switching system in this context can be a single switch or an ATM network in one administration domain. The NNI is called the "private NNI" or the "public NNI," depending on whether the switching systems are private or public.

Private NNI (P-NNI) specification is currently being worked on by the ATM Forum. The specification includes P-NNI routing and P-NNI signaling. P-NNI routing is used to find a path across a network between two end stations (point-to-point connection) or more than two end stations (point-to-multipoint connection). The switching system that a connection request originates across its UNI is responsible for finding the end-to-end path to the destination end station and is referred to as source routing. In determining the path, the originating switching system uses link-state routing in which each switching system advertises information about its P-NNI links to other switching systems. In this context, a P-NNI link connects one switching system to another (in a given direction) across a P-NNI. After finding the path, the originating switching system uses P-NNI signaling to request the establishment of a connection from the intermediate switching systems along the path. The switching systems along the end-to-end path are specified by the switch at which the connection originates in the designated transit list stack and included in the corresponding P-NNI signaling message. Each switching system along the path processes a received connection request message, makes a connection admission decision (i.e., accept or reject the connection request), and passes the signaling message to the next switching system along the path (if accepted) or denies the connection. Connection admission decisions (as well as path selection) are vendor-specific solutions and are not subject to standardization. However, the source switching system should be able to predict the outcome of a possibly unknown admission procedure at a switch (or a switching system or both) with "some" confidence so that a large portion of connection setup requests under normal operating conditions are successful (i.e., result in connection establishment). A generic call admission control procedure is specified in P-NNI for this purpose.

The choice of what internal state information to advertise, how often, and to where requires the specification of a multilevel hierarchical routing

model. In P-NNI routing, the number of hierarchical levels is not preset and may vary from one level (i.e., no use of hierarchy) up to 104. A very large corporate network may have between two to four hierarchical levels of routing. The P-NNI hierarchical model explains how each level of hierarchy works, how multiple nodes at one level can be summarized into the higher layer, and how state information among nodes within the same level and between different levels is exchanged.

ITU-T standardization on public NNI includes specifications for physical and ATM layers, control, user, and management planes, and NNI signaling (B-ISDN user part, i.e., B-ISUP). Unlike P-NNI, however, NNI specification does not include any routing protocol. B-ISUP provides signaling functions required to support basic bearer services and supplementary services for B-ISDN applications. It was developed for international applications as a network-to-network interface. However, it is also designed to support national applications as a network-node interface.

The B-ISUP procedure and its operation are rather involved and are not discussed here. The elements of signaling messages and their functions used by B-ISUP are currently defined to support basic bearer services and supplementary services for capability set 1, which includes the following features:

- Point-to-point demand (switched virtual) channel connections
- Connections with symmetric or asymmetric bandwidth requirement
- Single connection per call
- Basic signaling functions and procedures
- Class X, class A, class C ATM transport services
- Request indication of signaling parameters
- Virtual channel identifier (VCI) negotiation
- Out-of-band signaling for all signaling messages
- Error recovery
- Public UNI addressing formats for unique identification of ATM endpoints
- End-to-end compatibility parameter identification
- Signaling interworking with ISDN and provision of ISDN services

UNI specification

The UNI specification defines standards-based physical and ATM layers and the UNI signaling pro-

Figure 3 ATM protocol reference model




Figure 3 illustrates the UNI protocol architecture.

Physical layer. The physical layer transports ATM cells between two ATM entities. It guarantees, within a certain probability, the cell header integrity and multiplexes user cells to generate a continuous bit stream across the physical medium. Different types of physical media specified for ATM include fiber optic cable, shielded twisted pair, and coax at various speeds ranging from 1.5 Mbps to 622 Mbps.

ATM layer. The ATM layer specifications define various functions and services that include:


- Services provided to the ATM layer user
- Services required by the physical layer
- ATM cell structure and encoding

- Traffic and congestion control at the ATM layer
- ATM layer management specification

The ATM layer provides for the transparent transfer of fixed-size (48 bytes) *service data units* (ATM-SDUs) between ATM layer users. It generates a five-byte ATM cell header in the transmit direction and extracts ATM cell headers in the receive direction. Multiplexing of cells is supported by placing cells from individual connections into a single cell flow. The corresponding function in the receiving direction is the demultiplexing of the cell stream. The ATM cell format is illustrated in Figure 4.

ATM layer connections are either pre-established using internal network procedures or are set up as needed on demand. The former are referred to as *permanent virtual connections* (PVCs), whereas the

Figure 4 The ATM cell format




latter are referred to as *switched virtual connections* (SVCs).

An ATM connection is identified by the virtual channel identifier (VCI) and the virtual path identifier (VPI) fields at the cell header. VCI values are unique only in a particular VPI value, and VPI values are unique only in a particular physical link. Furthermore, the VPI/VCI has local significance only, and the VPI/VCI label is translated at every switch the cell traverses. This technique is known as label swapping. Accordingly, a connection within an ATM switch is a simple mapping that translates incoming labels to the outgoing labels.

An ATM routing framework defines a two-level-routing hierarchy: virtual path (VP) and virtual circuit (VC). A VP is a logical semipermanent connec-

tion established by using network management or control functions, or both, over a set of physical links in sequence. That is, routing table entries that map an incoming (physical link identifier, VPI) pair to a particular outgoing (physical link identifier, VPI) pair are predefined, and they are only changed when the network is reprovisioned. Note that this provisioning does not include VCIs. The use of VPs simplifies network control and management functions. A VP in the network can accommodate up to 65 536 (2^{16}) VC connections; that is, only one table entry is required to switch up to 65 536 individual connections (assuming that they are all following the same set of physical links in the same sequence). Based on this framework, the VP concept is used to simplify the routing table management, connection setup process, path selection, and bandwidth reservation.

Figure 5 The VCC and VPC hierarchy in ATM connections



In summary, a VP is a collection of VCs between two nodes in an ATM network. A predefined route is associated with each VP in the physical network (i.e., VPs are semipermanent connections), and routing tables for VP switching are preset in ATM networks using, for example, network management functions.

