Network Aspects of Broadband ISDN

Marius D. Soneru
Jon Anderson
Henry J. Kafka

The emerging Broadband Integrated Services Digital Networks (BISDNs) are expected to be flexible, supporting a wide spectrum of services and infrastructure applications, and responsive to evolving customer needs. The networks are also expected to be cost-effective, supporting diverse services on a common platform, and highly manageable, providing unified operations and customer control. Meeting all these expectations represents a major challenge in the design and implementation of BISDN networks. This paper introduces the basic structure of BISDN networks by describing a network reference model and the BISDN Protocol Reference Model. It also describes BISDN functions for signaling and call control, network management and operations and maintenance (OAM), traffic management and congestion control, and network performance and quality of service. The evolution of the BISDN networks is also presented, including an example of a BISDN network interworking with service-dependent networks.

Introduction

The BISDN networks will support a wide range of voice, data, and video applications based on a limited set of service and connection types and multipurpose network interfaces. The service types include constant bit rate/circuit emulation, variable bit rate, connection-oriented data, and connectionless data services. The connection types supported by BISDN networks include point-to-point and multipoint, both switched and non-switched. Multipurpose network interfaces have been defined for both private and public networks.

To meet these diverse needs, the asynchronous transfer mode (ATM), which provides standardized transport, multiplexing, and switching for the BISDN network, was selected as a flexible transfer capability common to all services.[1] It has many of the advantages of both circuit and packet switching, and it allows all services to be switched and transported in a common digital format.

Network Reference Model

The BISDN network reference model, shown in Figure 1, identifies the principal functional elements that make up the network, their relative positions in the network, and the reference points between these elements.[2] These functional elements generate, process, and terminate various layers of the BISDN Protocol Reference Model, discussed briefly in the next section. No physical implementation of the functional elements is implied in this model. However, physical interfaces, called user-network interfaces, are often associated with the U_B and T_B reference points. A network node interface is also depicted in Figure 1.

The equipment that resides on customer premises includes the following functional elements:

- Broadband network termination equipment (B-NT1 and B-NT2)
- Broadband terminal adapter (B-TA)
- Broadband terminal equipment (B-TE1,B-TE2) or other, non-broadband, terminal equipment (TE2).

The B-NT1 terminates the physical layer, handles interfacing at the U_B and T_B reference points, and provides OAM functions. Optionally, B-NT1 may multiplex several T_B interfaces into a single U_B interface. The

B-NT2 also performs adaptation functions for different media at the physical layer. It may also provide higher-layer functions, such as cell delineation, multiplexing and demultiplexing, signaling adaptation layer functions, signaling protocol handling, and local switching. Examples of B-NT2 equipment include a simple time-division multiplexer, a terminal controller or a local area network (LAN), and a private branch exchange (PBX). The B-NT2 is optional, allowing terminal equipment to access the broadband network directly at the T_B interface. To make this possible, the functional and physical specifications of the S_B and T_B interfaces need to be similar.

Terminal equipment with ATM service interfaces (B-TE1) are connected at the S_B/T_B interfaces, while those without an ATM service interface (broadband, B-TE2 or non-broadband, TE2) are connected to a B-TA. The B-TA adapts the native protocol at the R interface to the ATM-based protocol at the S_B/T_B interfaces.

Network elements include the remote multiplexer, the access node, and the service node. The remote multiplexer is an optional network element that multiplexes several subscriber links onto a single higher-speed link. The access node, another optional network element, concentrates traffic and can also perform local switching. The service node, the equivalent of the local switch in the ATM network, contains an ATM cell-switching fabric. It provides all functions required for broadband switching, exchange termination, interworking with non-ATM network elements, operations control, signaling processing, and call processing control.[3] Service nodes communicate over links using the network-node interface.

Protocol Reference Model

The BISDN Protocol Reference Model can be represented as three vertical planes — the user plane, control plane, and management plane. These are intersected by four horizontal layers — the physical layer, ATM layer, ATM adaptation layer, and higher-layer protocols and functions.[4]

The user plane facilitates the flow of user information. The control plane provides call and connection control functions, including signaling for call and connection control set-up, supervision, and release. Both planes, user and control, have layered structures that incorporate all four horizontal layers. The management plane handles layer management and plane management. Layer management provides management functions and operations and maintenance to the individual layers. Plane management, which has no layered structure, coordinates all the planes.

