Network Troubleshooting

by Othmar Kyas

10 ATM

An Agilent Technologies Publication

Agilent Technologies

ATM 10

"To err is human, but to really foul things up requires a computer."

ANONYMOUS

10.1 ATM: Specification and Implementation

Since the early 1980s, data communications have been divided into two separate areas with little in common: local-area and wide-area networking. For technical reasons, data transport methods have been fundamentally different in these two areas. In wide-area networks, data communication has been connection-oriented: before the first bit of user data is transmitted, a signaling process takes place in which a dedicated connection to a given remote station is set up. Local-area networks, by contrast, have used "connectionless" broadcast transmissions: every data packet is sent out over a medium shared among all stations without waiting for acknowledgment. It is the receiver's job to detect packets in the data stream that have its address as the destination and process them.

Asynchronous Transfer Mode (ATM) was originally conceived for wide-area data communication, but was soon adapted for local-area networks. ATM creates a unified system that does away with the historical distinction between local- and wide-area data communication techniques. In both ATM LANs and WANs, data is transported by means of switches according to principles that had been customary only in wide-area communications, such as telephony. Every ATM end system is connected to a dedicated switch port. Every data packet, or "cell", that is addressed to a given station is delivered by an ATM switch to a corresponding switch port. Consequently, all data packets no longer travel over the same shared broadcast medium. Pick-up and delivery of cells (data packets) is managed entirely by an ATM switch. The key advantage of this technique is that every station connected to an ATM switch is guaranteed a certain bandwidth on its port regardless of how many other nodes are connected. This is possible because the switch's internal data throughput is many times higher than the bandwidth of any switch port. In conventional LANs, the average bandwidth available to any node is inversely proportional to the total number of active nodes. Furthermore, the switching principle is combined in ATM networks with connection-oriented data transmission. The traffic parameters necessary for a given service can be negotiated during the signaling or provisioning procedure required for every connection. For example, an end system may request a transmission path to a destination with a certain bandwidth and a certain maximum transmission delay. If the switches along the transmission route have the capacity necessary to grant the requested connection, then these communication parameters are guaranteed for the duration of the connection and appropriate network resources are reserved.

10.1.1 ATM in Homogeneous Private Networks

If a network is built using only ATM components, then data communication can take place using classic ATM transmission techniques. Because ATM is connection-oriented, the data virtual channel used to transport actual user data must be defined before the data transmission by means of a signaling or provisioned connection. With signaling, special communication protocols govern the negotiation of the user data connection parameters (bandwidth, delay, routing, etc.). Data transmission takes place until the connection is terminated by an appropriate command on the signaling virtual channel. Different signaling protocols are used at the interface between the end system and the ATM switch (the user-to-network interface, or UNI) and at the interfaces between ATM switches (network-to-network interfaces, or NNI).

In a private network, NNIs (that is, interfaces between two ATM switches) use the Private Network to Network Interface (PNNI) protocol. It makes no difference whether the private network is a LAN or a WAN. In public ATM WANs, however, the NNI protocol commonly used is the Broadband ISDN User Part (B-ISUP) protocol. The reason for the use of different protocols in the private and public spheres is that the ITU standardization body has left the internal operating aspects of public data networks in the control of national telecommunications companies. For example, B-ISUP does not specify how routing information is propagated or how topology detection mechanisms should be implemented. In private networks, however, it is desirable for all data communication mechanisms to be precisely defined so that communication works automatically without custom additions. For this reason, PNNI also specifies appropriate routing mechanisms in addition to the signaling processes. Because PNNI has the basic capabilities necessary for use in public WANs, and because many ATM system manufacturers provide an implementation of PNNI, but not B-ISUP, PNNI is sometimes also used in public networks, especially in North America. In Europe and Asia, however, almost all public networks use B-ISUP.

At the interface between an ATM switch and an ATM end system, that is, the UNI, connection setup is governed by the ITU-T Recommendation Q.2931 (and Q.2971 for point-to-multi-point connections), or by one of the ATM Forum's UNI signaling protocols, UNI 3.0, 3.1 or 4.0.

When an ATM end system is installed or becomes active, it must first register with its assigned ATM switch. Information exchanged in this process includes the network address of the switch, the user address of the ATM end system, and the service characteristics of the ATM end system. All of this information is stored in the ATM switch in a defined format as the Management Information Base (MIB). In order for this registration to be carried out automatically, a separate protocol for the exchange of MIB parameters was created, called the Integrated Local Management Interface (ILMI).

Because most network applications today build on the Internet Protocol (IP), and few native ATM applications are available to date, ATM is also able to serve as a fully transparent transport layer for IP. However, the Internet Protocol was originally designed for connectionless Ethernet networks, and therefore uses the broadcast principle for a number of functions. In the Classical IP over ATM protocol defined for ATM networks, broadcasts cannot easily be translated to connection-oriented ATM; this is possible only with the LAN Emulation (LANE) protocol. Nonetheless, IP can still be transported transparently over pure ATM networks. Address resolution is performed by specially defined ATM Address Resolution Protocol (ATMARP) and Inverse ATM Address Resolution Protocol (InATMARP) functions; other broadcast types, however, are not supported.

10.1.2 ATM in Heterogeneous LAN Environments

ATM is used not only in new networks planned exclusively on the basis of ATM, but often in combination with evolved existing network structures. In order to integrate ATM smoothly in such heterogeneous networks, the protocols LANE and Multiprotocol over ATM (MPOA) have been defined for ATM. These protocols permit complete, transparent interoperability between ATM networks and end systems, on the one hand, and conventional LAN topologies such as Ethernet and Token Ring, on the other. When the LANE protocol is used in an ATM network segment, every ATM end system in this network can communicate directly with every other Ethernet or Token-Ring station connected through an ATM router. Furthermore, selected stations in the LAN can be assigned to an ATM LANE workgroup. That is, end systems can be grouped regardless of their location or network interface into “virtual LANs” (VLANs). MPOA provides transparent interconnection of several ATM LANE segments.

Note that, because LANE is a Layer 2 (link layer) protocol, emulated LANs (ELANs) are either Ethernet or Token-Ring based, not a mixture of both. Layer 3 (network layer) routing is necessary to provide communication between Ethernet and Token-Ring based ELANs via, for example, the MPOA protocol.

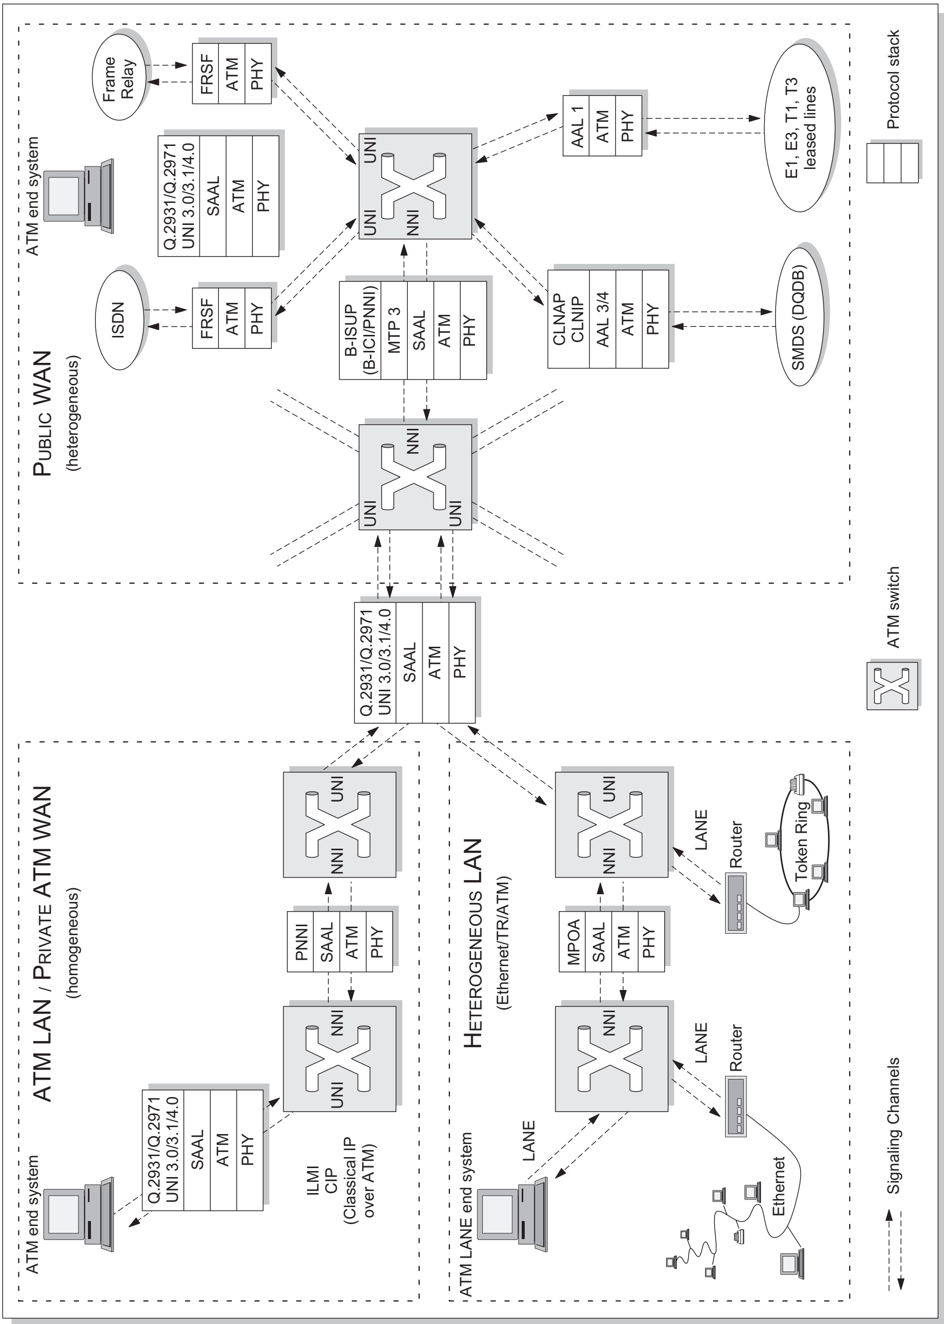

Figure 10.1 ATM in LAN and WAN: General view

10.1.3 ATM in Public Wide-Area Networks

When used for data communication in public wide-area networks, ATM can be transparently connected with Frame Relay and Switched Multimegabit Data Services (SMDS) networks, the latter being obsolescent. Transparent connection in this context means that not only user data but also protocol information is extracted from the foreign network topology and interpreted. In the case of Frame Relay, for example, the ATM network interprets and understands the Frame Relay Forward Explicit Congestion Notification (FECN) bit and the Discard Eligible (DE) bit, and represents them using the corresponding ATM protocol parameters (EFCI and CLP). In connection with Frame Relay networks, ATM emulates the Frame Relay UNI by means of the Frame Relay Service Function (FRSF). Because Frame Relay is also connection-oriented, it can be emulated relatively easily in ATM. Connecting SMDS networks to ATM is somewhat more complicated, because these networks, like conventional LANs, are based on the principle of connectionless communication. As in the case of LAN emulation, the connectionless transmission principle of SMDS must be translated to ATM's connection-oriented structure. This is done by means of the two protocols: Connectionless Network Interface Protocol (CLNIP) and Connectionless Network Access Protocol (CLNAP).

Another way to transport data over an ATM Network is to use the ATM Circuit Emulation Service (CES). With CES, the ATM network becomes a conduit for circuits such as E1, T1, E3, T3, etc., so any service that can be transported over these circuits is, by default also transported over ATM. Permanent virtual circuits are frequently used for CES and special interface equipment is needed to implement CES. More details are given later in the section dealing with the ATM Adaptation Layer (10.1.8).

As mentioned previously, the B-ISUP communications protocol developed by the ITU-T is used at the NNI between two ATM switches in a public network. The Broadband Inter-Carrier Interface (B-ICI) is an enhanced network-to-network interface protocol developed by the ATM Forum for connecting the networks of different carriers (communications service providers), and is used mainly in North America. A number of telecommunications providers use PNNI between their public ATM switches, however, for the sake of greater simplicity in operation.

10.1.4 Asynchronous Transfer Mode: ATM

ATM is defined as a packet-switched, connection-oriented data communication technique based on asynchronous time-division multiplexing and data packets of a fixed length. ATM data packets are called cells because of their constant

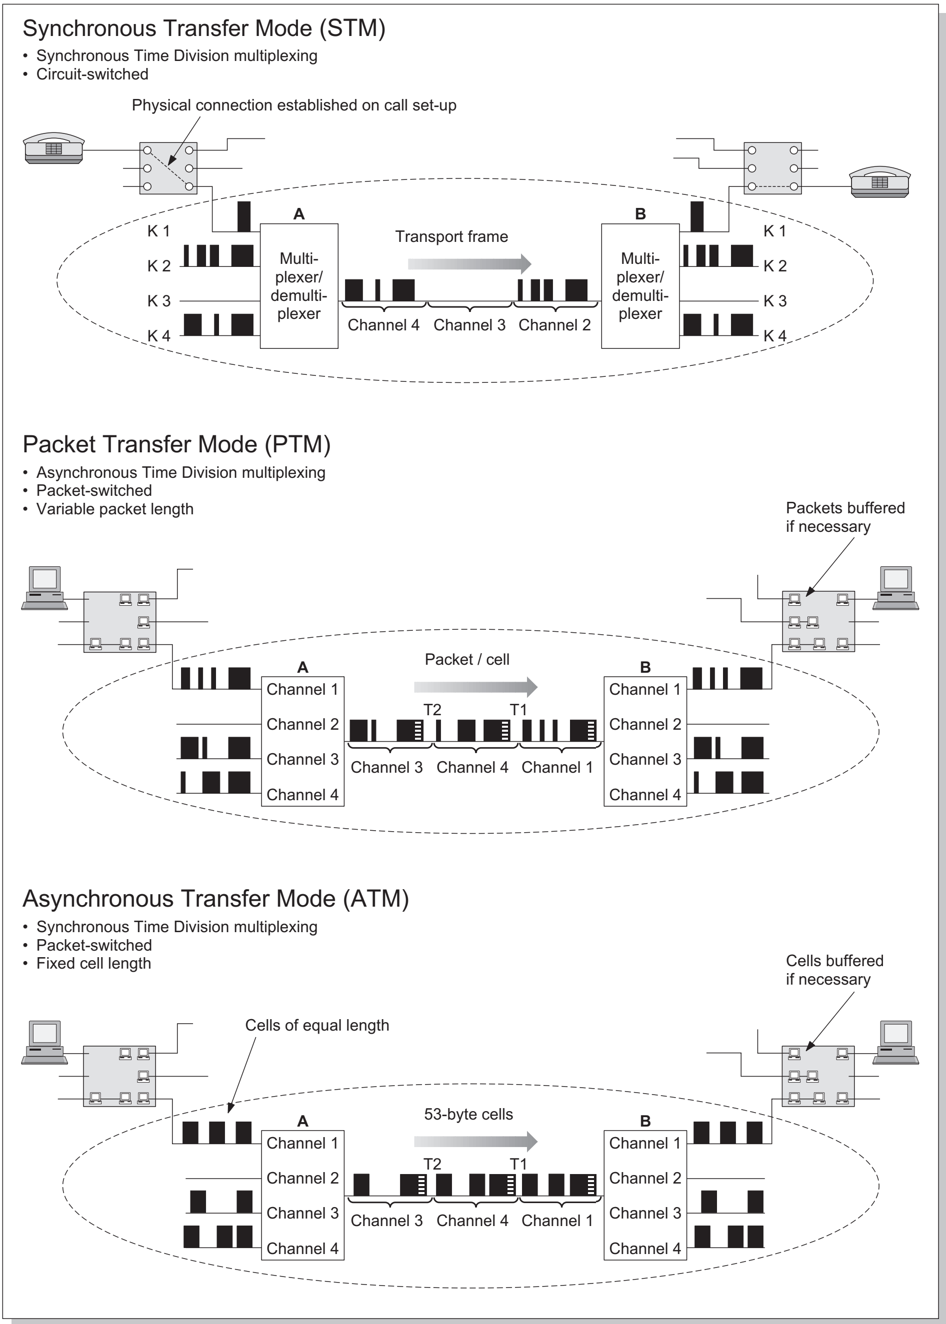

Figure 10.2 Principles of ATM

length of 53 bytes. Five of the 53 bytes form the cell header, which contains the Virtual Channel Identifier (VCI) and the Virtual Path Identifier (VPI). Note that ATM is “asynchronous” because cells are sent asynchronously, not because the octets (bytes) are asynchronous, as in RS-232C. All network nodes are connected by one or more ATM switches that forward the cells toward their destination point. The fixed cell size makes it possible for the ATM switches to forward multiple cells simultaneously with high efficiency, so that a very high throughput is attained in comparison to conventional routers. The network nodes do not share a communication medium as LAN nodes do, but simply give up their cells to the local ATM switch, without having to deal with media access algorithms. Before an actual data transmission begins, however, the transmission path for the user data cells is defined and set up by a signaling or provisioning procedure. A “traffic contract” guarantees that the traffic parameters granted, such as the maximum transfer delay, bandwidth or cell loss ratio, will be provided for the duration of the connection. This feature makes it possible to transmit many high-quality multimedia communication services over ATM networks. If the requested traffic parameters cannot be met by the network at call setup time, the call is rejected by the network so that the quality of service of other network users is not compromised.