VPIs are used to route packets between two nodes that originate, remove, or terminate the VPs, whereas VCIs (unique only within a VP) are used at the VP endpoints to distinguish between different connections. The VPI/VCI pair used at a switching node has a local meaning only (even the VCI does not change within a VPC; the VPI/VCI is translated as the VPI is translated at every switch).

Switching from incoming links to outgoing links is done by reading the routing field(s) of incoming cells, performing a table look-up to determine the outgoing link, placing the new routing identifiers used on the outgoing link, and delivering the cell to the corresponding outgoing link port with the new header.

An end-to-end connection between two end stations is referred to as a *virtual channel connection* (VCC). As illustrated in Figure 5, a VCC is composed of one or more VC links. A VC link consists of one or more physical links between a point where a VCI is assigned and the point where it is removed or translated or both. A VC link consists of one or more VP links in which a VP link is defined between

a point where a VPI is assigned, to the point where it is removed or translated or both. A virtual path connection (VPC) is a concatenation of one or more VP links.

A point-to-point VCC is a bidirectional connection consisting of a collection of VC or VP links connecting two endpoints that wish to communicate with each other. Similarly, a point-to-multipoint VCC is a unidirectional connection consisting of a collection of VC or VP links that connect one sender (root) to two or more receivers (leaves).

In general, the operations, administration, and maintenance (OAM) framework includes (1) configuration, (2) fault, (3) performance management, (4) security, and (5) billing functions. The current ATM layer OAM specification includes only the performance and fault management functions.

In particular, at the ATM layer, network resources are monitored for equipment faults and performance degradation. If a resource in the network is determined to be operating at a nondesired level, maintenance actions are taken to diagnose the cause and repair it. Operations activities involve continuous coordination of administration and maintenance of network resources based on feedback collected from network resources and user requirements.

The five OAM functions currently specified are:

- Performance monitoring: Normal functioning of the managed entity is monitored by continuous or periodic checking of functions, and maintenance event information is maintained.
- Defect and failure detection: Malfunctions are detected by continuous or periodic checking, and maintenance event information or various alarms are produced.
- System protection: Effect of failure of a managed entity is minimized by blocking or changeover to other entities, and the failed entity is excluded from operation.
- Failure or performance information: Failure information, alarm indications, and response to requests are given to corresponding management entities and planes.
- Fault localization: Determination by internal or external test systems of a failed entity for failure of information is insufficient.

The ATM layer management specification defines alarm surveillance and connectivity verification for

VPs and VCs. These functions are performed by using special management cells, known as *operation and maintenance* (OAM) cells. OAM functions are broken down into hierarchical levels associated with the physical and the ATM layers. The flows that result from the execution of OAM functions are referred to as *OAM flows* and are named F1, F2, F3, F4, and F5. F1, F2, and F3 are physical-layer flows, whereas the ATM-layer flows are respectively associated with the virtual-path-level (F4) and virtual-channel-level (F5) flows. F4 flows are terminated at the endpoints of a virtual path connection or at the connecting points terminating a VPC segment. The same applies to VCC F5 flows. Various functions provided by the F4 and F5 flows include the following:

1. Performance monitoring
2. Defect and failure detection
3. Failure or performance information
4. Fault localization

Fault management functions are used for identifying and reporting resource failures. Similarly, performance monitoring functions are designed to capture error conditions and troubles resulting from gradual deterioration of network equipment.

ATM adaptation layer. An ATM network provides an end-to-end ATM layer connectivity among end stations with the ATM layer dealing only with the functions of the cell header, independent of the type of information carried in the payload. This simplicity is needed to keep up with high-speed transmission links, and it is achieved by leaving out various services required by applications sent to the end stations. At the ATM layer, in particular, there is no:

- Information on the frequency of the service clock
- Detection for misinserted cells
- Detection for lost cells
- Means to determine and handle cell delay variation

The main reason for not providing these functions inside the network (i.e., at the ATM layer) is that not all of these services are required by every application. For example, data traffic does not require any information on the frequency of the service clock, whereas voice may not require any awareness of possible bit errors.

Architecturally, the ATM adaptation layer (AAL) is between the ATM layer and the next higher layer

Table 1 Four classes of applications defined by ITU-T

Class	Timing Relationship	Bit Rate	Connection Nature	AAL
A	Yes	Constant	Connection-oriented	AAL 1
B	Yes	Variable	Connection-oriented	AAL 2
C	No	Variable	Connection-oriented	AAL 3/4 and AAL 5
D	No	Variable	Connectionless	AAL 3/4 and AAL 5

in both the user plane and the control plane. AAL enhances the services provided by the ATM layer to support the functions required by the next higher layer. Accordingly, AAL is service-dependent. It isolates the higher layers from the specific characteristics of the ATM layer by mapping the higher-layer protocol data units into the ATM cell payload and vice versa. It is not feasible to address the requirements of all of the applications served by ATM networks either individually or in a single AAL framework. Instead, the functions required by various applications are grouped into a small number of classes based on the commonality of their service requirements and traffic characteristics. A different AAL is then defined for each class of service.

To meet this objective, ITU-T classified B-ISDN applications based on three parameters:

- Timing relationship between end stations (required or not)
- Constant bit rate (CBR) or variable bit rate (VBR)
- Connection-oriented or connectionless

Real-time services are governed by timing relationships. For example, pulse code modulated (PCM) voice requires playing out the voice frames at the called party at the same rate as they are generated at the calling party. This in turn requires the service clocks at the two end stations to be aligned. Nonreal-time services such as file transfer do not need such a timing relationship.

CBR applications may require continuous availability of a constant bit rate along their end-to-end paths. PCM voice or CBR video are examples of such applications. VBR applications such as VBR video and local area network interconnection traffic generate bit streams to the network at a varying rate.

Higher-layer protocols at ATM end stations may or may not require connections to be established between the peer protocols. For example, frame re-

lay is a connection-oriented protocol, whereas Internet Protocol is connectionless.

This framework gives rise to eight combinations of application classes, four of which are explicitly defined by ITU-T: class A, class B, class C, and class D. Based on these service classes, AALs 1, 3/4, and 5 are defined as illustrated in Table 1. Note that AAL 2 has not yet been specified.