The physical layer transports information (bits/cells) and delivers timing to upper-layer services. The ATM layer multiplexes/demultiplexes and routes cells. The ATM adaptation layer provides service-dependent functions to the layer above. Higher-layer protocols and functions are service-specific. In the user plane, they support the user application (e.g., video coding); in the control plane, they provide call and connection control; and in the management plane, they provide network supervision functions. The physical layer, ATM layer, and ATM adaptation layer are briefly described in the next three sections.
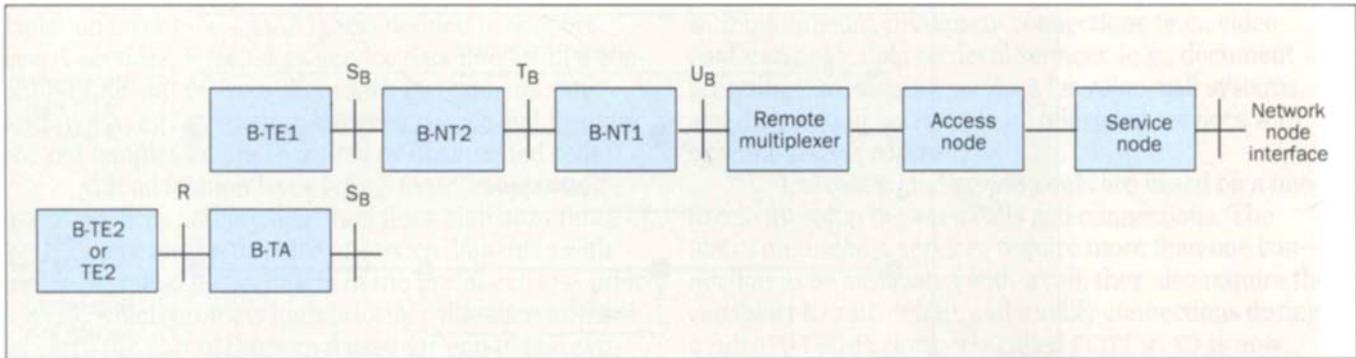
Figure 1. BISDN network reference model.

Physical Layer

Three transmission systems have been specified by the International Telecommunication Union–Telecommunication Standardization Sector (ITU-T), formerly CCITT: the Synchronous Digital Hierarchy (SDH), the Plesiochronous Digital Hierarchy (PDH), and cell based. In addition, the ATM Forum has specified the fiber-distributed data interface (FDDI) option. For the SDH option, the STM-1 signal has been defined for both optical and electrical interfaces with a total bit rate of 155.520 Mbits/s, in both directions. The STS-3c is an almost identical interface defined as part of the synchronous optical network (SONET) for use in the U.S. The PDH-based interfaces have been defined for carrying ATM cells using the existing transmission infrastructure (e.g., DS3). The cell-based interfaces are defined as a continuous transport of cells without time-related framing.

ATM Layer

A connection-oriented transfer mode, ATM uses an asynchronous time-division multiplexing technique. Connections can be established and released, either on a subscription basis or on demand using signaling procedures.

ATM Connections. The ATM layer has two levels of connections—virtual channel (VC) connections, which are the basic type of ATM connections; and virtual path (VP) connections, each of which contains a bundle of VC connections that are switched as a unit (see Figure 2). The VP connections are semi-permanent and transport a number of simultaneous VC connections between endpoints.

A two-part label within the ATM header uniquely identifies the VP, and the VC within the VP to which that cell belongs. The label contains a VP identifier and a VC identifier. Both the VP and VC identifiers have only local significance on the link between ATM nodes. Routing consists of translating the VC/VP identifier values of the incoming VC/VP link into the corresponding VC/VP identifiers for the outgoing VC/VP link. The translation is based on routing tables located in the VC switch (VC switching only), VP switch (VP switching only), or VC/VP switch (VC and VP switching).

Signaling and Metasignaling VC Connections. BISDN supports two types of access signaling configurations — point to point, and point to multipoint. In a point-to-point access signaling configuration, only one signaling entity (e.g., a TE or B-NT2/PBX, by subscription) on an interface communicates with the service node. A point-to-multipoint access signaling configuration may have multiple signaling entities on an interface communicating with the service node.

Connections can be set up and released using special connections between a TE and the service node, called the signaling VC connections. In a point-to-point signaling configuration, a point-to-point signaling VC connection is assigned at subscription time. In a point-to-multipoint signaling configuration, two signaling VCs are dynamically assigned to a TE, at activation time, using a predefined VC connection, called the metasignaling VC connection. One is the point-to-point signaling VC and the other is a broadcast VC used for call offering by the service node.