The fixed cell length of 53 bytes results from a necessary compromise between the requirements of analog voice (low bandwidth, delay sensitive) and digital data communication (often higher bandwidth and almost always delay insensitive). A cell length of 53 bytes is short enough to be well suited to the transmission of delay sensitive low bandwidth digital signals in ATM’s higher megabit/s and gigabit/s throughput ranges. The small cell size also permits quick and exact bandwidth allocations as well as parallel cell processing at high speed in ATM switches.

10.1.5 The ATM Layer Model

The transmission procedure commonly referred to as the ATM protocol actually consists of a number of protocol layers that build on one another, and map the various network services to ATM cells. Taken as a whole, this set of protocol layers is called the B-ISDN protocol reference model. This name is a reminder of the origins of the ATM specification, which has its roots in the ITU-T’s Broadband-ISDN project for a universal wide-area network. Since ATM technology in the data communications field has spread beyond the boundaries of conventional telecommunications, ATM has become accepted as the universal name for the B-ISDN layers as a whole, although in a strict sense ATM refers only to cell-based data transport, that is, the ATM layer within the B-ISDN model.

The logical B-ISDN network architecture was designed for four independent levels of communication, after the ISO OSI reference model (ITU-T Recommendation X.200). Whereas the OSI model provides for seven layers, however, the B-ISDN protocol reference model defines the Physical Layer, the ATM Layer, the ATM Adaptation Layer (AAL) and the application layer designated in the model as "Higher Layer Protocols".

The physical layer consists of two "sublayers", the Transmission Convergence (TC) sublayer and the Physical Medium Dependent (PMD) sublayer. The TC sublayer embeds the ATM layer cells in the transmission framework of the given transport medium. If ATM cells are transported over a 34 Mbit/s E3 link, for example, they must be fitted into the user data field of E3 frames. The transport medium could also be a Synchronous Digital Hierarchy/Synchronous Optical NETwork (SDH/SONET) or a Plesiochronous Digital Hierarchy (PDH) frame such as DS1, E1, E3, DS3 or E4. In the case of direct cell transfer over the physical medium with no intermediate transport frame ("cell-based physical layer"), this sublayer is not necessary (that is, it is "null").

The functions of the ATM layer are completely independent of the underlying physical layer. Their chief purpose is to bring the data received from the higher-order AAL to its destination. The 53-byte cells comprise the information units of the ATM layer. Each cell has an identification number in its header that assigns it to a certain connection. The cells belonging to various connections are multiplexed into a cell stream in each direction, unused bandwidth being filled with unassigned or idle cells. These cell streams are hierarchically structured in virtual channels (VC) and virtual paths (VP), which correspond to one or more virtual connections. A physical transport medium (such as an optical fiber) can transport a number of virtual connections. Each cell on the medium can be unambiguously attributed to a certain connection by the VPI and VCI in its header, comprising the identification number referred to earlier.

The AAL disassembles the higher-layer data streams into 48-byte information segments for transport in ATM cells (five of the cell's 53 bytes are used for header information, so that 48 bytes of user data are transported in each cell). At the receiving end, the original data streams must be reassembled from the individual ATM cells. The functions of the AAL are thus dependent to a great degree on the characteristics of the higher-order applications. For this reason, the AAL functions are performed by two sublayers, the Convergence Sublayer (CS) and the Segmentation and Reassembly (SAR) sublayer. To limit the number of different AAL implementations, four AAL types have been defined: AAL1, AAL2, AAL3/4 and AAL5. Each of these AAL types is defined for a certain class of applications. By far the most widely used AAL variant today is the simplest one to implement, AAL5.

The actual network services are transported on the basis of the appropriate AALs. The most common applications are:

- Ethernet, Token Ring (using LAN emulation)
- Classical IP over ATM
- MPEG Video (Moving Pictures Experts Group)
- Frame Relay
- Leased-Line Data Links
- Voice (Telephony)

10.1.6 The ATM Physical Layer

Three methods can be used to transport ATM cells: cell adaptation to the frame structure of the transport medium (such as PDH, SDH/SONET or DXI); transport in PLCP frames, or cell-based physical layer.

In order to continue using existing data communications infrastructures, techniques have been developed to transport ATM cells in the container frames of the most common communications interfaces. Cell transport in SDH containers or SONET Synchronous Payload Envelopes (SPEs) is the most widely used transport method for ATM cells today. Adaptation methods have also been developed to transport ATM cells in wide-area networks over the older but still widely used E1, E3, E4, T1 and T3 PDH interfaces. In this technique the cells are inserted directly into the user data field of the given transport frame, and the cell rate is adapted to the throughput capacity of PDH by inserting fillers, called “idle” cells. In North America, “unassigned” cells are often used (for most purposes idle and unassigned cells can be considered identical). Adaptation to the existing PDH frame is defined in ITU-T Recommendation G.804 for the following transfer rates: 1.544 Mbit/s, 2.048 Mbit/s, 6.312 Mbit/s, 34.360 Mbit/s, 97.728 Mbit/s and 139.264 Mbit/s. Cell transport over the T3 line type (44.736 Mbit/s), common in North America, was once defined only using the T3-PLCP frame format (Bellcore TR-TSV-000772, TR-TSV-000773), originally developed for metropolitan area networks. However a more efficient, direct mapping was introduced a few years ago and is now the preferred mapping, even though the PLCP mapping is still very common in older equipment.

In the cell-based physical layer, the cells are not inserted in an additional transport frame, but simply converted bit-for-bit into the given communication medium’s electrical or optical signals and transmitted. A different scrambling arrangement is performed from that used in other interfaces in order to provide sufficient clock transitions so that the clock timing can be recovered from the incoming bit stream; this “distributed sample scrambler” (DSS) uses a 31st order polynomial ($x^{31} + x^{28} + 1$) and is described in I.432.1. One advantage of a cell-based physical layer is the efficient utilization of bandwidth. Because the ATM cell is not inserted in another frame structure with its own overhead, the ratio of total overhead to user data is maintained at about 1:9 (5 bytes of header, 48 bytes of data). The disadvantage is that existing transport infrastructures can no longer be used. Another problem with direct cell transfer in wide-area networks is the need to transport monitoring and management information. In both classical PDH and modern SDH/SONET networks, the frames or containers in which data are transported also contain the monitoring and error handling information required by the switching equipment. For this reason, special operation control cells have been defined for the physical layer, called Physical Layer Operation and Maintenance (PL-OAM) cells. These cells transport information required for the monitoring and management of the cell stream.

Many of the existing infrastructures in wide-area networking are unable to support ATM cell-based physical layers, however. The use of this technique is therefore limited to local-area networks. In practice, transportation of ATM cells in SDH/SONET frames has also become accepted in LANs even though, as ATM

Transport frame	Communication medium	Data rates (Mbit/s)
SDH/SONET	Single-mode fiber	155, 622, 2460, 9960
SDH/SONET	Multimode fiber	155, 622
SDH/SONET	POF, HPCF	155
SDH	75Ω coaxial cable	155
SDH/SONET	UTP3, UTP5, 150Ω STP	12.96, 25.92, 51.84, 155
PDH	75Ω coaxial cable	2.048, 34.36, 44.73, 139.26
PDH	100Ω TP	1.544, 2.048
DXI (not cell, but AAL-5 or AAL-3/4 based)	V.35, EIA/TIA 449, HSSI	up to 52
Cell stream	Single-mode fiber	155, 622
Cell stream	Multimode fiber	155, 622
Cell stream	V.35, EIA/TIA 449, HSSI	up to 52
Cell stream	UTP3, 120Ω Cat.4, 150Ω STP	25.6
Cell stream	FDDI multimode/ single-mode (TAXI)(obsolete)	100

Figure 10.3 ATM transmission interfaces

transport containers, these frames with their relatively simple structures (no multiplexing hierarchies or spatial limitations) are actually superfluous. Cell-based physical layers are therefore currently only used in a few applications developed especially for this transport method, such as video.

10.1.6.1 ATM Data Rates

At present, a large number of transmission interfaces are defined both for LAN and WAN use, spanning a wide range of bandwidths and transport media. Data rates range from 1.544 Mbit/s to 9.9 Gbit/s or even higher. The differences between the LAN and WAN specifications are primarily associated with the transport medium. Whereas single-mode fiber optic and coaxial cable are the primary media for WAN interfaces, multimode fiber, plastic optical fiber (POF) and copper twisted pair are becoming increasingly common in local-area networking. Figure 10.3 lists the currently available transmission interfaces.

10.1.6.2 ATM in PDH Networks

ATM over DS1: 1.544 Mbit/s

The mapping of ATM cells to DS1 line framing is specified in ITU-T Recommendation G.804. Under this standard, ATM cells can be transported in a 24-frame multiframe (Extended Superframe, or ESF) with the cells occupying bits 2 to 193. The individual ATM cells are byte-aligned with the DS1 frame, but not frame-aligned. This means that an ATM cell can be split across two DS1 frames.

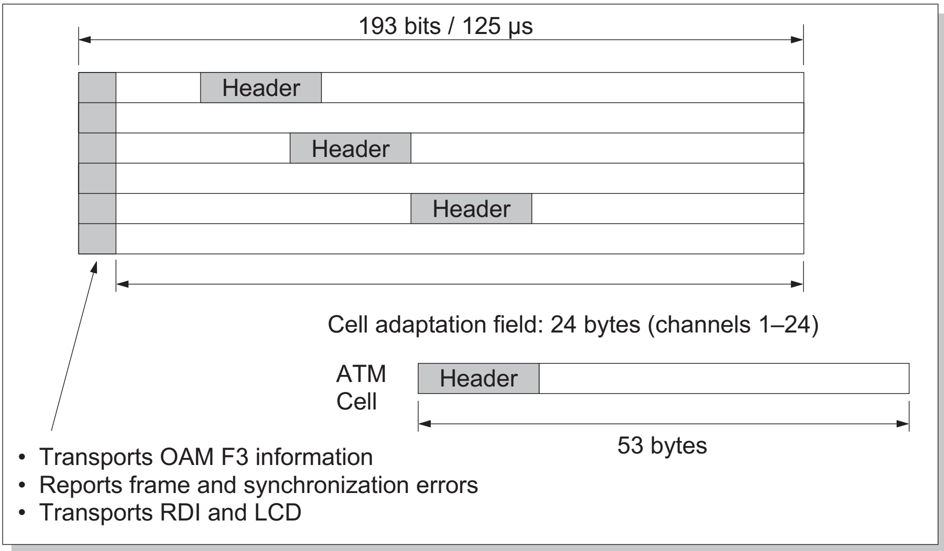

Figure 10.4 Direct cell adaptation to the DS1 frame format (G.804)

By inserting idle or unassigned cells when no valid ATM user data cells are available, the cell transfer rate is adapted to the user data bandwidth of DS1. Scrambling of the ATM cell's data field is optional, in contrast to the E1 ATM mapping. The self-synchronizing scrambling method is used with the generator polynomial x^{43}+1.

ATM over E1: 2.048 Mbit/s

Like DS1 mapping, the mapping of ATM cells to E1 frames is also described in ITU-T Recommendation G.804. The ATM cells can be transported in bits 9–128 and 137–256, which correspond to timeslots 1 through 15 and 17 through 31. The cell rate is adapted to the E1 frame payload bandwidth of 1.920 Mbit/s by means of idle cells when no ATM cells are queued for insertion. Once again, the ATM cells are byte-aligned with the E1 frame, but not frame-aligned.

The 48 data bytes of each ATM cell are scrambled before transport using self-synchronizing scrambling (SSS) with the generator polynomial x^{43}+1. This permits fast cell delineation, allowing the receiving station to recover quickly from a loss of cell delineation due to physical layer bit errors, for example.

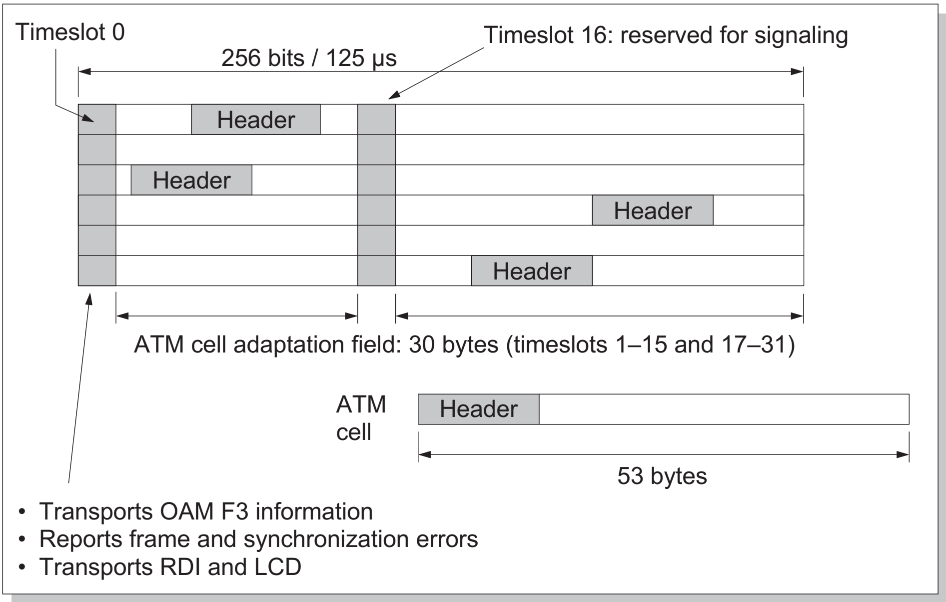

Figure 10.5 ATM cell adaptation to the E1 frame format

ATM over E3: 34.368 Mbit/s

The mapping of cells to the E3 frame format is described in ITU-T Recommendation G.804. The E3 frame format used is not the E3 frame format described in

G.751, however, but a modified frame format described in G.832. The mapping of ATM cells to the older G.751 frame structure is difficult to accomplish: the cells would have to be "nibble-aligned" because each G.751 subframe is an integer multiple of 4 rather than 8 bits. The newer G.832 frame consists of 537 bytes, of which 7 bytes are used for overhead information (see Figure 10.6). The remaining 530 user data bytes correspond exactly to the length of ten ATM cells, so that these can be byte- and frame-aligned (though the latter is not required). The ATM cell rate is adapted to the E3 user data rate by the insertion of idle cells when no ATM cells are queued for transport; in North America, unassigned cells are often used for this purpose—both idle and unassigned cells are discarded at the input to switches, etc. The 48-byte data field of the ATM cell (including idle/unassigned cells) is scrambled using self-synchronizing scrambling with the generator polynomial $x^{43}+1$.