UNI signaling. The UNI signaling architecture is a layer-3 protocol used to establish, maintain, and terminate demand-switched virtual channel connections. Figure 6 presents the B-ISDN signaling structure used at the ATM UNI. The user-network signaling protocol runs on top of the signaling ATM adaptation layer, which defines how to transfer the signaling information for call or connection control reliably using the cells of the ATM layer on a dedicated point-to-point signaling virtual channel (VCI = 5 and VPI = 0).

Signaling messages containing *information elements* (IEs) are exchanged between an end station and the network across the UNI for switched virtual circuit (SVC) connections. The IEs contain the information required by the network to provide a sequence-preserving cell transfer service with an agreed quality of service. IEs also contain information used to negotiate connection parameters among ATM endpoints. SVC connections are classified as point-to-point and point-to-multipoint.

Each ATM end station has a unique ATM address assigned to it. These addresses are used in signaling messages for the network to locate the destination end nodes.

The ATM Forum private ATM address is 20 bytes long. The first 13 bytes are referred to as the prefix and identify the switch to which the end station is

Figure 6 B-ISDN signaling structure at the ATM UNI



attached. The next 6 bytes, referred to as the end system identifier, uniquely identify a particular end station among all that are attached to the switch. The last byte is the selector field reserved for use by the end station.

It is necessary for both the end station and the network to know the full ATM address (i.e., first 19 bytes) that uniquely identifies the end station. ATM Forum UNI specification version 3.1 includes a client registration mechanism for the exchange of address information between an ATM end station and an ATM switch port across a UNI. This mechanism is defined within the interim local management interface (ILMI) framework. In addition to the client registration mechanism, ILMI provides the status, configuration, and control information about the link and physical layer parameters connecting the two sides of the UNI.

In its connection setup request, the ATM layer user requests ATM layer *quality-of-service* (QOS) class. The acceptance of the request indicates the com-

mitment from the network that the service requirements of the connection will be met for the duration of the connection as long as the traffic generated by the application conforms with the parameters agreed on at connection establishment time. Currently, the service requirement of an application is specified as a QOS class that can be either "specified" or "unspecified." A *specified* QOS class includes a set of performance parameter values defined by the network, which may include cell loss ratio, end-to-end delay, and delay variation. The current UNI framework requires one or more service profiles to be specified for each service class. These service profiles are provisioned (pre-configured and cannot be changed dynamically) and cannot be signaled explicitly. Furthermore, the parameters and their values are determined by switch vendors and agreed on between ATM end stations and the network at the time the end stations are attached to the network (i.e., provisioned at the switch, depending on the internal switch design). The new releases of UNI specifications are expected to address this problem by allowing end stations to explicitly signal the parameters and their values required for their applications.

The traffic characteristics of connections are signaled by using the connection traffic descriptor, which includes up to four parameters: peak cell rate (PCR), cell delay variation tolerance (CDVT), sustainable cell rate (SCR), and burst tolerance. The first two are mandatory parameters, whereas the last two are optional (i.e., may not be specified).

PCR defines the minimum time between the interarrival times of two cells of a connection. This time, in general, is defined by the physical interface speed. The ATM end station may also submit its traffic to the network at a slower rate than that of the interface speed by shaping (i.e., buffering) its traffic. As an example, if the minimum time between two cells at the physical link is 1 millisecond (ms), the PCR of this connection can be up to 1000 cells per second. Although the application may generate cells at the PCR, it may choose to submit its cells so that the minimum time between two cells is 5 ms. This minimum time may be achieved by buffering cells and allowing them to leave the buffer at the rate of 200 cells per second. This process is referred to as shaping.

Mainly because of multiplexing at the ATM and the physical layers and the generation of management cells, the minimum interarrival times of cells may

not be exactly equal to their specified values (i.e., $1/PCR$). Another parameter, CDVT, is defined by the network to address this problem. In particular, this parameter defines how much deviation is allowed by the network so that the cell traffic is considered to be staying within the negotiated parameters (i.e., conforming).

The average rate of a connection is equal to the total number of cells transmitted divided by the duration of the connection. On the basis of this definition, the network can know the average rate of a connection only after the connection terminates. Accordingly, the average rate of a connection cannot be used in a meaningful way by the network until after the connection terminates. Sustainable cell rate is an upper bound on the average rate of an ATM connection. SCR is used together with another metric, the burst tolerance (BT), and the peak cell rate. BT is the time during which the source is allowed to submit traffic at its peak rate. Given these parameters, the maximum number of cells that can arrive at the switch back-to-back at the PCR, referred to as the maximum burst size, is given as $(1 + BT/(1/SCR - 1/PCR))$.

The QOS requirements of connections are agreed on between the ATM end station and the network during connection establishment time. The network commits to meet the QOS requirements for the duration of the connection as long as the connection stays within the agreed-on connection traffic descriptor values. This agreement between the ATM end station and the network across the UNI is referred to as a traffic contract.

The UNI specification also includes the conformance definition used to classify each arriving cell as either conformant (within the negotiated parameters) or nonconformant.

Data exchange interface

The physical layer interfaces for ATM over wide area networks included at the time the data exchange interface (DXI) was developed were DS-3, OC-3, and OC-12. Although it is envisioned that in the near future a large portion of the physical infrastructure will be fiber optic cable, currently these interfaces are still expensive, and it is necessary for the success of ATM to enable ATM services within the current physical infrastructure. The main objective of DXI is to provide access to ATM networks for the installed equipment without

costly (hardware) upgrades. To meet this goal, DXI allows data terminal equipment (DTE), i.e., a router, and data communications equipment (DCE), usually called an ATM-data service unit (ATM-DSU), to cooperate to provide a UNI for ATM networks.

DXI defines a data link control protocol and physical layers that handle data transfer between a DTE and a DCE.

The DXI framework defines the protocols for a DTE to transport a DTE-service data unit (DTE-SDU) to a corresponding peer entity across an ATM network.

DXI defines a data link control protocol and physical layers that handle data transfer between a DTE and a DCE. A DXI local management interface (LMI) and management information base are also specified as part of the DXI. DXI supports V.35, RS449, and HSSI physical layer interfaces at speeds ranging from several Kbps to 50 Mbps. The data link layer defines the method by which the DXI frames and their associated addresses are formatted for transport over the physical layer between the DTE and the DCE.