ATM Adaptation Layer

The ATM adaptation layer provides the adaptation functions needed to implement two-way mapping between the higher-layer protocol data units and the information fields of ATM cells. As a result of the role that the ATM
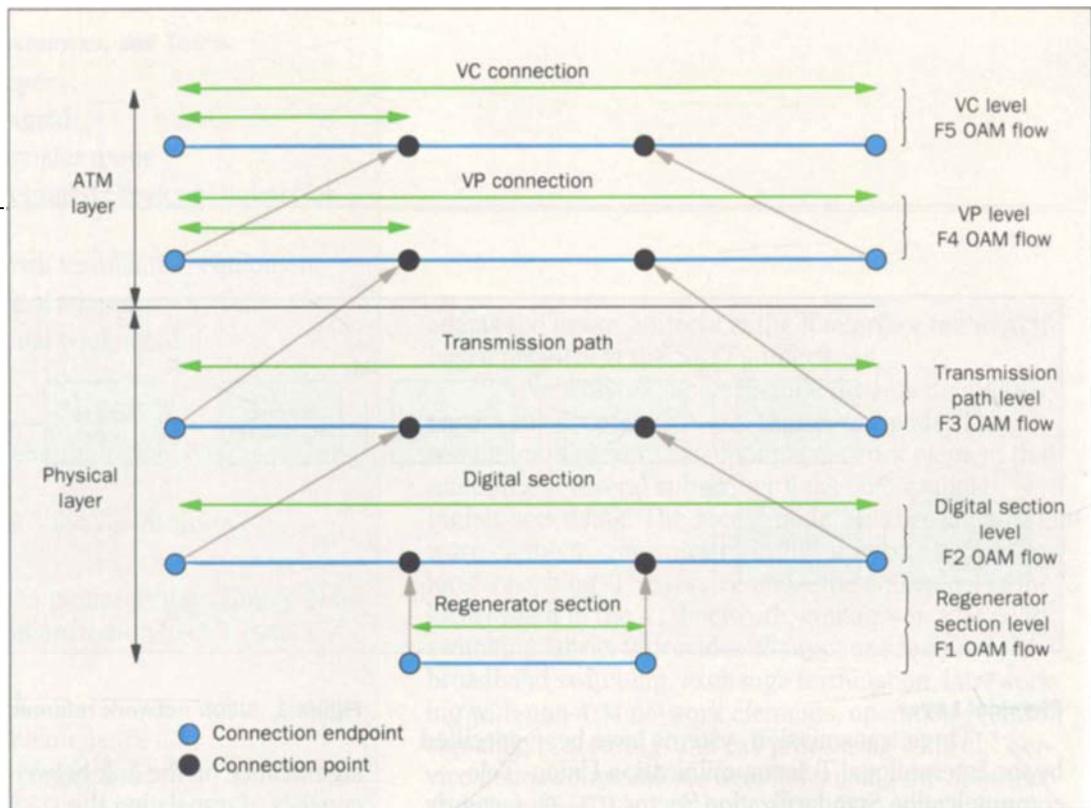
Figure 2. OAM architecture of the ATM network. The ATM layer has two levels of connections — VC connections, which are the basic type of ATM connections; and VP connections, each of which contains a bundle of VC connections that are switched as a unit. F flows manage the layers' functions. The F1, F2, and F3 flows correspond to the regenerator section, multiplex section, and path of the physical layer, respectively. F4 and F5 flows correspond to the ATM layer VP connection and VC connection levels, respectively. The F4/F5 flows are independent of the physical layers used in the BISDN. They provide the end-to-end, segment OAM and management capabilities of associated VP/VC connection-based services. The management plane coordinates the layer management capabilities and, in essence, provides the Q3 interface.

adaptation layer plays, its functions depend on the services supported. To minimize the number of ATM adaptation layer protocols needed for the different services, the ATM adaptation layer was partitioned into two sublayers, and the services to be supported were categorized into four service classes.

The lower ATM adaptation sublayer, called the segmentation and reassembly sublayer, segments the information received by the upper layer into ATM cells at the transmitting end, and reassembles the original information at the receiving end. The segmentation and reassembly sublayer can also provide additional functions, as needed, for some services (e.g., bit-error detection). The upper ATM adaptation sublayer, called the convergence sublayer, provides additional service-specific functions, such as multiplexing, timing recovery, and cell loss detection.

Services to be supported by the ATM adaptation layer are based on the following characteristics:
- Timing relationship between source and destination,
- Bit rate, i.e., constant bit rate, or variable bit rate, and
- Connection mode, i.e., connection-oriented or connectionless.

The four service classes and their characteristics are:
- Class A — Source–destination timing, constant bit rate, connection oriented (also called circuit emulation), such as constant-bit-rate voice and video.
- Class B — Source–destination timing, variable bit rate, connection oriented, such as variable-bit-rate audio and video.
- Class C — No source–destination timing, variable bit rate, connection oriented, such as signaling and connection-oriented data.
- Class D — No source–destination timing, variable bit rate, connectionless mode, such as Switched Multimegabit Data Services (SMDS).

A combination of one segmentation and reassembly sublayer protocol and one or more convergence sublayer protocols represents an ATM adaptation layer type. It provides a service access point to the higher layers. So far, ATM adaptation layer types 1, 2, 3/4, and 5 have been defined by ITU-T.