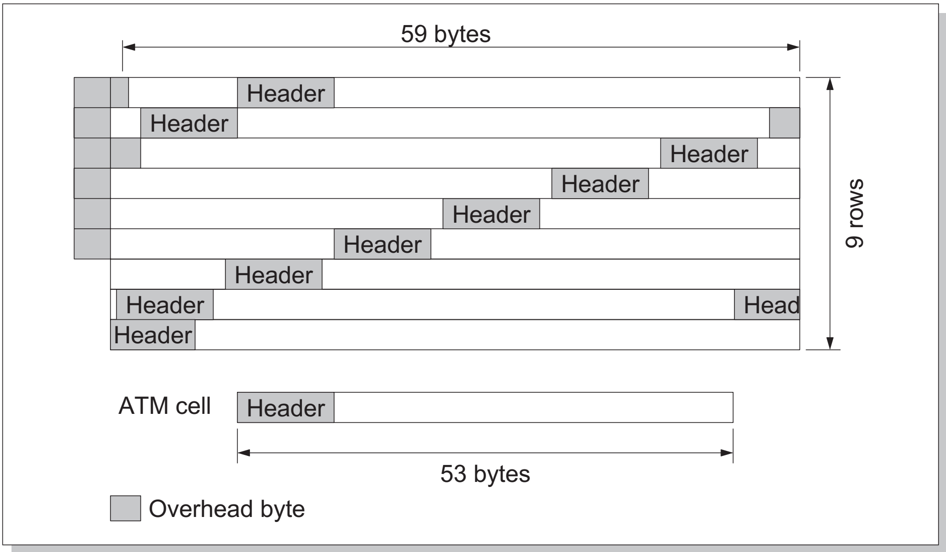

Figure 10.6 Cell adaptation to the E3 frame format

ATM over DS3: 44.736 Mbit/s

ATM cells can be transported over DS3 links using either the Physical Layer Convergence Protocol (PLCP) mapping found in older equipment or the direct cell mapping to the DS3 frame format preferred today. A DS3 PLCP frame consists of twelve rows of 57 bytes each. The last row contains an additional trailer of twelve or thirteen nibbles (half-bytes) to fill out the user data field of a DS3 multiframe. The DS3 PLCP frame has a transmission time of 125 µs, corresponding to a rate of 44.21 Mbit/s, and thus fits exactly in the user data field of a DS3 multiframe (Figure 10.7).

1	1	1	1	53 bytes	
A1	A2	P11	Z6	L2_PDU	
A1	A2	P10	Z5	L2_PDU	
A1	A2	P9	Z4	L2_PDU	
A1	A2	P8	Z3	L2_PDU	
A1	A2	P7	Z2	L2_PDU	
A1	A2	P6	Z1	L2_PDU	
A1	A2	P5	F1	L2_PDU	
A1	A2	P4	B1	L2_PDU	
A1	A2	P3	G1	L2_PDU	
A1	A2	P2	M2	L2_PDU	
A1	A2	P1	M1	L2_PDU	13 - 14 nibbles
A1	A2	P0	C1	L2_PDU	Trailer

125 μs

A1, A2....... Alignment bytes
P0 - P11.... Path overhead identification
C1............. Padding bit count
M2, M1...... SIP Layer 1 control information
G1...........PLCP path status
B1 ....Bit-interleaved parity (BIP-8)
F1 ....PLCP path user channel
Z1 – Z6 ...reserved for future use

Alignment bytes A1 and A2

Incoming signal or LSS code	Outgoing LSS code
connected	connected
rx_link_up	connected
rx_link_dn	rx_link_up
PLCP Out-of-Frame	rx_link_up
PLCP Loss-of-Frame	rx_link_dn

PLCP Path Status Byte

4 bits	1 bit	3 bits
Far-End Block Error	Yellow Signal	Link Status Signal

Path Overhead (POH) Identifiers

P8	00100000
P7	00011100
P6	00011001
P5	00010101
P4	00010000
P3	00001101
P2	00001000
P1	00000100
P0	00000001

Coding of the Link Status Field

LSS Code	LSS Name	Status
000	connected	Received Link Connected
011	rx_link_dn	Received Link Down, No Input, or Forced Down
110	rx_link_up	Received Link Up

Figure 10.7 The DS3 PLCP frame

The PLCP frame mapping of ATM cells is thus a two stage process that is complicated and also inefficient, because there is additional overhead associated with the PLCP frame structure; the cell payloads are not normally scrambled. The more efficient direct mapping of cells to the DS3 frame, sometimes referred to as "HEC" mapping, makes use of the same cell delineation process used for DS1, E1 and E3 mapping; cell payload scrambling again uses the self-synchronizing scrambling with the generator polynomial $x^{43}+1$.

ATM over E4: 139.264 Mbit/s

Similar to E3, ATM cell transport in the E4 transmission frame (ITU-T Recommendation G.804) does not use the E4 frame format described in G.751, but a modified frame described in G.832. The ATM cells are inserted byte-aligned in the E4 frame payload field of the 2,160 byte G.832-E4 frame. All other processing, such as the scrambling of the ATM cell payload and the cell rate adaptation, is performed as in the mapping of ATM cells to the E3 format.

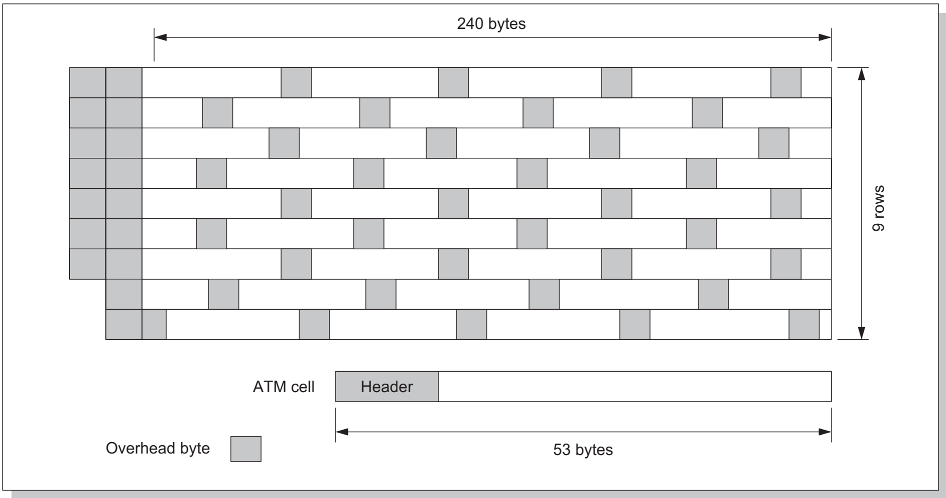

Figure 10.8 ATM cell adaptation to the G.832 E4 frame format

ATM over 6.312 Mbit/s and 97.728 Mbit/s

For the sake of thoroughness, we may mention at this point that ATM cell transport is also specified for the PDH bit rates 6.312 and 97.728 Mbit/s. These interfaces are only of regional importance, however, because they are almost never used outside Japan.

Bytes 1 - 5	ATM cell header: GFC=0, VPI=0, VCI=0, PTI=101, CLP=1, HEC
Byte 6	IMA label for compatibility with UNI 3.1 OAM cell format (OAM type field)
Byte 7	Cell ID: Bit 7: IMA OAM Cell Types (0: idle cell, 1: ICP cell) Link ID: Bits 6–5: not used; set to 0 Bits 4–0: logical ID for lines 0 to 31
Byte 8	IMA frame sequence number: 0–255, cyclical
Byte 9	ICP cell offset range (0 to M–1): indicates ICP cell position in the IMA frame
Byte 10	Link stuff indication: Bits 7–3: not used; set to 0 Bits 2–0: 111 = No imminent stuff event 100 = Stuff event in 4 ICP cell locations (optional) 011 = Stuff event in 3 ICP cell locations (optional) 010 = Stuff event in 2 ICP cell locations (optional) 001 = Stuff event at the next ICP cell location (mandatory) 000 = this is one of the two ICP cells compromising the stuff event (mandatory)
Byte 11	Status and control change indication: Bits 7–0: Status change indication: 0 to 255
Byte 12	IMA ID
Byte 13	Group status and control Bits 7–4: Group state 0000 = Group is active and ready to add further links 0001 = Group is inactive and not ready to add further links 0010 = Group start-up 0011 = Failure, protocol error 0100 = Failure, insufficient links 0101 = Failure, unsupported value of M 0110 = Failure, incompatible group symmetry 0111 = Failure, symmetry not supported Bits 3–2: Group symmetry mode 00 = Symmetrical configuration and operation 01 = Symmetrical configuration and asymmetrical operation (optional) 10 = Asymmetrical configuration and operation (optional) 11 = Reserved Bits 1–0: IMA Frame length (00: M=32, 01: M=64, 10: M=128, 11: M=256)
Byte 14	Transmit timing information Bits 7–6: not used; set to 0 Bit 5: Transmit clock mode (0: ITC mode; 1: CTC mode) Bits 4–0: Tx LID of the timing reference (0 to 31)
Byte 15	Tx test control Bits 7–6: not used; set to 0 Bit 5: Test link command (0: inactive; 1: active) Bits 4–0: Tx LID of test link (0 to 31)
Byte 16	Tx test pattern Bits 7–0: Tx test pattern (values 0–255)
Byte 17	Rx test pattern Bits 7–0: Rx test pattern (values 0–255)
Byte 18	Link 0 information Bits 7–5: Tx status Bits 4–2: Rx status Bits 1–0: Remote Defect Indicator
Bytes 19–49	Link 1–31 info status
Byte 50	not used; set to 6A hex (ITU I.432)
Byte 51	End-to-end channel (proprietary channel set to 0 if not used)
Bytes 52–53	CRC-10 (ITU I.432)

Figure 10.9 ICP cell format

Inverse Multiplexing for ATM (IMA)

ATM inverse multiplexing is the process of dividing an ATM cell stream into several component streams for transport, which are then reassembled into the original stream at the receiving end. This makes it possible to transport high bandwidths over several bundled data links of lower capacity. For example, bandwidths of 4 to 34 Mbit/s can be realized by bundling the appropriate number of E1 links, though it becomes uneconomical to consider bundles or "link groups" with more than about 6 - 8 links. The ATM cells are transported one at a time over each of the available lines in turn. For every M cell, a special IMA Control Protocol (ICP) cell is inserted. The ICP cell, together with the corresponding ATM user data cells, comprise an "IMA frame". Each ICP cell contains a Link Identifier (LID) indicating the individual line, an IMA frame sequence number, and various information fields used to monitor and synchronize the transmission (Figure 10.9).

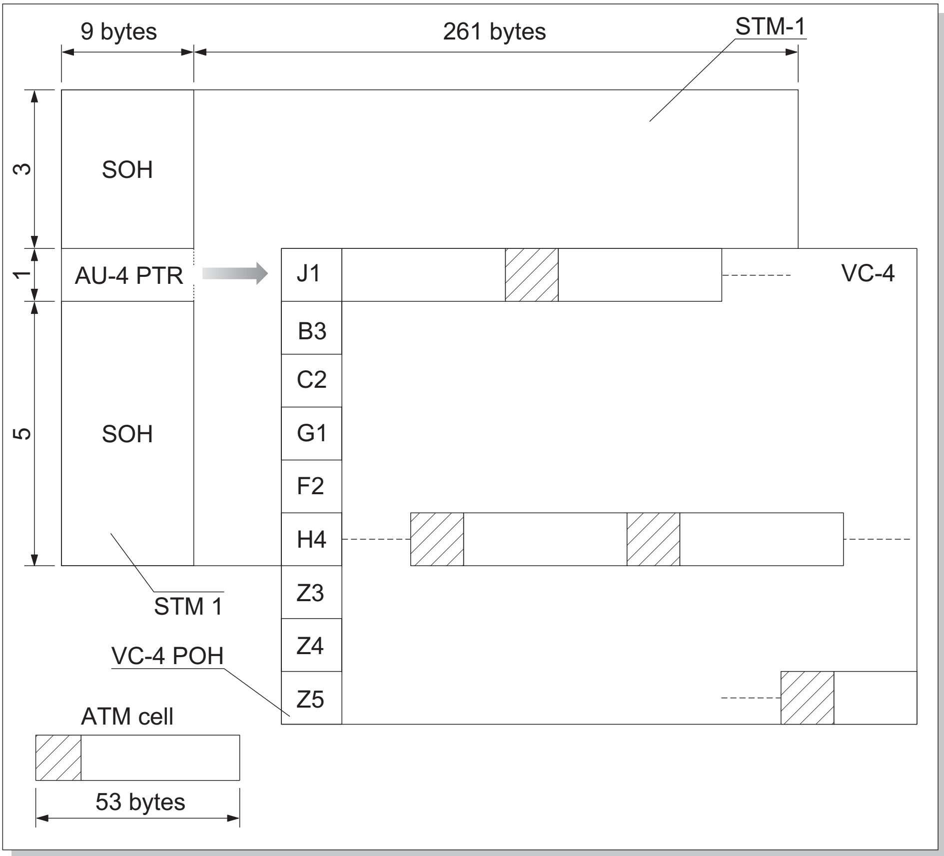

Figure 10.10 Transportation of ATM cells in the STM-1 transport module

10.1.6.3 ATM in SDH and SONET Networks

The transportation of ATM cells over SDH and SONET networks (ATM mapping) is specified in ITU-T Recommendation G.707 and ANSI T1.105 respectively; SONET is the North American equivalent of SDH and can be considered to be identical in most respects for equivalent interface rates. The cell stream is encapsulated with byte alignment in VC-x or in concatenated VC-xc containers (synchronous payload envelopes or SPEs in SONET). Because the container/SPE payload is not an integer multiple of 53 bytes, a cell can be split across two containers/SPEs. Before the ATM cells are inserted in the containers/SPEs for transport, the user data field of the cell is scrambled in order to facilitate delineation of the individual cell at the receiving station. The scrambling method used is self-synchronizing scrambling (SSS) with the generator polynomial x^{43}+1. If the cell transfer rate is different from the user data bandwidth of the

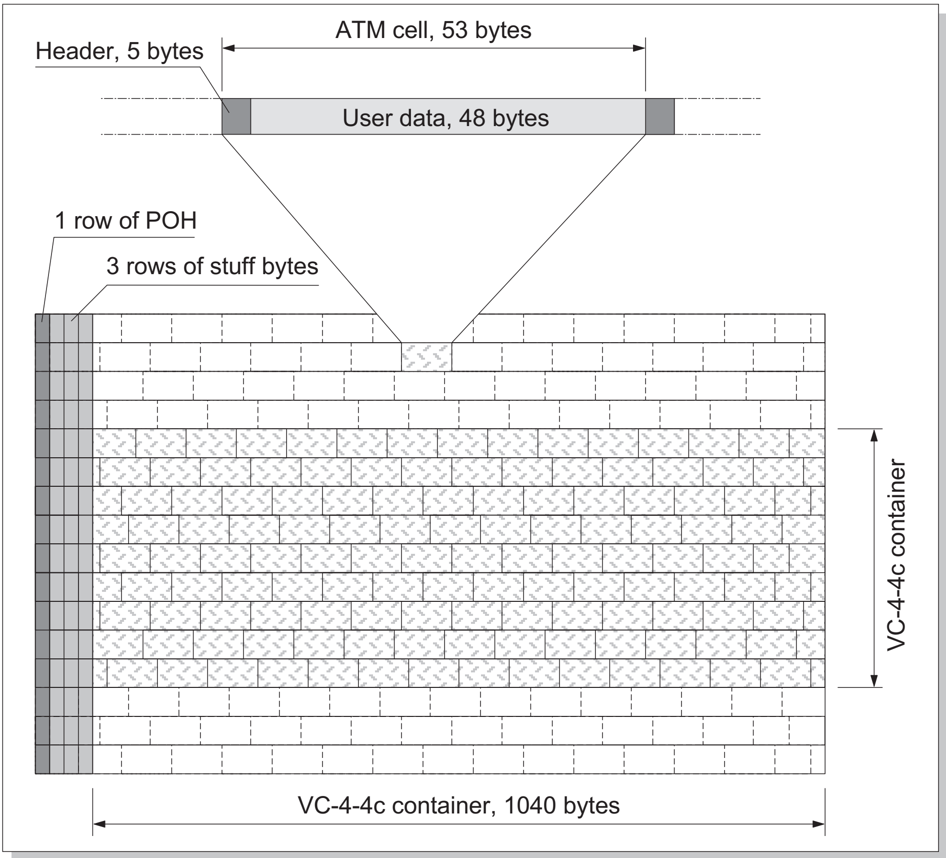

Figure 10.11 Transportation of ATM cells in the STM-4c transport module

SDH/SONET containers/SPEs, empty (idle/unassigned) cells are inserted. ATM cells that exceed the available bandwidth are discarded, with low priority cells being discarded preferentially. The resulting bit rate of the ATM cell stream is synchronous with that of the SDH container or SONET SPE. Note that, while Figure 10.10 shows cell mapping to an SDH STM-1 transport module, cell mapping to a SONET STS-3c transport system is essentially identical. Similarly, the mapping of cells to a SONET STS-12c transport system is identical to the mapping of cells to the SDH STM-4c transport module shown in Figure 10.11.