Broadband intercarrier interface

End-to-end national and international service requires networks belonging to different service providers to be interconnected. The broadband intercarrier interface (B-ICI) gives these networks the ability to interoperate to transport different services across each other. The ATM Forum B-ICI version 1.1 supports permanent connections only (i.e., PVCs).

B-ICI includes the specifications of the physical and ATM layers and service-specific functions above the ATM layer required to transport, operate, and manage different intercarrier services. The specification also includes traffic management, network performance, and operations and maintenance specifications.

The two B-ICI physical layer interfaces defined are 44.7 DS-3 and 155 Mbps SONET/SDH. The B-ICI is a multiservice interface that supports four inter-carrier services: cell relay service (CRS), circuit emulation service (CES), frame-relay service (FRS), and switched multimegabit data service (SMDS).

A service-specific non-ATM network is connected to an ATM network via an interworking unit (IWU) that allows non-ATM service to be mapped into ATM service and vice versa. The main functions performed at an IWU for each service supported are:

- CRS: receive ATM cells from one network and transmit to another over a permanent connection
- CES: receive DS-n frames, encapsulate them in AAL 1 PDUs, transmit over PVCs across a B-ICI, and reconstruct the original DS-n frames at the other end
- FRS: receive frame-relay frames, encapsulate them in AAL 5, transmit over PVCs across a B-ICI, and reconstruct the original frames at the other end
- SMDS: receive SMDS interface protocol L3 protocol data units (PDUs) encapsulated in intercarrier service protocol connectionless service (ICIP-CLS) PDUs, encapsulate them in AAL 3/4, transmit over PVCs across a B-ICI, and reconstruct the ICIP-CLS PDUs at the other end

The ATM cells of a particular service are multiplexed together and passed across B-ICI over one or more VPC or VCC, or both, that are preconfigured at subscription time. For each service supported, there is at least one connection across a B-ICI. Cells belonging to different services are not multiplexed onto the same connection.

The traffic management and congestion control framework across B-ICI follows very closely the framework developed in the UNI 3.1 specification. The UNI traffic parameters apply to B-ICI, with some simplifications. For example, the source traffic descriptor across B-ICI is required to include service type, conformance definition, peak cell rate, and cell delay variation tolerance; the inclusion of sustainable cell rate and burst tolerance in the traffic contract is an option.

How does it fit together?

When a connection request arrives at the network, the signaling message includes information for the network to characterize the call (i.e., its traffic characteristics and service requirements) and the

destination ATM address, among other call-related information. On the basis of this information, the network (or the switching system) needs to locate the destination, find a path that can provide the requested service guarantees without causing service degradation to connections already established in the network, and establish a connection in the network if the call is accepted.


UNI specification only includes the characterization of the connection request. Private NNI includes a routing framework, particularly the distribution of topological information in the network. This information is used by the switches that connections originate to find a path that has a high probability of accepting the call. Every switching system makes a call admission decision based on a call admission control procedure specific to the switch.

The other interfaces such as public UNI and NNI, ATM Forum B-ICI, and DXI do not provide any network control services to enable ATM networks.

NBBS is a comprehensive architecture for providing network services for high-speed, multimedia networks based on ATM standards and ATM Forum implementation agreements. In particular, NBBS provides a set of standards-based value-added services that include:

- A set of transport services that allow efficient integration of traffic generated by different sources with varying delay and loss requirements
- A bandwidth management framework that provides quality-of-service guarantees while achieving high resource utilization in the network
- Efficient distribution of control information in the network, minimizing the resources used for network control overhead while updating control information at the network nodes (at propagation delay speeds)
- Set management techniques that define and maintain logical groups and allow the network to learn about new resources and users automatically
- Nondisruptive path switching that minimizes disruption to users in case of network resource failures
- Call preemption that allows priority handling of connections and minimizes possible cascading of connection takedowns
- A path selection framework that addresses the requirements of different types of applications with varying delay and loss demands while achieving high utilization of network resources;

Figure 7 Signaling flows across a UNI



this framework includes both point-to-point and point-to-multipoint path construction

These services are described in detail in other papers in this issue.⁵⁻¹⁰ In this section, we describe how the NBBS architecture complements ATM standards using a point-to-point connection setup based on ATM Forum UNI 3.1 signaling, illustrated in Figure 7.

Consider a user application at an ATM end station that wants to establish an SVC to a partner application in another ATM end station (referred to as the destination node) across an ATM network. To initiate the process, the originating user sends a SETUP message across the UNI to the network. In the case of NBBS, this message is received by an NBBS ATM access agent (originating access agent) at the edge of the network. The SETUP message carries various information elements (IEs) used by the

network to establish the VCC and those that are used by the destination ATM end station. From the network point of view, the most important IEs are:

- Called party IE—identifies the destination ATM end station
- ATM traffic descriptors—a collection of traffic parameters that describe the traffic flow for the forward and reverse paths; the values for the forward and reverse path may be different
- Broadband bearer capability—describes the type of service that the user application requests for this connection
- QOS class of the connection

The first IE in the list is used by the network to locate the destination node, whereas the latter three are used to define the application requirements for a VCC.

Upon receipt of the SETUP message from the originating user, the first thing the originating access agent must do is locate the target end station. To accomplish this, the access agent employs NBBS *directory services*, a distributed directory function

Source routing requires the source node to know the topology of the network.

defined by an integrated set of query and database algorithms. After the target is located, the access agent uses the source traffic characteristics and application requirements in the SETUP message IEs to determine the amount of bandwidth necessary to support the user's application. Bandwidth calculation takes into consideration the cell loss requirement and the source traffic characteristics, and determines a guaranteed upper bound on the bandwidth required to support the application when this connection is multiplexed with the connections already established in the network. Then the agent runs the path selection algorithm to determine a path from the originating node to the destination nodes. In determining this path, the end-to-end delay requirement specified by the application is taken into consideration explicitly. In particular, the algorithm finds a minimum-hop path in the network that can support the bandwidth and end-to-end delay requirements of the application.