The following is a brief description of the functions performed by each ATM adaptation layer type. ATM

adaptation layer type 1 (AAL 1) was defined to support Class A services. It transfers service data units with a constant bit rate and delivers them with the same bit rate; transfers timing information between source and destination; and handles errors (e.g., lost or misinserted cells).

ATM adaptation layer type 2 (AAL 2) supports Class B services and provides functions similar to those of type 1, except for the transfer of service data units with variable bit rates. It also differs in the use of cell loss priority (CLP), which protects high-priority cells, such as those that carry the part of the signal used for end-to-end synchronization. CLP and its use are described later in this paper. The specification of AAL 2 is still at an early stage.

ATM adaptation layer type 3/4 (AAL 3/4) supports Class C and D data services. Its attributes are defined according to the particular service provided:

- Assured/non-assured — Assured service guarantees error-free delivery of all ATM adaptation layer service data units using flow control. The non-assured service may lose or mishandle some ATM adaptation layer service data units.
- Message/streaming mode — In the message mode, ATM adaptation layer service data units are received in their entirety, whereas in streaming mode the ATM adaptation layer service data units are received in one or more increments, called AAL interface data units.
- Multiplexing — Multiple ATM adaptation layer type 3/4 connections can be multiplexed onto a single VC connection.

The SMDS, defined by Bellcore, is an early connectionless data service that can be implemented over ATM using ATM adaptation layer type 3/4.

ATM adaptation layer type 5 addresses Class C services, including signaling, and provides functions similar to ATM adaptation layer type 3/4. The major difference is that ATM adaptation layer type 5 has less overhead and does not support the multiplexing or streaming mode.

Frame relay is an early connection-oriented data service that can be implemented over ATM using ATM adaptation layer type 5. In this case, ATM adaptation layer provides an underlying connectivity, and the ATM network provides channels for the frame relay service.

Signaling. The new services envisioned for BISDN and ATM's ability to flexibly allocate and de-allocate bandwidth increase the functions required for signaling in the BISDN network.[5] BISDN signaling is expected to provide call and connection control for conversational services with multimedia, multiparty connections (e.g., videoconferencing), data retrieval services (e.g., document browsing), messaging services for voice-mail systems, and distribution services (e.g., television services with or without user control).

Current signaling protocols are based on a one-to-one mapping between calls and connections. The BISDN multimedia services require more than one connection to be associated with a call; they also require the capability to add, delete, and modify connections during a call. ITU-T SG 11 (formerly called CCITT SG XI) is now extending the existing ISDN signaling protocols (Q.931 and ISUP) to BISDN (Q.93B and B-ISUP).

Call Control and Connection Control. Some BISDN services, such as multimedia, would benefit from the ability to control the connections associated with a call independently of controlling the call itself.[6] Call control is responsible for establishing and releasing calls, managing call parameters (reference value, permissions, call state) and call configuration, and interacting with connection servers in performing user-server negotiations. Connection control establishes connections that meet the quality of service, that is, performance attributes, requested by the user; modifies the attributes of active connections (e.g., changing the quality of service, adding an endpoint); manages information regarding connection state and configuration; and releases connections. The ability to control calls and connections independently would allow independent routing of connections within a call; addition, modification, and deletion of connections during a call; and modular implementation of services.

Network Management and OAM

Management of BISDNs is largely based on telecommunication management network concepts of fault, performance, configuration, security, and accounting management. Functions defined for these management categories are accessible at the network management systems interfaces (i.e., Q3) of BISDN equipment. ATM's flexibility in connection bandwidth allocation and routing provides some unique advantages to BISDN management.

The BISDN network management architecture can be depicted in terms of the BISDN Protocol Reference Model, discussed earlier. Associated with each layer is an OAM channel, or flow mechanism. These F flows provide layer management functions (see Figure 2). The F1, F2, and F3 flows correspond to the regenerator section,

Table I. OAM Functions in the ATM Layer

Fault management	Performance management	Configuration management
Alarm surveillance (e.g., AIS, RDI)	Bit-error performance monitoring and reporting	Network element configuration status
Connectivity verification	Cell loss/mis-insertion monitoring and reporting	Network topology status and update
Continuity monitoring	Cell transfer delay measurement and reporting	
Loopback		

multiplex section, and path of the physical layer, respectively. These OAM flow mechanisms are specific to the physical layer of the transmission system employed in the BISDN (e.g., a SONET/SDH lightwave system). F4 and F5 flows correspond to the ATM layer VP and VC connection levels, respectively. The F4/F5 flows are independent of the physical layers used in the BISDN. They provide the end-to-end or segment OAM and management capabilities of associated VP/VC connection-based services. The management plane coordinates the layer management capabilities and, in essence, provides the Q3 interface.