Byte	Function	Coding[1]
STM-1/STS-3c Section Transport Overhead		
A1, A2	Frame alignment	
C1	STM-1 identifier	
B1	Regenerator section error monitoring[2]	BIP-8
B2	Multiplexer section error monitoring	BIP-24
H1, H2	AU-4 pointer, path AIS[4]	A11 1s
H3	Action pointer	
K2 (bits 6–8)	Multiplexer section AIS / section FERF	111/110
Z2 (bits 18–24)	Multiplexer section error reporting (FEBE)[5]	B2 error count
VC-4 Path overhead		
J1	Path ID/verification	BIP-8
B3	Path error monitoring	ATM cell[3]
C2	Path signal level	B3 error count
G1 (bits 1–4)	Path error reporting (REI)	1
G1 (bits 5)	Path RDI	
FFS	Cell delineation supervision	FFS
FFS	Header error performance monitoring	FFS

1. Only the codes that are relevant for the monitoring function are listed.
2. The use of B1 for regenerator section error monitoring is optional.
3. The code for ATM cells is: C2 = 13hex
4. The use of H1 and H2 for path AIS is provisional.
5. The use of Z2 for multiplexer section error monitoring is provisional.

Figure 10.12 SDH/SONET overhead bytes in ATM cell transport

Alarm and Monitoring Signals

Two types of monitoring signal are defined for ATM cell transport over SDH/SONET networks: the Alarm Indication Signal (AIS) and the Remote Defect Indicator (RDI, formerly known as FERF for Far End Receive Failure) signal. An AIS is sent downstream to report an error. An RDI signal is sent upstream to report a receiving or transmission error. Both kinds of monitoring signal can be implemented using the Section/Transport Overhead (SOH/TOH) bytes of the STM-1/STS-3c frame, or with the Path Overhead (POH) of the VC4 container/SPE. Figure 10.12 shows the use of the corresponding SDH overhead bytes for ATM cell transport (the equivalent SONET overhead bytes are similar).

ATM over Single-Mode Fiber

ITU-T Recommendation G.957 defines six different single-mode fiber optic types that can be used to transport SDH/SONET in three different communication scenarios: in-house links, medium-range WAN links and long-range WAN links.

In-house links can attain a maximum range of 2 km. The optical transmitters may consist of light emitting diodes (LEDs) or multilongitudinal mode (MLM) lasers with a wavelength of 1,310 nm. The permissible loss is between 0 and 7 dB.

Medium-range WAN links span distances of up to 15 km. The optical transmitters may be either single longitudinal mode (SLM) or multilongitudinal mode (MLM) lasers with wavelengths of 1,310 or 1,550 nm. Permissible loss is between 0 and 12 dB.

For long-range WAN links of up to 40 km, lasers with a wavelength of 1,310 nm can be used. If high-power SLM (500 µW or –3 dBm) or MLM lasers are used at wavelengths of 1,550 nm, the range can be up to 80 km. Permissible loss is between 10 and 24 dB.

Parameter	Medium-range WAN links (15 km)	In-house links (2 km)	Units
Transmission characteristics			
Wavelength	1293 – 1334	1261 – 1360	nm
Spectral width: Max. RMS width	4	14.5 (MLM Laser35 (LEDs))	nm
Mean signal power	-15 to -8	-15 to -8	dBm
Minimum extinction rate	8.2	8.2	dB
Eye diagram	See T1.646	See T1.646	
Reception characteristics			
Minimum sensitivity	-28	-23	dBm
Minimum overload	-8	-8	dBm
Optical path power penalty	1	1	dB

Figure 10.13 Single-mode fiber-optic parameters for 622.08 Mbit/s interfaces (ITU-T Recommendation G.957)

Using single-mode fiber, the attainable ATM data rates are 155 Mbit/s with STM-1/OC-3, 622 Mbit/s with STM-4/OC-12, 2.4 Gbit/s with STM-16/OC-48, and 9.9 Gbit/s with STM-64/OC-192. In practice, ATM is transported over single-mode fiber optic links primarily in wide-area networks. Due to the steadily growing demands for range and capacity in local-area networks, however, these

Optical transmission parameters

Parameters		Units	Value
Baud rate		Mbaud	194.4
Baud rate tolerance		ppm	100
Optical signal power,	min.	dBm	-20
	max.	dBm	-14
Wavelength	min.	nm	1270
	max.	nm	1380
Spectral width		nm	< 200
Optical extinction rate		%	< 10
Pulse rise and fall time	min.	ns	0.6
(10%–90%)	max.	ns	2.5
Pulse overshoot		%	10
Pulse undershoot		%	10
Optical transmission jitter due to data		ns	± 0.5
Optical pulse phase shift		ns	± 0.5
Optical pulse jitter		ns	± 0.5

Eye diagram for output signal

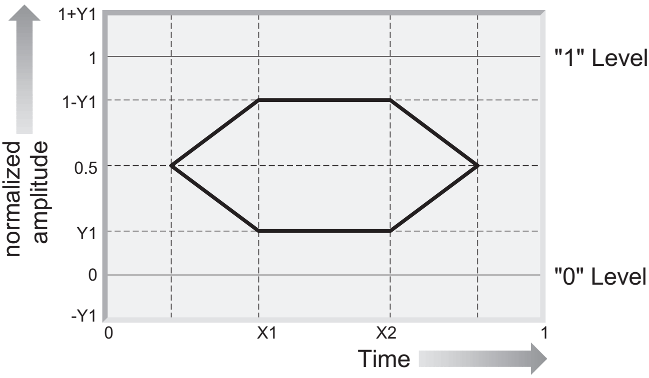

Optical reception parameters

Parameters		Units	Value
Baud rate		Mbaud	194.4
Baud rate tolerance		ppm	100
Optical signal power	min.	dBm	-29
	max.	dBm	-14
Wavelength	min.	nm	1270
	max.	nm	1380
Pulse rise and fall time	min.	ns	0.6
(10%-90%)	max.	ns	3.0
Optical transmission jitter due to data		ns	± 1.0
Optical pulse phase shift		ns	± 0.5
Optical pulse jitter		ns	± 0.5

Figure 10.14 Optical transmission and reception parameters for the 155 Mbit/s ATM multimode fiber-optic interface

media are found increasingly often in LANs as well. Figure 10.13 lists the single-mode fiber optic parameters as specified in ITU-T Recommendation G.957 for 622.08 Mbit/s interfaces.

ATM over Multimode Fiber

For the transmission of ATM cells in SDH/SONET frames in the LANs at data rates of 155 Mbit/s and 622 Mbit/s, the ATM Forum has also specified multimode fiber in place of the single-mode fiber prescribed in the ITU-T recommendations. The maximum range of 2 km is lower than the range for single-mode fiber, but this is seldom a problem in the LAN field, especially in indoor cabling structures with typical segment lengths of 100 to 1,000 m. The cost of the communication infrastructure can be significantly lowered, however, by avoiding expensive single-mode fiber and the high-quality laser sources it requires. In addition to ATM cell transmission over single-mode SDH/SONET WAN interfaces, multi-mode fiber also makes it possible to implement the STM-4c/OC-12 (622 Mbit/s), STM-1/OC-3 (155 Mbit/s) and STM-0/STS-1 (51.84 Mbit/s) interfaces in LANs.

The STM-1/OC-3 multimode fiber optic interface uses a wavelength of 1,300 nm over a 62.5/125 µm multimode fiber with a modal bandwidth of 500 MHz/km. 50/125 µm is also permissible. The data stream is 8B/10B-encoded for transmission, so that the physical medium's 194.4 Mbaud yields an effective data rate of 155.52 Mbit/s. Either SC or BFOC/2.5 (IEC 86B) connectors may be used. This specification, originally defined for UNI 3.0, has been rarely implemented, and is described here for the sake of completeness because it is still contained in the UNI 3.1 specification. Figure 10.14 lists the parameters for the multimode fiber optic interface.

The multimode fiber (MMF) interfaces are specified for 622 Mbit/s interfaces using LEDs or short-wave (SW) lasers. Both types of MMF interfaces support the following multimode fiber types:

- 62.5/125 µm IEC 793-2 Type A1b
- 50/125 µm IEC 793-2 Type A1a

The data stream is NRZ-encoded. LED-based MMFs can be used for segments of 300 to 500 m; SW MMFs have a range of about 300 m. Figure 10.15 lists the transmission and reception parameters for both MMF types at 622 Mbit/s.

ATM over Plastic Optical Fiber

As an alternative to both single-mode and multimode fiber optic media, the ATM Forum has also specified plastic optical fiber as a transport medium for ATM networks. For STM-1/OC-3 interfaces (155 Mbit/s) and segments of up to 100 m, Plastic Optical Fiber (POF) can be used, and segments of up to 100 m can use Hard Polymer Clad Fiber (HPCF). The POF medium is 1,000 µm multimode

LED-based MMF parameters			
Transmitter characteristics	62.5 µm MMF	50 µm MMF	Units
Wavelength	1270 to 1380	1270 to 1380	nm
Maximum spectral width	200	200	nm
Mean optical power	-20 to -14	-24 to -14	dBm
Minimum extinction rate	10	10	dB
Maximum rise and fall time (10% – 90%)	1.25	1.25	ns
Maximum systematic interface peak-to-peak jitter	0.4	0.4	ns
Maximum random interface peak-to-peak jitter	0.15	0.15	ns
Maximum overshoot	25	25	%
Receive characteristics			
Minimum sensitivity	-26	-26	dBm
Minimum overload	-14	-14	dBm
Maximum rise and fall time (10% – 90%)	1.6	1.6	ns
Maximum systematic interface peak-to-peak jitter	0.5	0.5	ns
Maximum random interface peak-to-peak jitter	0.15	0.15	ns
Minimum eye diagram aperture at receiver	0.31	0.31	ns
Software-based MMF parameters			
Transmission characteristics	62.5 µm MMF	50 µm MMF	Units
Wavelength	770 to 860	770 to 860	nm
Maximum spectral width	9	9	nm
Mean optical power	-10 to -4	-10 to -4	dBm
Minimum extinction rate	9	9	dB
Maximum rise and fall time (10% – 90%)	0.75	0.75	ns
Maximum systematic interface peak-to-peak jitter	0.35	0.35	ns
Maximum overshoot	25	25	%
Receiver characteristics			
Minimum sensitivity	-16	-16	dBm
Minimum overload	0	0	dBm
Maximum rise and fall time (10% – 90%)	1.2	1.2	ns
Maximum interface peak-to-peak jitter	0.55	0.55	ns
Minimum eye diagram aperture at receiver	0.31	0.31	ns

Figure 10.15 Optical transmission and reception parameters for the two 622 Mbit/s ATM multimode fiber-optic interfaces

plastic fiber (IEC 793-2 Section 4 A4d), and for HPCF, 225 µm multimode polymer fiber is used (IEC 793-2 Section 3 A3d). The maximum loss is 9.1 dB for POF media and 1.8 dB for HPCF. The minimum modal bandwidth for both fiber types must be at least 10 MHz/km at a wavelength of 650 nm. The data stream is NRZ encoded; the PN and F07 connectors used must conform to IEC 1753-BB (FFS). Figure 10.16 lists the transmission and reception parameters for POF and HPCF interfaces.

ATM over 75 Ω Coaxial Cable

SDH-based ATM transmission at 155 Mbit/s over 75 Ohm coaxial cable is used almost exclusively in European wide-area networks; this interface is rarely used in North America, where it is called "EC-3" (EC stands for "Electrical Carrier").

Transmission parameters for ATM plastic fiber interfaces	POF	HPCF	Units
Maximum spectral width (FWHM)	40	40	nm
Numerical aperture (transmitter)	0.2 to 0.3	0.2 to 0.3	
Mean optical power	-8 to -2	-20 to -14	dBm
Wavelength	640 to 660	640 to 660	nm
Minimum extinction rate	10	10	dB
Maximum rise and fall time (10% – 90%)	4.5	4.5	ns
Maximum overshoot	25	25	%
Maximum systematic jitter	1.6	1.6	ns
Maximum random interface jitter	0.6	0.6	ns
Reception parameters for ATM plastic fiber interfaces	POF	HPCF	Units
Minimum sensitivity	-25	-26.5	dBm
Minimum overload	-2	-14	dBm
Maximum rise and fall time (10% – 90%)	5.0	6.0	ns
Maximum systematic jitter	2.0	2.0	ns
Minimum eye diagram aperture (time interval reserved for clock regeneration after electrical/optical conversion)	1.23	1.23	ns
Maximum random interface jitter	0.6	0.6	ns

Figure 10.16 Optical transmission and reception parameters for POF and HPCF interfaces ATM over 75 Ohm Coaxial Cable

The physical medium and the encoding technique are the same as those specified in ITU-T Recommendation G.703 for the 140 Mbit/s E4 interface. Two 75 Ω coaxial cables are used: one for transmission and one for reception. The signals are coded using Coded Mark Inversion (CMI) at a voltage level of ± 0.5 volts.

ATM over 100 Ω Copper Cable (Cat. 3 Unshielded Twisted Pair)

Using a special coding technique called CAP-64, which obtains high data rates at low frequency bandwidths, ATM in SDH/SONET framing can be transported at a speed of 155 Mbit/s even over low-quality UTP-3 data cable, which is very common in North America. The maximum segment length is 100 m. In addition to the STM-1/OC-3 rate of 155 Mbit/s, the ATM Forum has also specified slower rates for this cable type, 51.84 Mbit/s, 25.92 Mbit/s and 12.96 Mbit/s. The coding of the corresponding bit stream is CAP-16 for 51.84 Mbit/s, CAP-4 for 25.92 Mbit/s, and CAP-2 for 12.96 Mbit/s. CAP stands for Carrierless Amplitude Modulation/Phase Modulation, and is an extremely efficient method for achieving high data rates in spite of low available frequency bandwidth. The encoding process divides the symbol stream to be transmitted into n data paths (where n is the symbol period). Of the resulting symbol streams, one is sent through an in-phase filter, and the others through phase-shift filters. The output signal of the in-phase filter is added to the inverted output signal of the phase-shift filter and then sent to the twisted pair through a low pass filter. In this way the information is coded in the form of phase shifts. Each phase now contains not just one data bit, but an entire bit sequence. In the case of CAP-16, a given signal level can represent any of 16 different values, depending on its phase. The information contained in a single signal amplitude thus corresponds to 4 bits. If each amplitude/phase state is to represent 6 bits, a total of 64 different phase states are necessary.

The ATM cells are framed in STM-1/STS-3c containers/SPEs in accordance with ITU-T Recommendation G.707/ANSI T1.646 (Section 7.4). At transfer rates of 51.84 Mbit/s, the ATM cells are transported in the payload field of the STM-0/STS-1 frame. The entire payload field of the STM-0/STS-1 frame is filled with cells, with the exception of columns 30 and 59, for a net bandwidth of 48.384 Mbit/s. The 25.92 Mbit/s and 12.96 Mbit/s data rates are obtained by reducing the STS-1 frame rate.

- 51.84 Mbit/s: frame period of 125 µs
- 25.92 Mbit/s: frame period of 250 µs
- 12.96 Mbit/s: frame period of 500 µs

As in the case of 155 Mbit/s network, the maximum length of network segments at 51.84 Mbit/s is 90 m, plus 10 m of flexible patch cable. If higher-quality Cat. 5 cable is used rather than Cat. 3, up to 160 m can be spanned at 51.84 Mbit/s. The lower rates, 25.92 Mbit/s and 12.96 Mbit/s, permit segment lengths of 320 and 400 m with Cat. 5 copper cabling. UTP-3 cabling is used with 8-pin RJ-45 connectors (IEC 603-7) that conform to the electrical specification ANSI/TIA/EIA-568-A.

ATM over 100 Ω Copper Cable (Cat. 5 Unshielded Twisted Pair)

When Category 5 UTP cabling is used to transmit data at up to 155 Mbit/s, the segment length must not exceed 150 m. As for Cat. 3 wiring, the standard connector is the 8-pin RJ-45 plug. With higher-quality receivers, segments can be up to 350 m long. The electrical signal is NRZ-encoded.

ATM over 150 Ω Shielded Twisted Pair (STP)

The use of 150 Ω shielded twisted pair also permits segment lengths of up to 350 m at 155 Mbit/s. Nine-pin D-sub connectors or IBM MIC connectors, as used in Token Ring networks, are recommended. Here too the bit stream is NRZ-encoded.

10.1.6.4 Cell-Based Physical Layer

When ATM cells are transmitted over data lines as a plain bit stream, using neither PDH nor SDH/SONET framing, this is called cell-based physical layer. Interface specifications exist for cell-based physical layer at 51.84 Mbit/s (with optional partial transfer rates of 25.92 Mbit/s and 12.96 Mbit/s), 155 Mbit/s and

622 Mbit/s. The cells are transmitted in a continuous stream of ordinary ATM data cells, OAM cells and idle/unassigned cells. Every 27th cell at most may be a physical layer cell, which is either an idle cell inserted when no ATM user data cells are queued for transmission, or a PL-OAM cell. The latter type is used to carry out the operation monitoring functions that are otherwise accomplished by SDH/SONET headers. At least one PL-OAM cell is required for every 513 cells.