This type of path computation in which end-to-end paths are determined by the node where the connection originates is referred to as *source routing*. As discussed previously, the ATM Forum private NNI routing framework also uses source routing between ATM switching systems. Source routing requires the source node (i.e., a node that supports ATM access service) to know the topology of the network and the utilization and the current reservation levels of network resources. The topology of the network may change at times because of link and node failures. Although the topology of the network does not change often, utilization levels change frequently. Such change necessitates an efficient means of distributing network control infor-

mation that maximizes the amount of information available to each node while minimizing the amount of network control traffic overhead. These two objectives are contradictory, and the design and development of this feature is a complex task. NBBS addresses this problem by using a reliable, hardware-assisted multicast over a spanning tree. The spanning tree is built and maintained using a distributed algorithm. In addition, a spanning tree minimizes the total number of links used in the network to provide a path from one network node with a control point function to another.

Assuming that a path between the originating and destination end stations is found that can support the connection, the access agent proceeds with the end-to-end connection establishment using the NBBS network connection (NC) services, a set of algorithms and messages that enable fast connection setup and real-time connection maintenance in a dynamic, high-speed environment. Using the NC services, the originating access agent multicasts an NBBS connection request message only to the NBBS nodes along the path chosen for the connection, including the destination access agent that serves the called user. The connection request message contains bandwidth and QoS requirements for the connection. It also includes the information from the SETUP message having an end-to-end significance. Each node along the path determines whether it can support the required service (bandwidth, delay, loss) and then responds directly to the originating access agent. If the destination access agent can support the connection, it takes the information received in the NBBS setup message (including IEs with end-to-end significance in the original SETUP message), builds a SETUP message similar to the one built by the source user, and sends it to the target user. If an intermediate node or the destination station or both reject the connection request, NBBS connection services inform all nodes along the path that the connection is rejected and needs to be torn down (for nodes that accepted the request to release associated resources).

The destination user receives and processes the SETUP message. If it accepts the call, it responds with a CONNECT message, which includes any end-to-end information that had to be negotiated. The destination access agent forwards this information back to the originating access agent and sends a CONNECT ACKNOWLEDGE message to the target

user to indicate that the CONNECT was received and processed. If every node along the path can support the connection, the originating access agent sends a CONNECT message back to the source user. The source responds with a CONNECT ACKNOWLEDGE message. At this time, the connection is fully enabled, and user ATM cells can start flowing.

In summary, the ATM UNI architecture defines the basic framework upon which different services for varying source characteristics and quality-of-service requirements can be supported in an integrated manner. UNI signaling essentially provides the basic language for an ATM end station and an ATM switch to interface with each other. Based on this standards framework, it is the responsibility of the network to provide a comprehensive set of control services to supply the ATM service.

Private NNI is the only interface that will include some of the network control services required to establish switched ATM connections across switching systems that include topology database formats and advertising, clustering, P-NNI link initialization, and so forth. Although these services are required for basic interoperability among switching systems, the capabilities currently defined in this framework are not rich (lack of group management, directory services, and so forth), and various services are not included within the scope of the specification (path selection, call admission control, congestion control, and so forth) or do not result in efficient use of bandwidth (flooding of the topology database).

Consider a network with traffic generated by a large number of connections. At the edge of the network, it is necessary to monitor the traffic generated by each source to ensure that each source stays within the parameters negotiated at the call establishment phase. Otherwise, nonconforming sources would cause the quality of service provided to conforming sources to degrade.

Let us now focus on the nonconforming traffic. There are various reasons why a source might not stay within the parameters agreed on at the call setup time: it may not be possible to characterize the traffic behavior accurately, equipment might malfunction, or the source might simply be cheating. Two mechanisms are in the NBBS architecture to provide some amount of forgiveness to nonconforming sources. One is the *adaptation function* that monitors the traffic and estimates the amount

of bandwidth required for connections by filtering the actual traffic generated by each source. The other mechanism is the use of cell marking, referred to as *red/green marking*. In this mechanism, conforming traffic is tagged green (high priority) before it is transmitted, whereas nonconforming cells are tagged red (low priority). The red/green marking uses an ATM cell header priority bit. Allowing (some) nonconforming traffic to enter the network provides a better service to applications, for example, that may not be able to characterize their traffic behavior accurately; otherwise their cells would be dropped at the interface. By allowing nonconforming traffic to enter, the network resources are better utilized as well. Doing so, however, requires the development of mechanisms so that the service provided to conforming traffic in the network is not affected by the nonconforming traffic at times when one or more nodes in the network become congested. To achieve that, NBBS uses a discarding mechanism at the intermediate switches so that low-priority cells (i.e., tagged cells) are discarded before they can negatively impact the service provided to high-priority cells.

Another service required in the network is to minimize the potential impact of resource failures to network applications. Nondisruptive path switching (NDPS) is used in NBBS to reroute connections established on a link (node) around failures in a way that the impact of the failure on the service provided to these affected connections is minimal.

NDPS is also used as a part of the framework developed in NBBS to support different connection priorities. Using this function, a high-priority connection may cause one or more low-priority connections to be rerouted in the network to make resources available along the path used for the high-priority connection. The main challenge here is to minimize the cascading effect that may occur when a connection that was taken down tries to reestablish a new connection. NDPS takes into consideration such cascading effects when connections are rerouted in the network.

PVC connections. From the perspective of a standard, a PVC (permanent virtual connection) VC (either a VPC or VCC) encompasses only the physical, ATM, and management aspects of the UNI. This means that the setup of a PVC VC is initiated without the user signaling across the UNI to its network interface. Accordingly, the processes of creating

and deleting PVCs are invoked mostly through the network management function.

NBBS supports permanent connections using NBBS network management.⁹ The message used to establish a permanent connection conveys the connection characteristics that include connection

NBBS supports permanent connections using NBBS network management.

bandwidth, its QOS, and an indicator to specify whether the connection is permanent or semipermanent. Semipermanent connections are set up and taken down on a time-of-day basis.

To make this process as dynamic as possible in NBBS, only the nodes that serve the calling and called-user-UNIs receive related connection setup network management flows. In particular, only one such node is selected by the network management to initiate the setup. This node then invokes the NBBS functions and services required to set up the connection inside the network. As discussed previously, these services include directory services to locate the remote edge node, path selection to find a suitable path across the network, and network connection services to reserve resources along the switches on the connection path. Note that the PVC VC and switched VC setup inside the network are very similar processes. The switched VC setup is reviewed in the previous section and not repeated here.