The F4/F5 flows are among the more interesting aspects of the BISDN OAM and network management.[7] These flows use dedicated cells to carry out fault management, performance management, and configuration management of the ATM layer. Table I presents a short list of these functions. Although these functions currently are defined only for point-to-point connections, in the future, similar OAM capabilities will be established for multipoint connections. For the most part, these functions can support non-intrusive in-service monitoring and testing, as well as out-of-service testing for established virtual connections, or specified segments of those connections. Most importantly, these ATM OAM channels are handled in the same manner as the user channels. There, OAM connections can be established and released on demand; channel bandwidth can be either asymmetric or symmetric, and can vary according to need; and routing to (distributed) OAM processors can be made more flexible. These features of ATM OAM add flexibility to the BISDN network management plan, greatly enhancing such capabilities as network protection and restoration, service management, and end-customer control.

Traffic Management and Congestion Control

The BISDN networks must be designed to transport a wide variety of traffic classes whose requirements for transfer capacity and network performance vary broadly. Traffic management and congestion control capabilities that are adaptable to a wide variety of traffic types are critical to achieving BISDN network performance objectives. To date, this topic has posed a significant challenge in the design of the BISDN network.

The generic functions that have been defined for managing and controlling traffic and congestion in BISDN (see Figure 3) include:

- Network resource management — Provisioning capabilities (e.g., reservation protocols) to allocate network resources for separate traffic flows based on service characteristics.
- Connection admission control — Procedures carried out by the network, when it establishes or renegotiates a call, to determine whether a VC or a VP connection will be accepted or rejected. These procedures use call source traffic descriptors, quality-of-service parameters (discussed in the following section), and current network resource allocation.
- Priority control — The user may generate different priority traffic flows by invoking CLP in the cell header. If the network is congested, or becoming congested, it may discard low-priority traffic to protect network performance for high-priority traffic.
- Usage/network parameter control — Actions taken by the network to monitor and control traffic coming into the user access and network access, respectively. These include detecting malicious, as well as unintentional, violation of negotiated parameters for the call, and policing actions (such as cell discard, cell tagging for lowering priority, or rescheduling) in the case of violations.
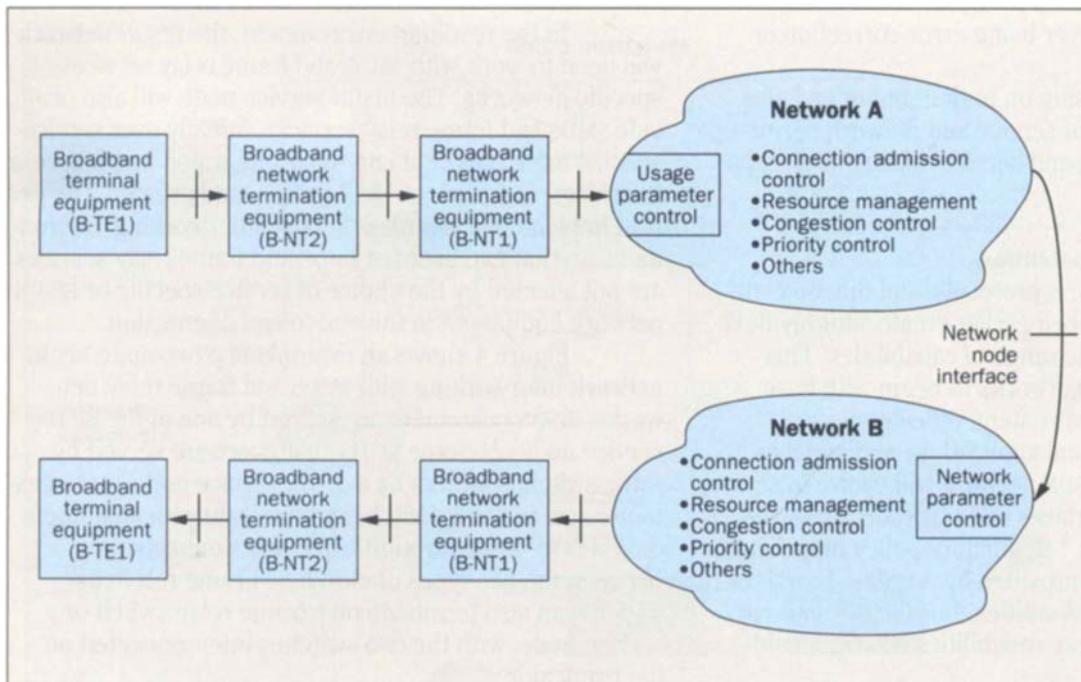
Figure 3. Traffic management reference configuration for the BISDN network.

Other traffic management and control functions, such as congestion detection and notification, and definition of traffic parameters, are being actively studied.