ATM Cell Streams over V.35, EIA/TIA 449/530, HSSI, and E1

The ATM Forum has specified a cell-based transmission convergence sublayer based on ITU-T Recommendation I.432 for “clear channel” interfaces. This term designates all interfaces that are capable of transporting any data stream without imposing bit stream encoding and framing restrictions. Examples include V.35, EIA/TIA 449/530, EIA/TIA 612/613 (High Speed Serial Interface or HSSI) and unframed E1. Any other clear channel interface can also be used, however. The cells are transferred in a continuous stream (of ordinary ATM cells, OAM cells, and idle/unassigned cells). No F1 or F2 OAM functions are specified for monitoring the network. F3 OAM functions can be optionally implemented using special physical layer OAM cells to monitor processes at the level of transmission paths. In this case, the following parameters can be analyzed:

- The number of included cells (NIC) per OAM cell: 128
- The number of cells for which transmission error values are calculated (Monitoring Block Size or MBS): 16
- The number of blocks monitored per OAM cell: 8
- The number of monitored blocks received by the remote station: 8

The cell rate is decoupled from an interface’s data rate by inserting idle/unassigned cells. The remote station is synchronized with the individual ATM cells using the Header Error Control (HEC) mechanism described in the next section. Cells may be transmitted either in scrambled or in unscrambled form.

10.1.6.5 ATM Cell Streams over FDDI Infrastructures (TAXI)

TAXI is a special variant of cell-based ATM transmission is TAXI. The TAXI interface was defined to support the use of existing FDDI infrastructures to transport ATM cells. In this way FDDI rings can be converted into ATM networks while conserving most of the existing FDDI hardware. The name TAXI originated with the first commercially available chipset for FDDI-based ATM. The fiber optic media and signal characteristics specified for TAXI communication are exactly the same as those defined in the FDDI standard ISO 9314-3. The ATM cells are 4B/5B-encoded and transmitted without any additional framing. In 4B/5B encoding, 4 bits of data are transmitted as symbols of 5 bits. This is due to the requirement that no more than three ‘1’ bits occur in a row. The bit sequence 1111 is coded as 11101, for example. Of the 32 possible 5-bit symbols, only 16 are used to transmit data (a byte is represented by a symbol pair). Some of the remaining 16 symbols are used as line state or control symbols in FDDI and some others are unusable as they have long runs of 1s or 0s. The symbol pair sequence JK, for example, announces the beginning of an FDDI frame; the I or Idle symbol is used as a continuous padding-bit stream for clock synchronization. The beginning of an ATM cell is marked by the control symbol pair sequence TT. In TAXI, the JK sequence is used as an idle symbol and is sent when there are no assigned cells to send (unassigned and idle cells are not used in the TAXI interface – this is a truly "asynchronous" interface). In case of noise, interfaces resynchronize only when the next idle signal is received. If higher cell losses are tolerable, fewer idle symbols may be required, but not less than one every 0.5 seconds. The MIC connector described in ISO 9314-3, customary in FDDI networks, is also used in TAXI. The TAXI interface is now obsolete, as STM-1/OC-3 interfaces have become commonplace in the local area, offering a higher bandwidth.

10.1.6.6 ATM Cell Streams at 25.6 Mbit/s

Originally developed to allow the economical connection of workstation computers to ATM networks but more recently associated with ATM over ADSL modem interfaces, a further specification for a 25.6 Mbit/s ATM interface was developed independently of the cell-based physical layer interfaces described previously. The ATM Forum specification calls for Cat. 3 100 Ω UTP (unshielded twisted pair), 120 Ω Cat. 4 (ISO/IEC11801) or 150 Ω STP (shielded twisted pair). As for the TAXI interface, the bit stream is 4B/5B-encoded (the actual coding is different from TAXI) and transmitted asynchronously without further framing. The transfer rate of 25.6 Mbit/s with 4B/5B encoding implies a line rate of 32 Mbaud. This interface is closely based on an original specification developed by IBM and is derived from Token-Ring technology.

The specified maximum segment length for all three cable types is 100 m (90 m plus 10 m for patch cables). The cables must also conform to the values specified in EIA/TIA-568-A or ISO/IEC 11801 for attenuation and near-end crosstalk (NEXT). The connectors are RJ-45 for Cat. 3 UTP or STP-MIC for shielded twisted pair.

10.1.6.7 ATM and DXI Interfaces

The Data Exchange Interface (DXI) is a simple transmission protocol developed for interfaces such as V.35, EIA/TIA 449/530 or EIA/TIA 612/613 (HSSI). Originally, DXI was conceived as a data tributary protocol for metropolitan area networks (SMDS-DXI, SMDS Interest Group Technical Specification SIG-TS-005/1993). For lack of an ATM-specific DXI to connect their systems to ATM networks, many manufacturers began by implementing the simpler SMDS-DXI as an interface protocol. The ATM Forum subsequently drafted a DXI variant tailored to the requirements of ATM (af-dxi-0014).

The functions of ATM DXI are divided into user-end processes (Data Terminal Equipment or DTE, Modes 1a and 1b) and network-end processes (Data Communications Equipment or DCE, Mode 2). In Mode 1a, user interfaces support up to 1,023 virtual connections with AAL5 Protocol Data Units or PDUs and data packets (Service Data Units or SDUs) of up to 9,232 bytes. The DTE SDUs, which correspond to AAL5 PDUs, are encapsulated in DXI frames for transmission. Mode 1b also supports AAL3/4 with a maximum SDU length of 9,224 bytes. The

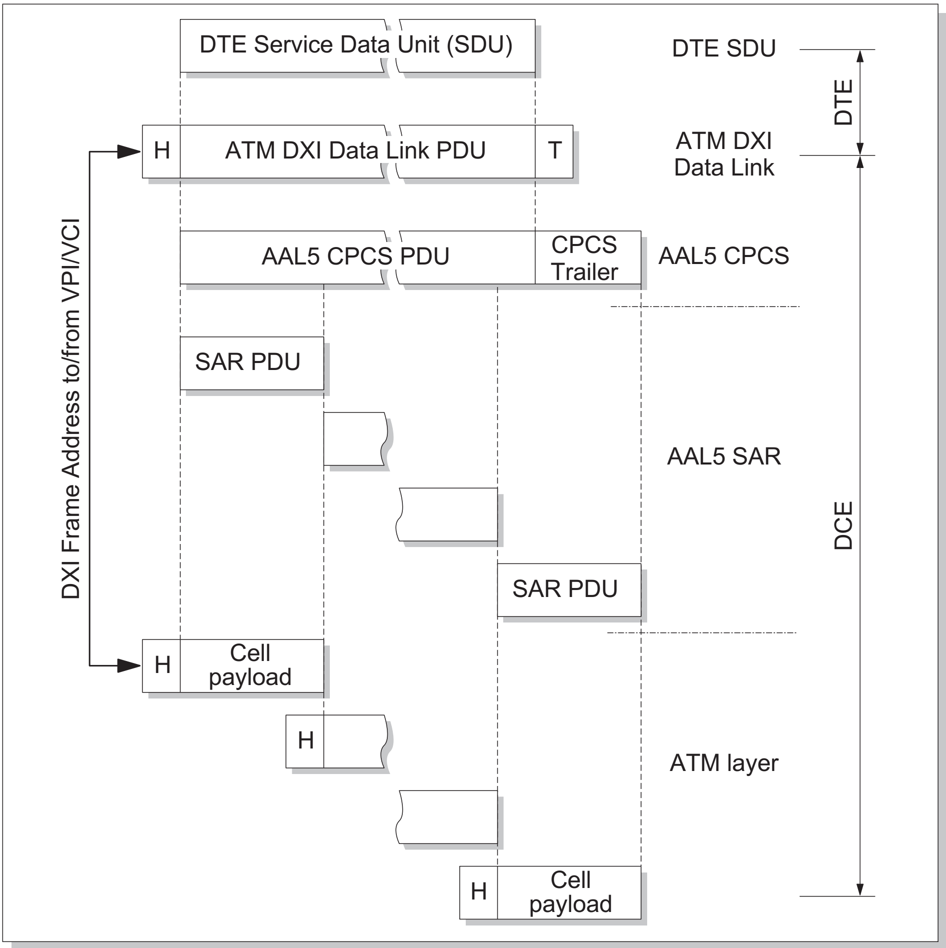

Figure 10.17 Encapsulation of ATM cells in DXI PDUs (Modes 1a and 1b)

AAL3/4 PDUs are encapsulated by the DTE, but all segmentation and reassembly (AAL3/4 SAR) must be performed by the DCE (Figure 10.17).

In Mode 2, network interfaces support up to 16,777,215 virtual connections for AAL5 and AAL3/4. SDUs can be up to 65,535 bytes long.

The DXI frame address (DFA) is carried in the DXI header. It is used to convey the VPI and VCI information between DTEs (Data Terminal Equipment) and DCEs (Data Communication Equipment). The DFA is 10 bits long in Modes 1a and 1b, and 24 bits long in Mode 2 (Figure 10.18). Note that no ATM cells are involved at this interface, only AAL frames.

Flag	DXI header	DTE SDU	DXI FCS	Flag
1	2	$0 < N \leq 9232$ (Bytes)	2	1

Figure 10.18 DXI data packet for Modes 1a and 1b (DTE, AAL5)

10.1.6.8 ATM Physical Layer Operations and Maintenance (OAM F1- F3)

A total of five information flows are defined for operations and maintenance of ATM networks, called the OAM flows. Each of these flows, numbered F1 to F5, is responsible for the monitoring of a certain part of a connection. Flows F1 to F3 are concerned with the physical layer of the B-ISDN protocol model; F4 and F5 with virtual paths (VP) and virtual channels (VC) in the ATM layer. The OAM parameters are determined by different mechanisms depending on the network topology used for ATM cell transport (SDH/SONET, PDH, cell-based physical layer). A transmission path is defined as the path between two network components that insert the user data into the transport medium and extract it again from the medium. Examples of transmission paths include the link between a B-NT2 and a switch (VP cross-connect), or between a B-NT2 and the connection end point. A transmission path is composed of several digital sections, each of which may lead through one or more regenerator sections.

OAM F1 – F3 for SDH/SONET-Based ATM Systems

In SDH/SONET-based ATM networks, the OAM flows F1 and F2 are transported in the SOH/TOH of the transport frames, and F3 in the POH of the given virtual container /SPE. Part of the F3 information can also be transported in specially indicated physical layer OAM cells (PL-OAM).

OAM F1 – F3 for PDH-Based ATM Systems

In PDH networks, F1 and F3 information can be transported in the PDH header. The frame alignment byte of the header is analyzed for F1 functions, while F3 parameters are taken from the remaining header bytes. There is no provision for an F2 flow.

OAM F1 – F3 for Cell-Based Physical Layer

In cell-based physical layer networks, the OAM information flows are transmitted by means of special physical layer OAM cells (PL-OAM cells). As described in I.432.2, evaluation of OAM-F1 and OAM-F3 is provided for, but there is no provision for an F2 flow. The corresponding parameters are communicated as part of F3. PL-OAM cells may be inserted in the cell stream no more than once every 27 data cells, and no less than once in 513 cells. The F1 cell contains OAM parameters for the regenerator section; F3 cells are used to monitor the trans-

OAM flow	Byte 1	Byte 2	Byte 3	Byte 4	Byte 5
F1	0000000	0000000	0000000	0000011	Valid HEC = 0101110
F3	0000000	0000000	0000000	0001001	Valid HEC = 0110101

Because the PL OAM cells are not passed to the ATM layer, the field values are of no significance as ATM cells.

Figure 10.19 PL-OAM cell headers

Byte	Field	Byte	Field	Byte	Field
1	R	17	R	33	EB4 EB3
2	AIS	18	R	34	EB6 EB5
3	PSN	19	R	35	EB8 EB7
4	NIC	20	R	36	R
5		21	R	37	R
6	MBS	22	R	38	R
7	NMB-EDC	23	R	39	R
8	EDC-B1	24	R	40	R
9	EDC-B2	25	R	41	R
10	EDC-B3	26	R	42	R
11	EDC-B4	27	R	43	R
12	EDC-B5	28	R	44	R
13	EDC-B6	29	R	45	R
14	EDC-B7	30	RDI (1)	46	STET
15	EDC-B8	31	NMB-EB	47	CEC (10)
16	R	32	EB2 EB1	48	

Figure 10.20 PL-OAM cell format for F1 and F3 cells

mission path. OAM cells have a special header that makes them easily recognizable (Figure 10.19). Note that the lower limit of 27 exists to allow easier interworking between cell-based interfaces running at 155 Mbit/s and 622 Mbit/s and the corresponding SDH/SONET interfaces (STM-1/OC-3 and STM-4c/OC-12) because the overhead in SDH/SONET (SOH/TOH + POH) occupies exactly 1/27 of the bandwidth (that is, 9+1 columns out of 270).

Figure 10.20 illustrates the structure of PL-OAM F1 and F3 cells. The following fields are reserved for the F1 and F3 flows:

OAM-F1 Cell Fields

PSN	PL-OAM Sequential Number (8 bits, modulo 256)
NIC	Number of Included Cells (maximum value: 512)
MBS	Monitoring Block Size (maximum value: 64)
NMB-EDC	Number of Monitored Blocks (recommended value: 8)
EDC	Error Detection Code (BIP 8 value calculated from the cells of the MBS)
NMB-EB	Number of Monitored Blocks at the Far End (recommended value: 8)
REI	Remote Error Indication (number of bit parity errors in each block)
AIS	Alarm Indication Signal (an AIS sent downstream to report an upstream error) (AIS-L for SONET)
RDI	Remote Defect Indicator (sent upstream to report a reception failure downstream. This occurs when frame synchronization or the data signal is lost, for example.)
CRC	CRC-10 Checksum
R	Reserved (set to 01101010)

OAM-F3 Cell Fields

PSN	PL-OAM Sequential Number (8 bits, modulo 256)
NIC	Number of Included Cells (maximum value: 512)
MBS	Monitoring Block Size (maximum value: 64)
NMB-EDC	Number of Monitored Blocks (EDC octets) (recommended value: 8)
EDC	Error Detection Code (BIP 8 - Bit Interleaved Parity - value calculated from the cells of the MBS)

NMB-EB	Number of Monitored Blocks at the Far End (recommended value: 8)
REI	Remote Error Indication (number of bit parity errors in each block) (REI-P for SONET)
AIS	Alarm Indication Signal (AIS-P for SONET)
TP-RDI	Transmission Remote Defect Indicator (formerly TP-FERF) (RDI-P for SONET)
CRC	CRC-10 Checksum
R	Reserved (set to 01101010)

10.1.7 The ATM Layer

The actual transport of ATM cells takes place at the ATM layer. In order to provide for the varying connection quality requirements of different applications, Quality-of-Service (QoS) parameters are first negotiated in the signaling process or provisioned by the network operator. After a successful connection setup transmission begins with the ATM cells of virtual channel connections (VCCs) and virtual path connections (VPCs) multiplexed in a continuous cell stream. During the connection, monitoring and control mechanisms operate to ensure that the connection parameters agreed upon in the ATM switching negotiation are maintained.

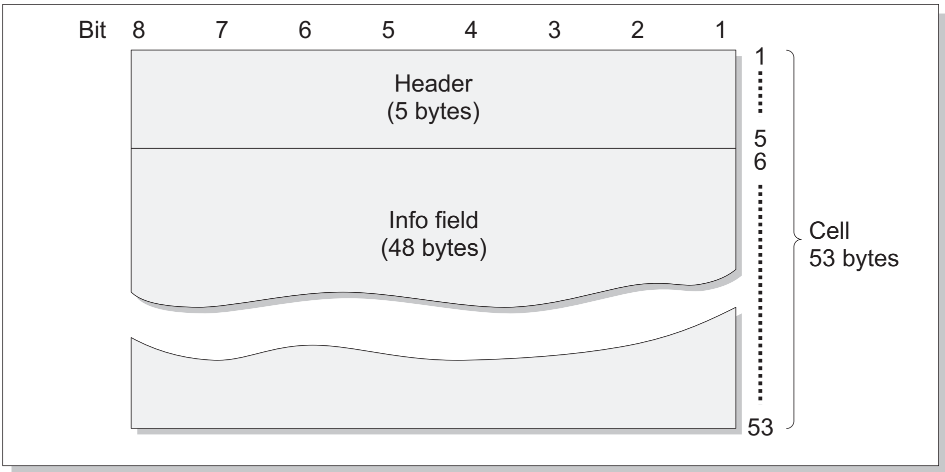

Figure 10.21 ATM cell

An ATM cell consists of a 5-byte header and a 48-byte information or user data field. There are two basic types of cells, UNI cells and NNI cells. UNI cells are transferred at user–network interfaces, NNI cells at network–node interfaces. The two cell types differ only in 4 header bits, which are specified for flow control in UNI cells (though, in reality not used for this) and to extend the virtual path identification (VPI) field to 12 bits in NNI cells.