Since the details of these NBBS services are presented in other papers in this issue, we next proceed with a review of various extensions to the architecture to support ATM in the network in a manner that conforms to ATM standards.

Implementation of ATM standards in the network infrastructure

So far, we have reviewed ATM standards and discussed briefly how NBBS services complement

these standards to enable ATM networks to meet application needs while addressing them efficiently in the network. In this section, we review how ATM connections are mapped into NBBS network connections (thereby obtaining NBBS network services) while providing ATM service in the network, and we discuss various ATM-related extensions to the architecture.

ATM connections. An ATM network supports connections with different service requirements. To integrate applications with different requirements in the network, connections supported in NBBS are classified into four categories: real-time-1, real-time-2, nonreal time, and best effort. Real-time-1 corresponds to ATM CBR service, real-time-2 to real-time VBR, nonreal time to nonreal-time VBR, and best effort to unspecified bit rate or available bit rate services.

Real-time traffic has stringent delay and loss requirements and requires bandwidth reservation to guarantee its QOS metrics. The difference between the two real-time classes is that real-time-1 corresponds to CBR service, whereas real-time-2 corresponds to VBR service. Nonreal-time traffic also requires bandwidth reservation but only for the loss guarantee; i.e., this type of traffic can tolerate moderate delay. Nonreserved traffic in contrast can tolerate both greater loss and larger delay. Accordingly, no bandwidth is reserved for this type of traffic.

NBBS provides both virtual path connection (VPC) and virtual channel connection (VCC) ATM bearer services between two or more UNIs. In either case, the connectivity is provided between two or more access agents supporting the ends of the ATM UNIs. An NBBS network connection emulating a VPC allows the user of the VPC to multiplex VCCs onto a single network connection. In the case of a network connection emulating a VCC in the network, there is a one-to-one correspondence between a VCC and a network connection.

When a network connection emulates a VPC, intermediate switches along the path of the connection swap only the VPIs in the cell headers, whereas both the VPI and VCI values may be swapped for network connections that emulate a VCC.

Given this mapping from ATM connections to the NBBS connections, all NBBS services discussed

throughout this issue apply to ATM service in a relatively straightforward manner.

VP trunk groups. An ATM bearer service VPC is a virtual path connection leased from an ATM bearer service provider. In NBBS these connections are referred to as virtual path (VP) trunks since they are essentially transmission links connecting two NBBS networks to each other.

In NBBS, information on network resources, i.e., link, nodes, bandwidth availability, etc., are kept in a topology database that is replicated in every network node. In an ATM network, there can potentially be a large number of VP trunks between two network nodes for a large number of switch pairs. Representing them individually in the topology database would require a large number of database entries. In addition, if each VPC is managed separately, a large amount of control traffic in the network is generated, thereby increasing the amount of bandwidth used for the control traffic.

A VP trunk can be used by one virtual path connection (VPC) only. In order to allocate a connection on a VPC, it is necessary to know its bandwidth and QOS metrics. Representing each VPC individually with its characteristics in the topology database (TDB) would allow the nodes that are attached to ATM end stations to know the characteristics of each VPC and establish an end-to-end network connection that emulates a VPC over a VP trunk meeting the requirements of the connection. Though simple and flexible, this solution would require one entry in the TDB for each VP trunk in the network.

In NBBS, we introduced the concept of *VP trunk group* to address this problem. A virtual path trunk group is a collection of VP trunks. More precisely, there is a controller, referred to as a *link manager*, and a VP trunk group is the collection of VP trunks between two link managers.

A VP trunk group in NBBS is represented as a single entry in the topology database. The immediate consequence of this approach is that the QOS supported by an individual VP trunk is no longer distinguishable. However, the VP trunk group representation reduces the number of entries representing the VPCs in the topology database to a minimum. In addition, the bandwidth management procedure developed for the VP trunk group minimizes the amount of control traffic generated for

the trunk group by managing individual VPCs locally without any impact on resource utilization in the network.

On the basis of this framework, each VP trunk is represented individually in a local storage, referred to as a VP trunk group database (VPDB). A local management procedure monitors and manages the

Bidirectional NBBS connections are not forced to follow the same physical path in each direction.

VPDB. The management procedure also filters the changes in the VPDB and updates the VP trunk group entry in the TDB, causing topology database updates to be sent to all other nodes in the network.

As discussed previously, in order for switches at which connections originate to determine the end-to-end paths, they need to have information on the network topology and availability of network resources. The amount of bandwidth advertised for the VP trunk group is the bandwidth of a single VP trunk that is not currently allocated to a connection with the maximum bandwidth. VP trunks are allocated to VP connections based on the *best-fit* rule. With this rule, the chosen VP trunk is the one that has available bandwidth greater than or equal to the bandwidth requirement of the connection and minimizes the difference between the VPC bandwidth and the requested bandwidth.

Same physical path selection. Bidirectional NBBS connections are not forced to follow the same physical path in each direction. Because a connection can have different bandwidth and delay requirements in each direction, allowing each unidirectional connection to be established independently permits the network, in general, to support more connections than if they were constrained to the same physical path in both directions. Bidirectional ATM connections, however, are required to follow the same physical path in the network (i.e., at the public NNIs). This requirement is provided in NBBS

by extending the NBBS connection protocol and the path selection algorithm.

In a manner similar to the NBBS connection protocol for connections with separate physical paths, the originating access agent (AA) owns the bidirectional ATM connection with the same physical path requirement. In order to establish the connection, assuming that a path is found, the originating AA sends two connection requests along the path to the destination AA: one for the forward path and the other for the return path. The destination AA owns the return path. Similarly, both AAs perform network connection maintenance, i.e., connection refresh, bandwidth adjustment, and connection termination functions in the directions they own.