Network Performance and Quality of Service

The user traffic parameters discussed earlier also include quality-of-service parameters. These parameters, which comprise cell loss ratio, cell transfer delay, cell delay variation, and cell block error ratio, are measured on an end-to-end connection basis. They represent the overall network performance with respect to a given connection.

A user may request, at most, two different quality-of-service classes for a single virtual connection, which may only differ with respect to the *cell loss ratio* objectives. The CLP bit in the ATM cell header can accommodate two cell loss ratio objectives for a given virtual connection. For example, on a variable-bit-rate connection carrying video, the cells carrying synchronization information must be protected. Within the connection, these cells can be designated to have a higher CLP (i.e., $CLP = 0$) than cells that do not contain synchronization information (i.e., $CLP = 1$). The CLP bit can also be used inside the network to provide different qualities of service for different connections. Here, for example, the

network can use the CLP bit to protect signaling connections against network congestion.

The *cell transfer delay* is the time it takes for a cell to travel from one end of the connection to the other. This time can vary. Variability can be added at the customer premises or in the network (e.g., due to variable queueing delays in intermediate ATM multiplexers and switches). Services over virtual connections can be distinguished and controlled by their tolerance to delay and delay variation. For example, a circuit emulation service would have tight bounds on delay and delay variation, and the ATM virtual connection supporting this service would have to meet these delay bounds. On the other hand, a data service would have less stringent delay requirements. Cell transfer delay can be accommodated by segregating traffic classes in switch node queues with constant-bit-rate traffic receiving higher priority.

The *cell block error ratio* is the ratio, for a given time interval, of the total number of cell blocks that contain an error to the total number of cell blocks transferred in the end-to-end connection. Using this example, a circuit emulation service may be able to tolerate a relatively high errored block ratio compared to a data service, where a very low errored block ratio is mandatory. In the latter case, received bit errors are handled in the

service ATM adaptation layer using error correction or retransmission.

Work is progressing on both defining and also measuring other quality-of-service and network performance parameters, to expand service management capabilities in the BISDN.

Evolution of Network Capabilities

The basic structure, protocols, and functions of the BISDN network described earlier create a highly flexible framework with a wide range of capabilities. This framework allows BISDN networks to begin with basic services and evolve in a consistent, efficient manner along the path of broadband applications and services.[8] The capabilities of the BISDN network will evolve in several areas, including interfaces and transport protocols, and signaling and control.[9] Regulatory policy may dictate that some capabilities be provided by regulated carriers and others by unregulated entities. For the purposes of this section, however, these capabilities will be considered together.

Interface and Transport Capabilities. The needs of the earliest broadband applications, such as wide-area networking and LAN interconnect, are being met today by the initial broadband services, such as frame relay and SMDS. While these may not fit the classic definition of BISDN services, they are sometimes carried over ATM backbone networks, and thus provide the first end-user interfaces for ATM backbone networks. Frame relay and SMDS users typically have access lines whose rates range from 56 kbits/s to 45 Mbits/s.

As network interfaces and protocols evolve, networks will offer ATM services directly to end users. These early ATM services will complement frame relay and SMDS by providing access line rates in the range of 30 to 155 Mbits/s. This will permit the combination of frame relay, SMDS, and ATM services to cover a wide range of user needs.

Early ATM services will include a range of protocol adaptation functions for frame relay, SMDS, and other packet-based and circuit-based protocols.[10] Initially, ATM networks will transport these additional protocols through the network transparently. As the network evolves, more sophisticated internetworking functions will permit communications between network customers subscribing to different protocols, or even terminating on different types of networks.

In the resulting environment, the BISDN network will need to work with SMDS and frame relay service-specific networks. The BISDN service node will also provide SMDS and frame relay services, initially over service-specific interfaces, and later over integrated BISDN access interfaces, at speeds of 155 Mbits/s, and greater. The goal is to implement a seamless service interworking environment, so that end users of SMDS and frame relay services are not affected by the choice of service-specific or BISDN network equipment in the end-to-end connection.

Figure 4 shows an example of a two-node BISDN network interworking with SMDS and frame relay networks. BISDN customers are served by one of the BISDN service nodes.[11] Some SMDS customers are served by SMDS switches, others by a BISDN service node. SMDS customers on an SMDS switch communicate with customers on a service node through trunk-side connections between the two types of switches. Frame relay customers can also terminate on a frame relay switch or a service node, with the two switches interconnected on the trunk side.

Network capabilities are continually evolving. Direct ATM access line rates will extend to both higher and lower speeds. The extension to higher speeds, such as 622 Mbits/s, and beyond, will be driven by the continuing growth in broadband traffic and the amount of bandwidth required for end-user applications. The extension to lower speeds will be driven by the move to ATM as a native protocol for end-user equipment, and also by the desire to eliminate the expense of protocol interworking devices.