Bit 8 is the most significant bit in all fields. The bits within each byte are therefore transmitted beginning with bit 8. The bytes in turn are transmitted in ascending order, that is, beginning with byte 1. Figure 10.22 illustrates the header structure of UNI cells. The header consists of the six fields GFC (4 bits), VPI (8 bits), VCI (16 bits), PT (3 bits), CLP (1 bit) and HEC (8 bits).

<table>
<tr><th>Byte \ Bit</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th></tr>
<tr><td>1</td><td colspan="4">Generic Flow Control (GFC)</td><td colspan="4">Virtual Path ID (VPI)</td></tr>
<tr><td>2</td><td colspan="4">Virtual Path ID (VPI)</td><td colspan="4">Virtual Channel ID (VCI)</td></tr>
<tr><td>3</td><td colspan="8">Virtual Channel ID (VCI)</td></tr>
<tr><td>4</td><td colspan="4">Virtual Channel ID (VCI)</td><td colspan="3">Payload Type (PT)</td><td>CLP</td></tr>
<tr><td>5</td><td colspan="8">Header Error Control (HEC)</td></tr>
</table>

CLP: Cell Loss Priority

Figure 10.22 UNI cell header

The Generic Flow Control Field (GFC)

The generic flow control field, which consists of 4 bits, is used to control local functions and to manage access and transmission rights in ATM networks. Its contents are not forwarded beyond the range of the local ATM switch because this field is replaced with VPI data at the NNI. Its significance is thus limited to the local ATM network segment. In practice, although specified by the ITU-T in Recommendation I.361, generic flow control has seldom, if ever, been implemented, and the ATM Forum in particular does not support this function.

The ATM Routing Label Field (VPI-VCI)

The UNI header contains a total of 24 bits for ATM layer routing purposes: 8 bits for the Virtual Path Identifier (VPI) and 16 bits for the Virtual Channel Identifier

(VCI). The VPI-VCI value is a label that has only local significance, that is, between two ATM interfaces in a single transmission link. This label has no end-to-end significance, which is why the VPI-VCI field is referred to as a label, not an address. There is such a thing as an ATM Address that is used in connection with signaling (see the section on *Connection Setup at the Caller's End*), and some publications do refer, incorrectly, to the VPI-VCI as an address. The value of the VPI-VCI field changes as cells of a given virtual connection pass through ATM switches. The "translation" of the VPI-VCI label value is determined by the contents of a "translation table" whose contents are determined by the signaling process, in the case of switched virtual connections (SVCs); network management provisioning, in the case of (semi-) permanent virtual connections (PVCs); or a label distribution protocol, in the case of multiprotocol label switching (MPLS).

An ATM virtual channel (VC) refers to a two-way transmission link for ATM cells, although both directions may not always be used. All cells in a given virtual channel have the same VCI value in a given link. Several ATM virtual channels may be carried within an ATM virtual path (VP). As with the virtual channel, all cells transported by a given ATM virtual path are identified by a particular VPI. Rather than the physical channels and paths familiar from older telecommunications techniques, the ATM concept of virtual paths permits substantially more efficient use of available bandwidth.

Figure 10.23 Virtual channels (VCI) and virtual paths (VPI)

The Payload Type Field (PT)

Three bits of the ATM header are used to identify the type of data the cell payload contains. This makes it possible to distinguish between ordinary user data and various special types, such as traffic or network management information, for example. The default value of the PT field is 000.

Cell Loss Priority (CLP)

The cell loss priority bit can be used to assign cells a relative priority. If the CLP bit is set to 1, the cell has low priority; 0 indicates normal (higher) priority. If the capacity of a connection is exceeded, or if other transmission problems arise, cells with low priority are discarded first.

Header Error Control (HEC)

Before an ATM cell is transmitted, a cyclic redundancy check (CRC) of the entire cell header is calculated and placed in the 8-bit header error control field. At the receiving end of each link, the header is checked to see if the HEC value is correct; it would normally only be correct if there were no bit errors anywhere in the header. If a single bit error has occurred somewhere in the header, including the HEC itself (because of physical layer bit errors, for example), there is sufficient redundancy in the HEC to allow that bit error to be corrected. Similarly, two bit errors will always be detectable. However, the detection of three or more bit errors in the header may not be detectable. For example, multiple bit errors may change the header such that it becomes a header with a different (corrupt) VPI-VCI value and for which the (corrupt) HEC is "correct". If the new VPI-VCI value happened to be a "legal" value, that is, one belonging to an existing virtual connection, the cell with the corrupt header would be misinserted into this other virtual connection (and, of course, lost from the

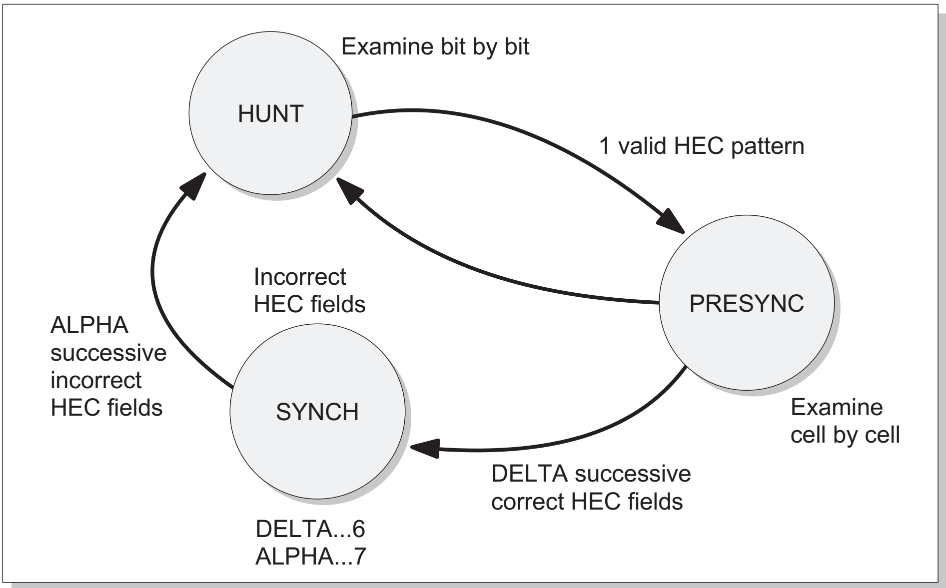

Figure 10.24 State diagram of the HEC cell synchronization algorithm

connection to which it really belongs). However, if the new VPI-VCI value was not supported, the switch would not forward it but would discard it because no translation table entry for it would exist. Note that some interfaces restrict the use of the HEC to detecting errors but not correcting them; this is because certain types of physical layer scrambling can give rise to error multiplication, making reliable correction impossible.

In addition to bit-error detection (and correction), the HEC field also allows receiving stations to synchronize with the beginning of the cell. This process is called "cell delineation".

The following algorithm is applied in order to detect the beginning of an ATM cell in a bit stream:

- In the HUNT state, the incoming signal is analyzed bit-by-bit to determine whether it could be part of an HEC pattern. As soon as a potential HEC pattern has been found, the receiver changes to the PRESYNC state.
- Under the assumption that a cell has been detected, the PRESYNC state examines subsequent "candidate cells". If the HEC fields of the next DELTA candidate cells are also possible checksums, then the receiver assumes that it has synchronized with the cell stream, and switches to the SYNC state and cell synchronization or "delineation" is declared. If less than DELTA consecutive cells meet the HEC test, the receiving station reverts to the HUNT state.
- If ALPHA consecutive cells fail the HEC test, the receiver in SYNC state assumes it has lost its cell delineation and switches to HUNT state. Loss of Cell Delineation (LCD) is then declared.
- For SDH/SONET physical layers the value of ALPHA and DELTA are 7 and 6 respectively; for cell-based physical layers these values are 7 and 8 respectively (see ITU-T Recommendation I.432.1).

Data Field Scrambling

The data field (payload) of the ATM cells is scrambled in order to optimize cell delineation by HEC patterns. In the HUNT state, potential HEC sequences are easier to detect when the data is scrambled. In the PRESYNC and SYNC states, an unscrambling function is activated for the 48 user data bytes of the cell, but is not applied to the cell headers. In all framing types except cell-based physical layer, the scrambling of the data field is always based on the self-synchronizing scrambler (SSS) $x^{43}+1$; scrambling is optional in some interface types. Due to the relatively poor overall communication characteristics of the SSS process, cell-based physical layer uses Distributed Sample Scrambling (DSS) with the generator polynomial $x^{31}+x^{28}+1$. Note that, because the scrambling takes place within the transmission convergence sublayer of the physical layer for non cell-based interfaces, all cells at the interface are scrambled using the SSS process, independent of the header routing label value (VPI-VCI) and including any idle/unassigned and OAM cells, etc.

Reserved Header Values

Certain header values are reserved for cells with special operation, management or signaling functions in the ATM network. Such cells include broadcasts, meta-signaling, resource management cells and, in cell-based physical layer, PL-OAM cells. All special cells have a VCI value in the range 0 through 31; the ITU-T reserves to itself all values of VCI below 16 and makes available all values between 16 and 31 for proprietary purposes; the ATM Forum has used several values in this range for its purposes (for example, VCI = 16 is used for ILMI, VCI = 17 is used for LANE, VCI = 18 is used for PNNI). With the exception of these defined values, all values (VCI = 32 and above) can be used for user traffic by the ATM layer. Figure 10.25 lists the reserved header byte values (ITU-T Recommendation I.361). Furthermore, the ITU-T recommendations and ATM Forum specifications define certain UNI cell header values as reserved. These are listed in Figure 10.26.

Reserved Header Bytes in UNI ATM cells (ITU I.361, I.432)

	Byte 1	Byte 2	Byte 3	Byte 4
Reserved for physical layer [1, 2]	PPPP 0000	0000 0000	0000 0000	0000 PPP1
Physical layer F1 OAM cell	0000 0000	0000 0000	0000 0000	0000 0011
Physical layer F3 OAM cell	0000 0000	0000 0000	0000 0000	0000 1001
IMA ICP cell	0000 0000	0000 0000	0000 0000	0000 1011
Idle cells	0000 0000	0000 0000	0000 0000	0000 0001
Unassigned cells	AAAA 0000	0000 0000	0000 0000	0000 AAA0

A ...... Bit available for use by the ATM layer
P ...... Bit available for use by the physical layer
1 .... In cells with the VPI-VCI value equal to 0-0, the four bits normally used to represent the PT and CLP fields are reinterpreted to distinguish between different types of unassigned and physical layer cells.
2 .... Cells identified by header information as physical layer cells are not passed to the ATM layer.

Reserved VPI and VCI Values in UNI ATM cells (ITU I.361)

	VPI	VCI
Meta-signaling channel	0000 0000	0000 0000 0000 0001
Broadcast signaling channel	0000 0000	0000 0000 0000 0010

Figure 10.25 Reserved header bytes and VPI/VCI values

Use	Value[1,2,3,4]			
	Octet 1	Octet 2	Octet 3	Octet 4
Unassigned cell indication	00000000	00000000	00000000	0000xxx0
Meta-signalling (default) [5, 7]	00000000	00000000	00000000	00010a0c
Meta-signalling [6, 7]	0000yyyy	yyyy0000	00000000	00010a0c
General Broadcast signalling (default) [5]	00000000	00000000	00000000	00100aac
General Broadcast signalling [6]	0000yyyy	yyyy0000	00000000	00100aac
Point-to-point signalling (default) [5]	00000000	00000000	00000000	01010aac
Point-to-point signalling [6]	0000yyyy	yyyy0000	00000000	01010aac
Invalid Pattern	xxxx0000	00000000	00000000	0000xxx1
Segment OAM F4 flow cell [7]	0000aaaa	aaaa0000	00000000	00110a0a
End-to-End OAM F4 flow cell [7]	0000aaaa	aaaa0000	00000000	01000a0a

[1] ..... "a" indicates that the bit is available for use by the appropriate ATM layer function.
[2] ..... "x" indicates "don't care" bits.
[3] ..... "y" indicates any VPI value other than 00000000.
[4] ..... "c" indicates that the originating signalling entity shall set the CLP bit to 0. The network may change the value of the CLP bit.
[5] ..... Reserved for user signalling with the local exchange.
[6] ..... Reserved for signalling with other signalling entities (e.g. other users or remote networks).
[7] ..... The transmitting ATM entity shall set bit 2 of octet 4 to zero. The receiving ATM entity shall ignore bit 2 of octet 4.

Figure 10.26 Header bytes reserved by the ATM Forum

The NNI Header

Unlike the UNI header, the NNI header provides 28 bits for addressing: 12 bits for the VPI and 16 bits for the VCI. This permits the definition of a greater number of virtual paths at the network-node interface. The payload type field, the priority bit and the HEC field correspond with the same fields in the UNI cell header. Figure 10.27 lists the reserved byte values in the NNI header:

	Byte 1	Byte 2	Byte 3	Byte 4
Reserved for physical layer [1, 2]	0000 0000	0000 0000	0000 0000	0000 PPP1
Physical layer OAM cell	0000 0000	0000 0000	0000 0000	0000 1001
Idle cells	0000 0000	0000 0000	0000 0000	0000 0001
IMA ICP	0000 0000	0000 0000	0000 0000	0000 1011
Unassigned cells	AAAA 0000	0000 0000	0000 0000	0000 AAA0

A ...... Bit available for use by the ATM layer
P ...... Bit available for use by the physical layer
[1] .... In cells with the VPI-VCI value equal to 0-0, the four bits normally used to represent the PT and CLP fields are reinterpreted to distinguish between different types of unassigned and physical layer cells.
[2] .... Cells identified by header information as physical layer cells are not passed to the ATM layer.

Figure 10.27 Reserved header bytes in the NNI ATM cell

10.1.7.1 ATM Cell Types

In addition to user ATM cells (cells whose VCI > 31), there are several types of cells that are used not to transport user data, but to perform certain operational functions. These include idle cells, unassigned cells, PL-OAM cells, ICP cells, RM cells and VP/VC-OAM cells. Note that in cells having a VPI-VCI = 0-0, the least 4 bits in the 4th byte normally used to convey the PT and CLP fields are re-assigned for other purposes (see Figure 10.27).

Idle Cells

Idle cells are physical layer cells and carry no useful information. They are used to adapt the cell rate to the bandwidth of the transmission medium. If there are not enough ordinary cells to fill the allocated bandwidth, idle cells are inserted by the transmission convergence Sublayer. This permits alignment of the ATM cell stream with the throughput of the physical medium (such as an SDH-VC4 container/SONET SPE). Unlike unassigned cells, idle cells are not passed to the ATM layer. The ATM Forum, following the Bellcore definition, calls the idle cell header pattern an "invalid header", revealing some confusion in the minds of Bellcore (and hence ATM Forum) specifiers (see *Unassigned Cells* for more on this). See Figure 10.27.

Unassigned Cells

Unassigned cells (VPI-VCI = 0-0) are ATM layer cells and carry no useful payload information. They are used when no assigned cells are available to send from the ATM layer. At the receiver, they are treated similarly to idle cells and discarded. In North America (and implementations elsewhere based on certain ATM Forum specifications), unassigned cells are used where the ITU-T would specify idle cells; in particular, they are specified for rate adaption by the ATM Forum, adopting Bellcore usage which, strictly, conflicts with ITU-T usage. Despite this confusion, idle and unassigned cells can normally be considered equivalent. See Figure 10.27.

Physical Layer OAM Cells

In the cell-based physical layer, special cells can be inserted up to once in every 27 cells to transmit operation and maintenance information concerning the physical layer. These are called PL-OAM cells. When received, these cells are used by the physical layer and not passed along to the ATM layer. Their purpose is to convey some of the information (such as alarms) normally carried by the overhead of frame-based physical layers, such as SONET.

IMA Control Protocol (ICP) Cells

These are special control cells used with implementations of the ATM Forum's Inverse Multiplexing for ATM specification.