So far we have assumed that a path in the network supporting the connection along the same physical links is found. However, in general, it cannot be guaranteed that such a path exists, although the network may support the connection if separate paths are used. Even if there is such a path, finding the path cannot be guaranteed in real time. In particular, when the links have different delay values in each direction, finding a path that minimizes the number of hops while satisfying end-to-end delay constraints in both directions simultaneously has a nonpolynomial time complexity, i.e., any algorithm that would find such a path if one exists has an exponential time complexity. It requires the use of a heuristic approach that would work fairly well in most cases. The heuristic approach chosen to address this problem in NBBS minimizes the changes to the current NBBS path selection algorithm. If this extended path selection algorithm cannot find a single path that meets the end-to-end delay and bandwidth requirements of the connection in both directions, either the connection request can be rejected or the path selection algorithm can be used to find separate physical paths for a (predetermined) number of times.

In general, the transmission links in ATM networks are expected to have the same delay values in each direction. Although the proposed algorithm should address the problem with asymmetric transmission links, it should not have any added complexity to address the case in which a majority (or all) of the transmission links in the network have similar delay characteristics in each direction. The simplest extension is to use a single link metric that is the maximum of the two delay metrics in each direction of a transmission link and apply the NBBS path

selection algorithm. This approach would also work without any performance penalty if the links have the same delay values in each direction. However, this approach may be too restrictive. In addition to keeping track of the total delay from the originating node to a network node in the original NBBS algorithm, the extended algorithm keeps track of the total delay from the originating node to a network node and from the network node to the source. The NBBS path selection algorithm is then applied with the maximum of each metric.

Bandwidth allocation in ATM standard mode. The NBBS bandwidth allocation procedure for bursty sources guarantees service requirements of applications while achieving high multiplexing capability in the network. However, its development is based on a particular assumption about the behavior of the traffic. ATM source parameters, i.e., peak cell rate, sustainable cell rate, and burst size, are deterministic parameters defined more for ease of characterization than for modeling real source behavior. Thus, the NBBS bandwidth allocation procedure must be extended to model the behavior of sources that define their traffic characteristics with deterministic parameters. For ATM sources, the extended NBBS bandwidth allocation procedure allows a larger number of sources to be multiplexed in the network while guaranteeing their service requirements.

Summary

ATM provides the basic standards framework for integrated voice, video, and data services in high-speed networks supporting multimedia applications. Much work has been done in defining the B-ISDN standards in which ATM transport is featured. Both international standards bodies and the ATM Forum continue the prolific work to enhance the standards and implementers' agreements fundamental to realizing customer expectations.

In order to enable ATM networks, it is necessary to provide a comprehensive set of network services that complement the ATM standards while achieving high resource utilization in the network.

NBBS is IBM's architecture for high-speed, multimedia networking. The architecture not only provides network control services required in ATM networks, but it also provides an efficient solution to implementing ATM standards and specifications. This collection of NBBS services is now and will

remain complementary to existing standards and forum agreements.

Various NBBS services that are essential to providing ATM services are reviewed in detail elsewhere in this issue. In this paper, we discussed how these services are used in ATM networks and presented a high-level review of various NBBS services extended to support ATM services in a standards-compliant manner.

Acknowledgments

The authors gratefully acknowledge John Drake for valuable comments on early versions of this paper and the members of the NBBS team for their valuable contributions to the NBBS ATM framework.

Appendix

In this appendix, we summarize a set of features specified in different ATM standards or specifications. NBBS is compliant with all of these requirements, including the functions that are listed as optional in Tables 2, 3, 4, and 5.

Table 2 ATM layer specification

Function
ATM cell structure and encoding at the UNI
Multiplexing among different ATM connections
Traffic shaping
F4 OAM flows
F5 OAM flows
Reserved VPI/VCIs
Traffic and congestion control (traffic contract that includes connection traffic descriptor, requested service class, for each direction of the ATM layer connection and conformance definition)
QOS classes
Generic cell rate algorithm for peak cell rate with cell delay variation tolerance and sustainable cell rate with burst tolerance
Allowable combinations of traffic parameters
Usage Parameter Control (UPC) function
Selective cell discard

Table 3 ATM bearer services

ATM Bearer Service Attribute	Private UNI	Public UNI
Point-to-point VPCs	Optional	Optional
Point-to-point VCCs	Required	Required
Point-to-multipoint VPCs	Optional	Optional
Point-to-multipoint VCCs, SVC	Required	Required
Point-to-multipoint VCCs, PVC	Optional	Optional
Permanent virtual connection	Required	Required
Switched virtual connection	Required	Required
Specified QOS class	Optional	Required
Unspecified QOS class	Optional	Optional
Multiple bandwidth granularities for ATM connections	Optional	Required
Peak rate bandwidth enforcement	Optional	Required
Sustainable cell rate traffic enforcement	Optional	Optional
Traffic shaping	Optional	Optional
ATM layer fault management	Optional	Required
Interim local management interface	Required	Required

Note:

1. Public ATM network equipment conforming to this interface specification shall be capable of providing ATM users with either a VPC service, or a VCC service, or combined VPC/VCC service.
2. Public ATM network equipment conforming to this interface specification shall be capable of providing ATM users with either support for VPC or SVC capability or both.
3. Only one of the specified QOS connection categories is required at the public UNI.

Table 4 Interim local management interface

ILMI Functions
<i>ILMI service interface</i> Simple Network Management Protocol for monitoring and control operations of ATM management information across UNI that include physical and ATM layer, ATM layer statistics, VP and VC connections
Address registration
Default VCC (VPI = 0 VCI = 16)

Table 5 UNI 3.1 signaling

Feature
Support of demand switched connections (point-to-point and point-to-multipoint)
Support of connections with symmetric and asymmetric bandwidth
Class A, class C, and class X service classes
Support of a signaling VC
Three private and one public ATM address
Client address registration using ILMI
UNI 3.0 signaling for point-to-point demand switched connections
UNI 3.0 signaling for point-to-multipoint demand switched connections