While some early applications may consist of traffic that is primarily bursty (such as LAN interconnect traffic), or primarily continuous bit rate (such as transparent circuit emulation), many applications will quickly move to a mixture of bursty and continuous-bit-rate traffic. ATM networks will be required to handle this mixture of traffic efficiently, maintaining low delay for continuous-bit-rate traffic, while allowing high use for bursty traffic, with very low cell loss rates in both cases. Developers of ATM networks must pay careful attention to quality-of-service parameters, congestion control, and BISDN switch fabric design.

Signaling and Control Capabilities. The BISDN signaling and control framework supports a wide range of advanced and sophisticated services. The network will deploy these signaling and control capabilities as end-
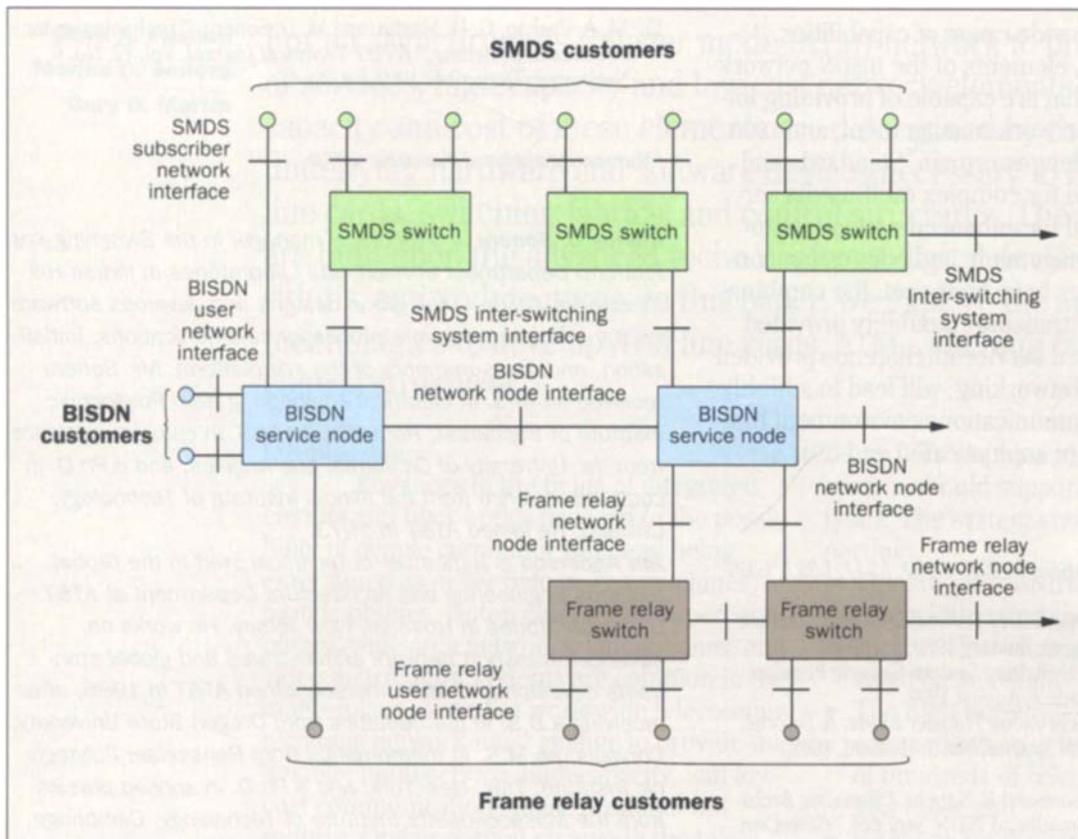
Figure 4. A broadband network with BISDN, SMDS, and frame relay. BISDN customers are served by a BISDN service node. Some SMDS customers are served by SMDS switches, while others are served by a BISDN service node. Frame relay customers can also terminate on a frame relay switch or a service node. Trunk connections between the BISDN, SMDS, and frame relay nodes permit interworking between customers served by different types of nodes.

user applications demand them and as they become better defined in standards. Initial ATM-based signaling and control capabilities will combine connectionless services provided by SMDS and connection-based services provided by frame relay and ATM. Also known as permanent virtual circuits, the earliest connection-oriented ATM services will be based on semi-permanent provisioned connections, changed through service orders rather than through signaling messages.

Early ATM services will be accompanied by customer network management capabilities. These end-user network management capabilities will use the network's internal telecommunication management network-based functions, but they will provide additional interfaces directly to end users. These interfaces will allow end users to monitor the status of their subnetworks and to request provisioning changes electronically. End-user network management capabilities will be part of standard protocol interfaces, such as Simple Network Management Protocol, rather than broadband signaling protocols.