VP-OAM and VC-OAM Cells

VP-OAM and VC-OAM cells are used to transport the F4 and F5 maintenance flows respectively. This information allows the network to monitor and test the capacities and availability of ATM virtual paths and virtual channels. VP-OAM cells use reserved values of VCI = 3 and 4 for each VPI value for segment and end-to-end F4 OAM flows respectively; VC-OAM cells use a reserved value of PT field = 100 (binary 4) and 101 (binary 5) for each value of VPI-VCI for segment and end-to-end F5 OAM flows.

VP/VC-RM Cells

These are resource management cells used, so far, for implementations of the available bit rate (ABR) traffic type. VP-RM cells have a reserved value of VCI = 6, VC-RM cells have a reserved PT field value of 110 (binary 6).

VP/VC Cells

VP/VC cells are common cells used for communication within ATM virtual channels (VCs) and virtual paths (VPs). VP/VC cells can be of the following five types:

- User data transport cells (any VPI, VCI > 31)
- Metasignaling cells (VPI = 0, VCI = 1)
- Broadcast signaling cells (VCI = 2)
- Point-to-point signaling cells (VCI = 5)
- ILMI cells for ATM network management (VPI-VCI = 0-16)

User data are assigned to a specific connection by their VPI/VCI values and transport data for higher-layer services in the 48-byte payload field. Meta-signaling cells were originally conceived to select and define signaling virtual channels. Because meta-signaling is not used in present-day networks, however, these cells do not occur in actual practice. Broadcast signaling cells are used to send signaling information to all network stations. Network nodes that do not support broadcast signaling simply ignore all cells in the broadcast virtual channel, VCI 2. VCI 5 cells are used for point-to-point signaling. For this reason the virtual channel VCI 5 is also called the signaling virtual channel. Finally, ILMI cells are used for ATM network local management tasks at the UNI, which include registration of new active stations with the ATM switch, querying ATM MIBs or configuring network components.

10.1.7.2 ATM Connections

In order to manage the wide variety of communication situations in ATM networks, the ATM layer provides for different connection types with different characteristics. Both virtual channel and virtual path connections can be structured as point-to-point or point-to-multipoint. The bandwidth allocated to a connection can also be asymmetric: that is, the bandwidth for transmission can be lower than for reception, or vice versa. Furthermore, a number of Quality-of-Service (QoS) parameters can be negotiated for each connection.

Virtual Channel Connections (VCCs)

ATM virtual channels represent the lowest level in the structural hierarchy of ATM data streams. All virtual channel connections have the following four properties:

- The ATM switch assigns each virtual channel connection QoS parameters, that define properties such as cell loss ratio or cell delay.
- Virtual channel connections can be either dynamically switched (SVC) or (semi-) permanent (PVC).
- The sequential order of cells in a virtual channel is preserved during transportation through the ATM network (the virtual channel is “connection-oriented”).
- For each virtual channel connection, traffic parameters, such as the maximum bandwidth available for the connection; the result is the “traffic contract”. Cells sent to the network by the user are monitored to ensure conformance with the traffic contract.

A VCC can be set up by using four different signaling methods:

- No signaling (for example, a (semi-) permanent connection is set up through network management)
- Meta-signaling
- User-to-network signaling
- Network-to-network signaling

Setting up permanent or semi-permanent VCCs is a good idea especially when a relatively low number of network nodes need to be interconnected over statically defined connections, or when dynamic signaling processes are undesired for security reasons. In ATM networks with a large number of nodes, connections are set up dynamically over individual transmission sections using signaling protocols:

UNI<>NNI Signaling from an end system into an ATM network that consists of several networked switches (user-to-network)

NNI<>NNI Signaling within an ATM network that consists of several networked switches (network-to-network)

Meta-signaling is not used in practice because the use of the reserved signaling virtual channel, VCI 5, has been found to be sufficient even in large ATM networks.

Virtual Path Connections (VPCs)

Virtual path connections are a hierarchical level above virtual channel connections. In other words, a virtual path can contain several virtual channels (note, however, that in terms of the layered model, the VC is above the VP in the stack). VPCs have the same properties as VCCs. Path connections can likewise be set up manually as permanent virtual paths, or on request by means of signaling processes or network management functions. One aspect of VPC worth noting is that, while the VPI of a VPC changes at every switching node in the network, the VCI values of all VCs within the VP are preserved end-to-end.

10.1.7.3 Quality-of-Service (QoS) Parameters

In ATM networks, services with widely differing communication requirements are transported concurrently. Real-time applications with varying bit rates are much more exacting with regard to average cell delay, for example, than simple data transfers at a constant bit rate. For this reason, the ATM layer assigns each connection QoS parameters that specify certain characteristics of the connection when the connection is set up. ITU-T Recommendation I.356 defines seven cell transfer performance parameters:

- Cell Error Ratio
- Severely Errored Cell Block Ratio
- Cell Loss Ratio
- Cell Misinsertion Rate
  (the proportion of cells with a valid but incorrect header)
- Cell Transfer Delay
- Mean Cell Transfer Delay
- Cell Delay Variation

On connection setup, any user can request a “traffic contract” that specifies a particular QoS class for each direction of communication. Once the traffic contract has been negotiated, the ATM network, or rather the ATM switches along the transfer path, guarantee the QoS parameters granted for as long as the user respects the traffic contract. A distinction is made between service classes with defined QoS parameters and service classes without performance parameters. QoS classes with performance parameters must specify at least two such parameters. If a QoS class contains two parameters for cell loss ratio, then one value applies to cells with a cell loss priority (CLP) of 1, and the other to cells with CLP=0. Performance parameters that can be specified in a QoS class include:

- Maximum Cell Transfer Delay
- Cell Delay Variation
- Cell Loss Ratio for cells with CLP = 0
- Cell Loss Ratio for cells with CLP = 1

Quality-of-service classes and traffic parameters are defined for each of the six service classes defined in the ATM Forum’s Traffic Management specification:

- CBR Constant Bit Rate
- rt-VBR Real-Time Variable Bit Rate
- nrt-VBR Non-Real-Time Variable Bit Rate
- UBR Unspecified Bit Rate
- ABR Available Bit Rate
- GFR Guaranteed Frame Rate

A service class can be requested without QoS parameters, such as in a request for a connection with the “best possible network service”, for example. Such a request, in which no QoS parameters are specified, may be made when no explicit network performance guarantee is required. The load level and error frequency in ATM networks have a direct influence on the QoS parameters.

ATM Payload Types (PT)

Ordinary user data cells can be distinguished from special-purpose, non-user cells by means of the payload type field. The PT values 0, 1, 2 and 3 identify user data cells; 4, 5 and 6 indicate virtual channel segment OAM, end-to-end OAM, and resource management (RM) cells respectively. Figure 10.28 lists the possible values of the payload type field.

Payload type field	Meaning	
000	User data cell, no congestion detected	SDU type = 0
001	User data cell, no congestion detected	SDU type = 1
010	User data cell, congestion detected	SDU type = 0
011	User data cell, congestion detected	SDU type = 1
100	Segment OAM F5 cell	
101	End-to-end OAM F5 cell	
110	VC resource management	
111	Reserved for future use	

Figure 10.28 Payload type field

Cell Loss Priority and Selective Discarding of Cells

To protect the network from users who produce unauthorized traffic loads or otherwise violate their traffic contract, each station's data stream is monitored by the Usage Parameter Control (UPC). This entity analyzes and regulates the cell stream on each virtual path and virtual channel connection. The UPC can act in three ways at the cell level to regulate the data stream: each cell can be passed along, tagged or discarded. Cell passing is the normal transfer of all cells that conform to the traffic contract. Cell tagging is performed on traffic that does not conform to the sustainable cell rate for one particular traffic contract type (VBR.3/SBR3). When such cells are tagged, their CLP value is changed by the UPC from 0 (normal priority) to 1 (low priority). In case of network congestion, these cells will then be among the first to be discarded. If cell tagging is not supported, then cells that do not conform to the traffic contract are discarded immediately.

Traffic Shaping

The traffic contract negotiated during the connection setup includes a Connection Traffic Descriptor that defines the parameters for permissible traffic, including peak cell rate, duration of the peak cell rate (maximum burst size, MBS), etc. A transmitting station can then force its communication to conform to the negotiated traffic contract by means of the optional traffic shaping function. Another user strategy would be to send all queued cells as they occur, and simply tolerate tagging of high-priority cells (if supported) and discarding of low-priority cells.

10.1.7.4 Monitoring in ATM Networks (OAM Flows)

Of the five ATM network operations and maintenance information flows F1 through F5, F4 and F5 are situated in the ATM layer. All F4 and F5 OAM information is gathered and transmitted by means of special OAM cells. The F4 information flow is used for segment or end-to-end management at the virtual path level. A segment refers to one section of a connection, such as the link between two ATM switches (more than one switch may exist within a section); end-to-end refers to the entire communication path between the two endpoints for the complete virtual path. The information flow F5 is used for segment or end-to-end management at the virtual channel level; in this case, end-to-end refers to the entire communication channel between the two endpoints for the complete virtual channel, which may be longer than the virtual path. The purpose of F4 and F5 OAM cells is to provide the measurement data necessary to monitor the availability and performance of a given virtual channel or virtual path. Figure 10.29 lists the individual OAM functions of the ATM layer.

OAM function	Use
AIS (Alarm Indication Signal)	Error reports in the direction of transmission
RDI (Remote Defect Indication)	Error reports from the receiving system
Continuity check	Continuous monitoring of cell transport
Loopback	• Monitoring connections as needed • Localizing errors • Testing connections before making a link available
Forward performance monitoring	Estimating performance in the direction of transmission
Backward reporting	Returning results of forward performance monitoring
Activation/Deactivation	Activating or deactivating performance monitoring and continuity checking
System management	Managing by the given end systems
APS	ATM protection switching

Figure 10.29 ATM layer OAM functions

OAM Cell Format

The F4 OAM cells that transport management and monitoring information for virtual path connections have the same VPI value as the user cells on the virtual path being monitored. They are identified as F4 cells by the reserved VCI value 3 for F4 segment cells or 4 for F4 end-to-end cells. F5 OAM cells have the same VPI and VCI values as the user cells of the virtual channel connection to which they are related. They are identified by the payload type field (PT). The PT value 100 (decimal 4) denotes F5 segment OAM cells; 101 (decimal 5) indicates an F5 end-to-end OAM cell. End-to-end OAM cells must pass unchanged through all network nodes between the two connection endpoints. Only the endpoints may remove them from the cell stream. Segment OAM cells must be removed from the stream at the end of the given segment. There are five types of ATM-layer OAM cell: fault management, APS coordination protocol, performance management, activation/deactivation cells and system management. The OAM cell type is indicated by the first (4 bit) field in the OAM cell payload.

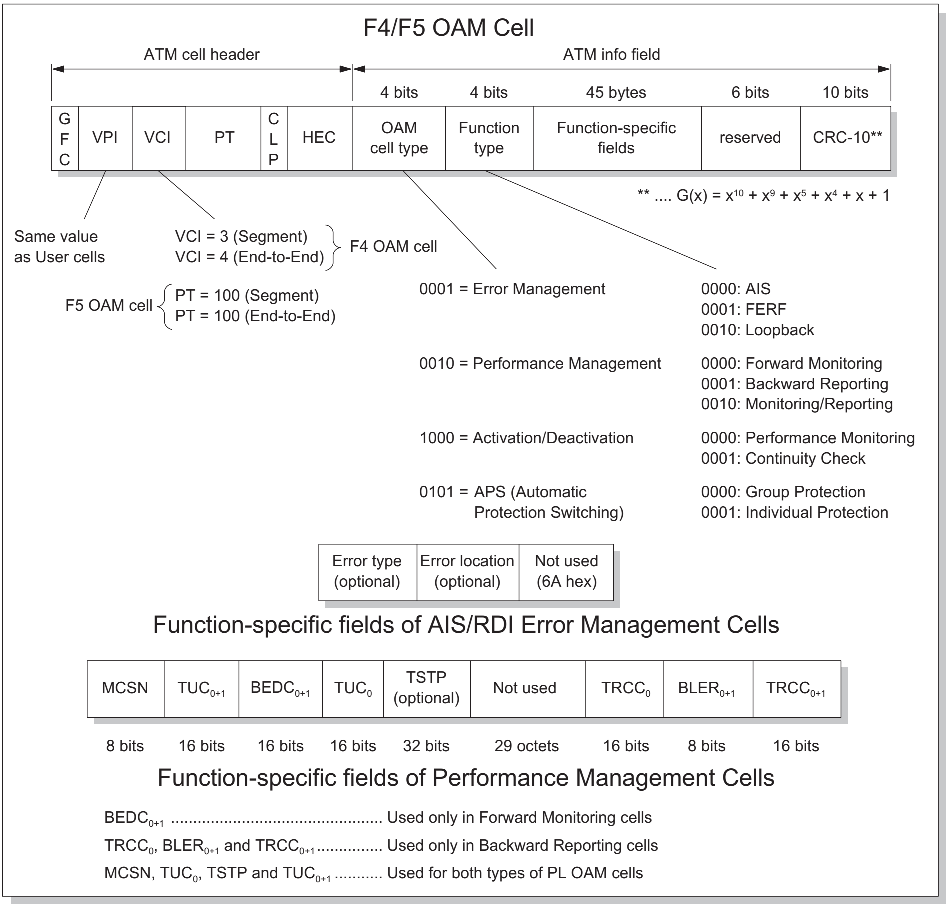

Figure 10.30 F4 and F5 OAM cell format

Fault management OAM cells are used to detect and localize communication faults and report them to the stations concerned. Performance management OAM cells detect parameters such as the cell block ratio, cell loss ratio, or the cell misinsertion rate, and thus provide information about the performance of a connection. APS coordination protocol OAM cells are used to manage ATM protection switching. Activation/deactivation cells are used to start and stop OAM functions such as error or performance management. The fourth OAM cell type, system management cells, has no defined purpose in the specifications: its usage is left to individual manufacturers' implementations.

OAM Fault Management (AIS/RDI)

In analogy to *SDH/SONET*, fault management in the ATM layer involves two kinds of alarm signals, Alarm Indication Signal (AIS) and Remote Defect Indication (RDI). The VP-AIS or VC-AIS is sent by the virtual path or virtual channel node that detects the fault downstream to all network nodes directly affected. The VP/VC-AIS OAM cell is transmitted periodically (approximately 2 second intervals) until the fault is eliminated. Immediately after the AIS has been sent, a VP-RDI or VC-AIS OAM cell is sent upstream to the endpoint of the virtual path or virtual channel connection concerned. This OAM cell is likewise repeated periodically as long as the AIS condition persists. VP-AIS and VP-RDI messages are always transmitted by means of cells with VCI = 4, while VC-AIS and VC-RDI messages are transported in cells with PT = 101.

A fault at the virtual path level inevitably affects the virtual channels contained within that virtual path. This faulty virtual path may terminate at a switch before the endpoint of these virtual channels, which will continue in one or more new virtual paths determined by switching. It is therefore necessary that the virtual channels affected continue to be notified of the fault condition. This is not done by propagating AIS or RDI on the F4 flows of the new virtual paths, which would be misleading because the new virtual paths probably are not themselves faulty. Instead, the fault management system causes the fault condition to propagate upwards to the F5 (virtual channel) flow level at the endpoints of the faulty virtual path and VC-OAM cells now carry RDI or AIS fault signals over all virtual channels that were contained in the original faulty virtual path. This can obviously result in an explosion of fault indications, of course, but this is inevitable if full fault management is to be achieved.

Fault management uses two mechanisms to detect fault conditions: continuity checks (CC) and loopback tests. Continuity check cells can be inserted in the user data stream at regular intervals to provide continuous verification of the availability of a connection. The ATM network nodes along the connection path can then monitor the presence of these cells. If expected CC cells are not received, loss of continuity (LOC) is signaled by AIS OAM cells. The insertion of

CC cells is useful where user traffic is intermittent because the absence of user traffic does not necessarily mean that the virtual connection has been terminated. Without CC cells, there would be no cells at all belonging to that virtual connection for perhaps long periods so, if a fault had occurred during this period, there would be no other mechanism to recognize this.

Loopback Cells

Loopback cells are used to verify connections to specific sections of the ATM network. There are five distinct kinds of loopback tests: end-to-end, access-line, inter-domain, network-to-endpoint, and intra-domain.