Cited references

1. ITU-T Recommendation I.151, *B-ISDN Asynchronous Transfer Mode Functional Characteristics*, International Telecommunication Union, Geneva.
2. R. A. Sultan and C. Basso, "ATM: Paving the Information Superhighway," *IBM Systems Journal* **34**, No. 3, 375-389 (1995).
3. R. O. Onvural, *ATM Networks: Performance Issues*, Artech House Inc., Norwood, MA (1994).
4. G. A. Marin, C. P. Immanuel, P. F. Chimento, and I. S. Gopal, "Overview of the NBBS Architecture," *IBM Systems Journal* **34**, No. 4, 564-589 (1995, this issue).
5. M. Peyravian, R. Bodner, C.-S. Chow, and M. Kaplan, "Efficient Transport and Distribution of Network Control Information in NBBS," *IBM Systems Journal* **34**, No. 4, 640-658 (1995, this issue).
6. H. Ahmadi, P. F. Chimento, R. A. Guérin, L. Gün, B. Lin, R. O. Onvural, and T. E. Tedijanto, "NBBS Traffic Management Overview," *IBM Systems Journal* **34**, No. 4, 604-628 (1995, this issue).
7. T. E. Tedijanto, R. O. Onvural, D. C. Verma, L. Gün, and R. A. Guérin, "NBBS Path Selection Framework," *IBM Systems Journal* **34**, No. 4, 629-639 (1995, this issue).
8. C. P. Immanuel, G. M. Kump, H. J. Sandick, D. A. Sinicropo, and K. V. Vu, "Access Services for the Networking BroadBand Services Architecture," *IBM Systems Journal* **34**, No. 4, 659-671 (1995, this issue).
9. S. A. Owen, "NBBS Network Management," *IBM Systems Journal* **34**, No. 4, 725-750 (1995, this issue).
10. *User-Network Interface Specification: Version 3.1*, ATM Forum Specification, ATM Forum, 480 San Antonio Road, Suite 100, Mountain View, CA 94040-1219 (1994).
11. *Data Exchange Interface*, ATM Forum Specification, ATM

Forum, 480 San Antonio Road, Suite 100, Mountain View, CA 94040-1219 (1994).

12. *B-ISDN Inter-carrier Interface*, ATM Forum Specification, ATM Forum, 480 San Antonio Road, Suite 100, Mountain View, CA 94040-1219 (1994).

General references

- ITU-T Recommendation I.121, *Broadband Aspects of ISDN*, International Telecommunication Union, Geneva.
- ITU-T Recommendation I.327, *B-ISDN Functional Architecture*, International Telecommunication Union, Geneva.
- ITU-T Recommendation I.361, *B-ISDN ATM Layer Specification*, International Telecommunication Union, Geneva.
- ITU-T Recommendation I.371, *Traffic Control and Congestion Control in B-ISDN*, International Telecommunication Union, Geneva.
- ITU-T Recommendation I.610, *B-ISDN Operation and Maintenance Principles and Functions*, International Telecommunication Union, Geneva.
- ITU-T Recommendation Q.2931: *B-ISDN Signaling*, International Telecommunication Union, Geneva.

Accepted for publication July 24, 1995.

Raif O. Onvural *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: onvural@vnet.ibm.com).* Dr. Onvural is a senior engineer at IBM's Research Triangle Park facility in the Networking Architecture organization and manages the Networking Technology Architecture department that develops network control services for ATM networks. He is also IBM's venue owner for the ATM Forum and has been attending the forum meetings since February 1993. Dr. Onvural organized several international conferences on high-speed networks, in general, and ATM networks, in particular. He has published in various journals and conferences, and has edited five books. He is also the author of the book *Asynchronous Transfer Mode Networks: Performance Issues*.

Haldon J. Sandick *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: sandick@vnet.ibm.com).* Mr. Sandick is an advisory programmer in the Networking Hardware Division's Networking Architecture (NA) group. He joined IBM in 1987 to work on the VTAM[®] design of APPN Directory Services. In 1990, he moved to the Networking BroadBand Services architecture team in NA, where he has worked on the architecture required to support various access technologies. During the last few years he has focused on supporting ATM interfaces. He received a B.S. in communication from Hunter College, City University of New York, an M.S. in education from Fordham University, and an M.S. in computer science from New York University. In 1994 Mr. Sandick received an Outstanding Technical Achievement Award for his work on the NBBS architecture and an IBM First-Level Invention Achievement Award. His current interests include multiprotocol support in high-speed networks and design of intelligent network services.

Garry M. Kump *IBM Networking Hardware Division, 800 Park Offices Drive, Research Triangle Park, North Carolina 27709.* Mr. Kump is a staff engineer in the Networking Architecture organization, where he works with the Networking BroadBand Services group. After serving in the United States Air Force from 1976 to 1984 as an electronic communications technician, he received his B.S.E. degree in electrical engineering from the University of Nebraska at Lincoln, where he became a member of IEEE and was accepted into the Eta Kappa Nu Electrical Engineering Honor Society. Coming to IBM in 1989, he first worked in unit testing of Interconnect Communication products, moving shortly thereafter into product development and then to communication architecture. His main interests lie in software development, especially in the area of high-speed communications.

Elizabeth A. Hervatic *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: hervatic@ralvm6.vnet.ibm.com).* Ms. Hervatic is an advisory programmer in ATM Market Development. She joined IBM in 1989 and has worked in the Networking Architecture organization. Among other positions, Ms. Hervatic has been the technical leader of the NBBS base architecture team. Ms. Hervatic received her B.S. in computer science from North Carolina State University in 1989.

Phillip F. Chimento *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709.* Dr. Chimento received the A.B. degree in philosophy from Kenyon College in 1972, the M.S. degree in computer science

from Michigan State University in 1978, and the Ph.D. degree in computer science from Duke University in 1988. He worked for IBM from 1978 to 1994, holding various positions in design, development, test, and architecture. Most recently, he was a member of the core team that developed IBM's Networking BroadBand Services architecture for high-speed packet and cell switching. In 1994, Dr. Chimento took a leave of absence from IBM to accept a visiting faculty position at the University of Twente in the Netherlands. There, as a member of the Centre for Telematics and Information Technology (CTIT) and the Tele-Informatics and Open Systems (TIOS) group, he is working on B-ISDN signaling and resource allocation issues and participating in Dutch and European telecommunications projects. He has had papers published in *IEEE Transactions on Computers*, *Operations Research*, and various conferences. He is a senior member of the IEEE and a member of the ACM and ORSA (INFORMS).

Reprint Order No. G321-5588.