At a later stage, switched virtual channel capabilities and services will be added to ATM networks, permitting end users to set up and tear down virtual circuit connections using signaling messages in real time. Initial switched virtual channel capabilities, based on the ITU-T Q.93B and B-ISUP specifications, will support basic point-to-point calls. In later phases, more sophisticated capabilities will be introduced to take advantage of ATM's flexibility. These capabilities include dynamically configurable, point-to-multipoint (or multicast) connections, add/drop, and multimedia calls with multiple underlying connections. Sophisticated calling features similar to those now associated with voice calls will also be introduced. Advanced Intelligent Network® services, including flexible service creation and execution, will be extended to cover BISDN calls.

Conclusions

The basic structure, protocols, and functions of the BISDN network have been defined to create a highly

flexible framework with a wide range of capabilities. Based on this framework, elements of the BISDN network are now being deployed that are capable of providing initial broadband services, network management, and OAM. However, significant challenges remain. Standards and strategies must be defined for complex multimedia services and applications, and for enhanced network management, OAM, traffic management, and congestion control. Once these challenges have been met, the combination of the bandwidth and transport flexibility provided by ATM, and the customized service intelligence provided by Advanced Intelligent Networking, will lead to a highly flexible, customizable communications environment that can support a wide range of sophisticated end-user services and obligations.

References

1. *CCITT BISDN Recommendations I.121R, I.150, I.311, I.321, I.327, I.361, I.362, I.371, I.610.*
2. W. R. Byrne, et al., "Broadband ISDN Technologies and Architectures," *IEEE Network Magazine,* January 1989, pp. 23-28.
3. Bellcore, "Broadband ISDN Switching System Generic Requirements," TA-NWT-001110, Issue 1, August 1992.
4. J.-Y. Le Boudec, "The Asynchronous Transfer Mode: A Tutorial," *Computer Networks and ISDN Systems,* North-Holland, 1992, pp. 279-309.
5. M. Veeraraghavan, D. M. Rouse and R. Kapoor, "Signaling Architectures and Protocols for Broadband ISDN Services," *GlobeCom 1992 Proceedings,* Orlando, December 6-9, 1992.
6. J. I. Cochrane and M. D. Soneru, "Switching, Signaling, and Control for Multimedia Services," *ComForum'90 Proceedings,* December 1990.
7. J. Anderson and M D. Nguyen, "ATM Operations and Maintenance Implementation Aspects," *IEEE Communications,* Vol. 29, No. 9, September 1991, pp. 79-81.
8. M. DeMaio and H. J. Kafka, "Broadband Services and Applications," *AT&T Technical Journal,* Vol. 72, No. 6, November/December 1993, pp. 23-28.
9. M. DeMaio, et al., "AT&T Network Systems Service Net-2000 Broadband Networking Plans," *GlobeCom 1992 Proceedings,* Orlando, December 6-9, 1992.
10. W. R. Byrne, et al., "Evolution of MANs to Broadband ISDN," *IEEE Communications Magazine,* Vol 29, No. 1, January 1991, pp. 69-82.
11. M. A. Pashan, G. D. Martin, and M. D. Soneru, "Technologies for Broadband Switching," *AT&T Technical Journal,* Vol. 72, No. 6, November/December 1993, pp. 39-47.

Marius D. Soneru is a technical manager in the Switching ATM Platform Department at AT&T Bell Laboratories in Indian Hill (Naperville), Illinois. His group designs and develops software for the infrastructure, interprocessor communications, initialization, and measurements of the ATM platform. Mr. Soneru received an M.S. in electrical engineering from Polytechnic Institute of Bucharest, Romania, an M.S. in computer science from the University of California, Los Angeles, and a Ph.D. in computer science from the Illinois Institute of Technology, Chicago. He joined AT&T in 1973.

Jon Anderson is a member of technical staff in the Global Systems Engineering and Architecture Department at AT&T Bell Laboratories in Holmdel, New Jersey. He works on advanced transport network architectures and global standards development. Mr. Anderson joined AT&T in 1986, after receiving a B.S. in mathematics from Oregon State University, Corvallis; an M.S. in mathematics from Rensselaer Polytechnic Institute, Troy, New York; and a Ph.D. in applied physics from the Massachusetts Institute of Technology, Cambridge.

Henry J. Kafka is a technical manager in the New Business Opportunities (NBO) and Visual Communications Department at AT&T Bell Laboratories in Indian Hill (Naperville), Illinois. He is responsible for the systems engineering and architecture of interactive video products, and their applications to emerging broadband ATM technology. Mr. Kafka received a B.S. from Northwestern University, Evanston, Illinois, and an M.S. from the University of Illinois, Urbana-Champaign, both in electrical engineering. He joined AT&T in 1979.