End-to-end loopback cells sent to one endpoint of a VP or VC connection are sent back to the originating endpoint. The connection may cross several subnetworks or "operator domains". In this way the entire connection is tested from end system to end system. By contrast, access line loopback cells are returned by the first ATM network node that receives them. This tests exactly one connection section or segment. Inter-domain loopback cells are reflected by the first network node of a neighboring operator domain. This makes it possible to test the connection to the neighboring network. Network-to-endpoint loopback cells can be used by a network operator to test the connection out of the network to an endpoint in an adjacent network. Finally, intra-domain loopback cells can be sent from a segment node to any other node in another segment within the same operator domain. This tests the data flow across a certain sequence of segments within a domain.

Loopback cells are F4 or F5 OAM cells that contain the value 0100 in the function type field (see Figure 10.30). The function-specific fields are the Loopback Indication, the Correlation Tag, the Loopback Location ID and the Source ID. All loopback cells other than end system to end system loopback cells are segment cells; the point at which they loop back is determined by the Loopback Location ID field. If a station receives a loopback OAM cell with a value other than 0 in the Loopback Indication field, it must return the cell within one second. Stations that transmit loopback cells must do so at such an interval that the management cell traffic is less than one percent of the capacity of each virtual channel or virtual path involved in the connection. It should be obvious from the above that only the loopback OAM cells are looped; related user cells are unaffected, so loopback test can be performed safely at any time.

OAM Performance Management

In addition to faults, the performance of the individual VPs and VCs can also be monitored by the periodical insertion of special performance management OAM cells in the user cell stream. Analysis of the special measurement data contained in these cells (cell sequence number, total user cell count, time stamp, cell loss count) yields direct information about the operating condition of the given ATM connection. OAM performance management involves two distinct functions known as forward performance monitoring and backward reporting. If both functions are activated, then information is transmitted in both directions when OAM performance information about a given cell block is determined. The individual fields of performance management cells are the following:

- Monitoring cell sequence number (MCSN, 8 bits).
- Total user cell count for the CLP_{0+1} user cell flow (TUC_{0+1}, 16 bits). This is the number of cells transmitted with cell loss priority of 0 or 1 up to the time of OAM cell insertion.
- Total user cell number for the CLP_0 user cell flow (TUC_0, 16 bits). Identical to TUC_{0+1}, but counting only cells with CLP = 0.
- Block error detection code for the CLP_{0+1} user cell flow ($BEDC_{0+1}$, 16 bits). This field is used only in forward monitoring cells. It contains a BIP-16 checksum of the information fields of the cells transmitted since the last forward monitoring cell.
- Time Stamp (TSTP, 32 bits). This is the time at which the OAM cell was inserted in the cell stream. Currently the time stamp is optional; its use has yet to be fully defined.
- Total received cell count for the CLP_{0+1} user cell flow ($TRCC_{0+1}$, 16 bits). This field is used only in backward reporting OAM cells. It contains the number of cells received before the corresponding forward monitoring cell.
- Total received cell count for the CLP_0 user cell flow ($TRCC_0$, 16 bits). This field is used only in backward reporting OAM cells. It contains the number of cells with CLP = 0 received before the corresponding forward monitoring cell.
- Block error result (BLER, 8 bits). This field is used only in backward reporting OAM cells. It contains the number of incorrect parity bits determined on the basis of the BIP-16 code in the corresponding forward monitoring cell.

Performance monitoring is done over blocks of cells, the size of the blocks N being related to the peak cell rate of the virtual connection being monitored. A forward performance monitoring OAM cell is sent after each N user cell providing information about the cell block. At the receiving end, the same calculation is performed on the cells of the block and the results compared with the information in the forward performance monitoring OAM cell. The results are reported to network management and, optionally, via a backward reporting OAM cell to the originating end.

The reason for having both forward performance monitoring and backward reporting cells is that the endpoints of the monitored virtual path or virtual channel may lie in the domains of different network operators. Normally, one service provider does not have access to the information in the network management system of another network operator, so the backward reporting cell provides the means to transfer the results back across the domain boundary/boundaries between the relevant network operators. When no domain boundaries are crossed, backward reporting cells need not be sent because network management can access all nodes in their networks.

Activation and Deactivation of OAM Functions

The performance monitoring and continuity check functions are started and stopped by means of special activation and deactivation cells. Figure 10.31 shows the format of these cells.

Message ID	Direction of action	Correlation tag	PM block size A-B	PM block size B-A	Not used
6 bits	2 bits	8 bits	4 bits	4 bits	336 bits

Figure 10.31 Function-specific fields of activation/deactivation cells

The message-ID field contains the various commands of the OAM activation/deactivation cell:

Activate	000001
Activation confirmed	000010
Activation request denied	000011
Deactivate	000101
Deactivation confirmed	000110
Deactivation request denied	000111

The correlation tag serves to correlate commands with the corresponding responses. The direction-of-action field specifies the direction of transmission of the activated OAM cells. A–B indicates the direction away from the activator or deactivator; B–A indicates transmission toward the activator. The fields PM block size A–B and PM block size B–A specify the cell block length to be used for a performance measurement.

10.1.8 The ATM Adaptation Layer (AAL)

The AAL maps data structures of higher application layers to the cell structure of the ATM layer and provides appropriate control and management functions. In order to meet the different requirements of various services, four AAL types were originally defined: AAL Type 1 the real-time sensitive services with constant bit rates; AAL Type 2 for real-time sensitive services with variable bit rates; AAL Types 3 and 4 for connection-oriented and non-connection-oriented transmission of non-real-time sensitive data. Later, AAL Type 5 was defined by the ATM Forum as a simplified version of AAL3 and was quickly adopted by the ITU-T. Early on, the distinction between connection-oriented and non-connection-oriented data communication in the AAL was found to be unnecessary, and AAL Types 3 and 4 were accordingly merged into AAL Type 3/4. Note that AAL types are never mixed within a single virtual channel for reasons that will become obvious.

Parameter \ Service	Class A	Class B	Class C	Class D
Time Compensation	required		not required	
Bit rate	constant	variable		
Communication mode	connection-oriented			connectionless
Example	circuit emulation	wireless voice communication	connection-oriented data communication	connectionless data communication
AAL type	AAL1	AAL2	AAL3, AAL5	AAL4

Figure 10.32 Service classes and AAL types

10.1.8.1 AAL Type 0

The AAL Type 0 indicates the absence of any AAL capability. The application data is inserted directly in the payload fields of the ATM cells and transmitted. Strictly speaking, AAL0 is thus not an AAL type at all; the communication mechanisms are already cell-based and the adaptation layer functions don't exist.

10.1.8.2 AAL Type 1

The Type 1 ATM Adaptation Layer (ITU-T Recommendation I.363.1) serves to transport data streams with constant bit rates (these include all the interfaces in the PDH hierarchy: T1, E1, T3, etc.), and provide them to the destination node in synchronization with the original transmission clock. This requires that the

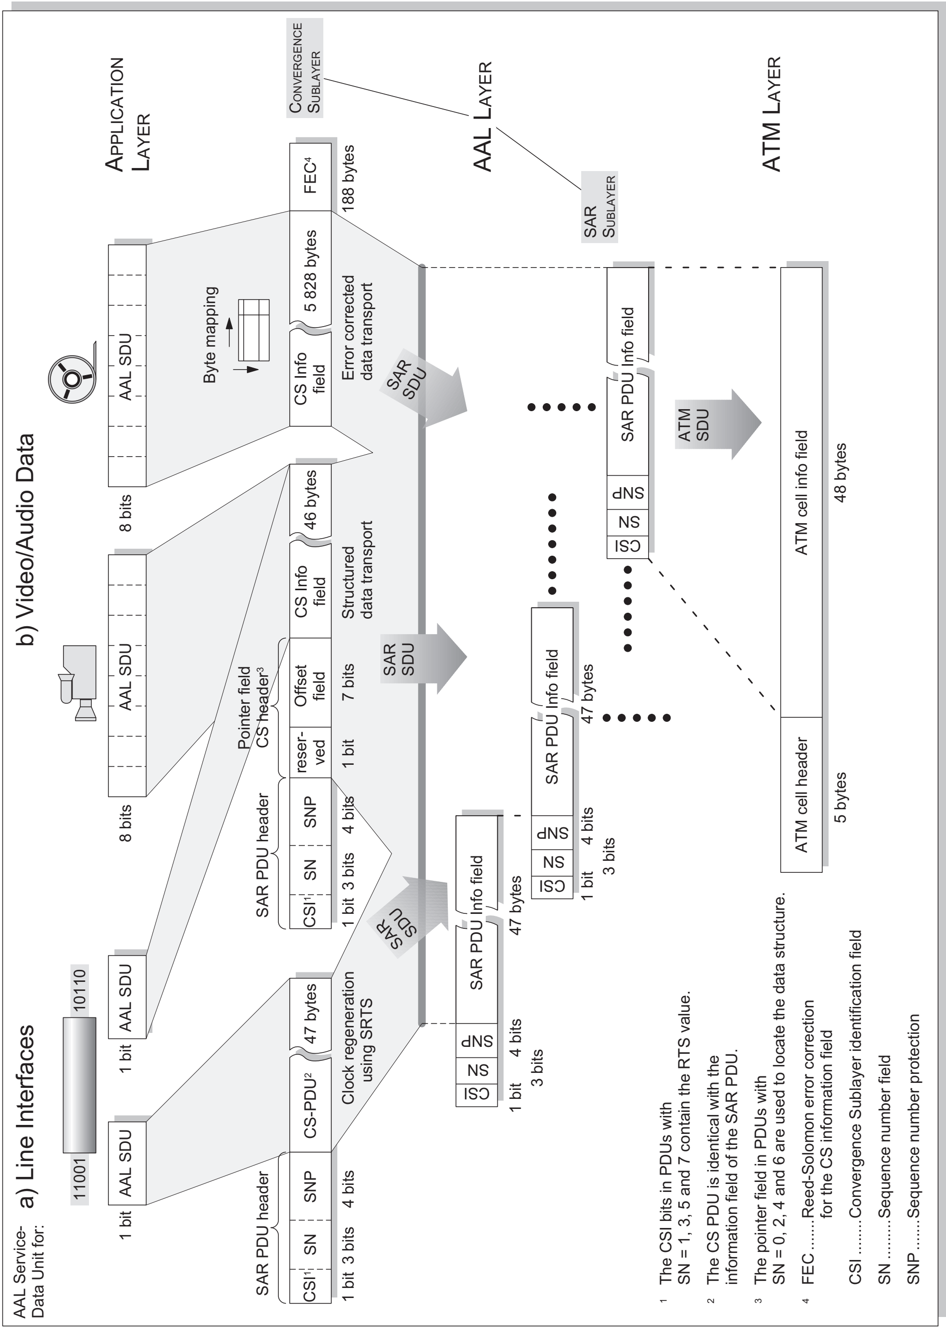

Figure 10.33 Structure of AAL Type 1

ATM network transport not only the data, but also the clock information. For this reason the AAL1 protocol is capable of transporting both continuous bit streams and byte-structured data, such as data based on an 8-kHz sampling interval. Lost or erroneous data is not repeated or corrected. Events such as cell loss or the transmission of incorrect service data units (SDUs), loss of synchronization or clock signal, buffer overflow or the occurrence of invalid AAL header information (AAL Protocol Control Information, or AAL-PCI) are passed from the user layer to the management layer. The AAL1 is composed of two sublayers, the Segmentation and Reassembly Sublayer (SAR) and the Convergence Sublayer (CS). The data to be transported is first broken into 47-byte data blocks (CS PDUs). Each such data block is given a one-byte header that contains a 3 bit sequence count and CRC-3 protection with parity plus a bit (the CSI bit) that is used to carry a multi-cell "synchronous residual time stamp" (SRTS). The 48-byte SAR PDU is then transported in the data field of an ATM cell. At the receiving station, the original data rate of the transmitting station can be synchronously regenerated using the clock information in the CS PDU header. In simple terms, the SRTS process compares the clock rate of the encapsulated service (that is the constant bit rate service carried over AAL1) with the physical layer clock (for example, the SDH/SONET clock) at the source and derives the SRTS, which is transported to the receiver. The physical layer clock and the SRTS are then processed to recover the service clock. An assumption is that there is a fixed relationship between physical layer clocks at both ends of the virtual channel carrying the AAL1 service, normally the case within national networks at least.

10.1.8.3 AAL Type 2

The Type 2 ATM Adaptation Layer (ITU-T Recommendation I.363.2) is used for efficient transportation of delay-sensitive, narrow-band applications with variable bandwidth (such as telephony). The network must guarantee certain QoS parameters, such as maximum cell delay or cell loss ratio, for each connection while making varying bandwidth available. AAL2 guarantees the QoS parameters for each connection using the QoS mechanisms of the underlying ATM layer. AAL2 itself only specifies the format of the short AAL2 SDUs (or "minicells") optimized for real-time applications, and their transport within ATM cells. The bi-directional AAL2 connections can be set up either as PVCs or as SVCs, but provide only non-guaranteed transport services. Corrupt or lost CPS SDUs are neither corrected nor repeated. The data to be transported is first filled into Common Part Sublayer (CPS) packets, which consist of 3 header bytes and 1 to 45 or 64 bytes of user data. AAL2 connection setup, the assignment of Channel IDentifiers (CIDs) and the negotiation of CPS service parameters takes place on AAL2 channel 1 using the ANP protocol (AAL2 Negotiation Procedure).

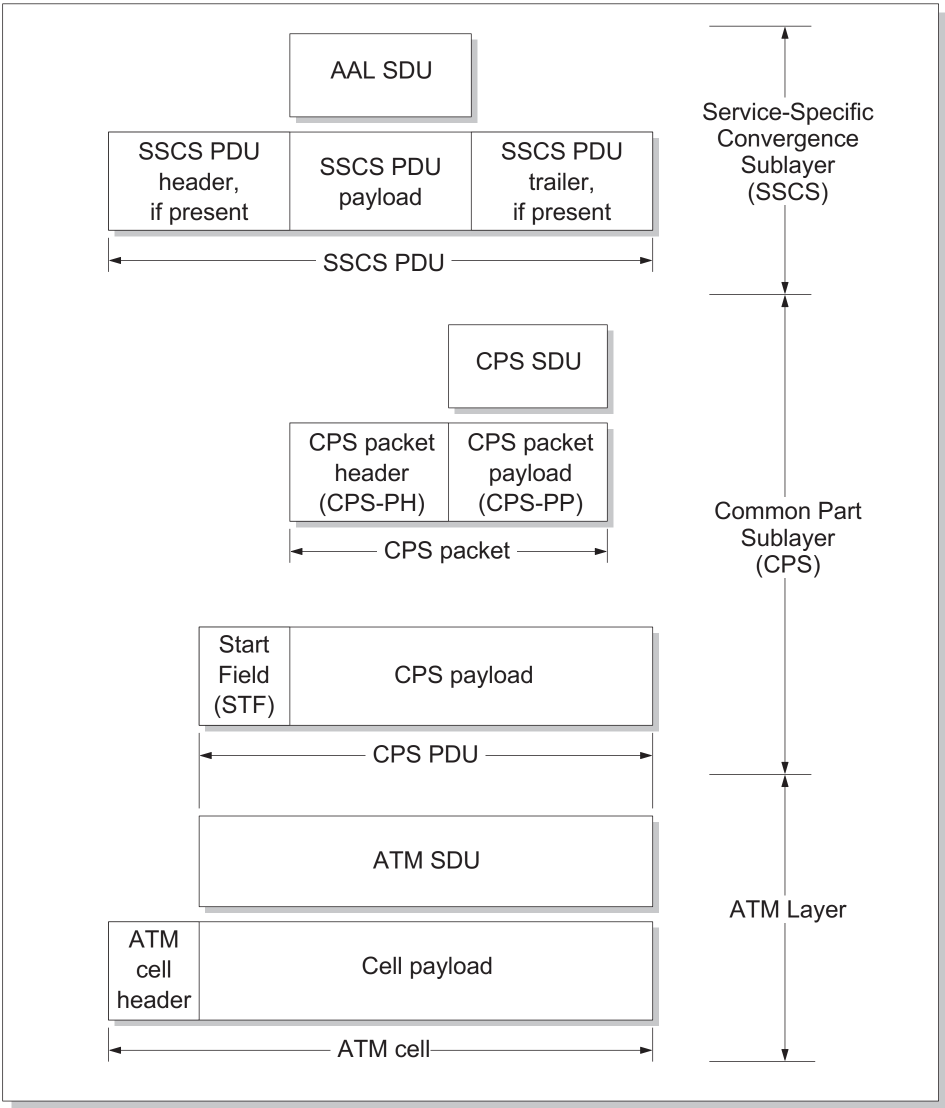

Figure 10.34 Structure of AAL2

The CPS packets (mini-cells) are then inserted in CPS PDUs, which consist of 1 header byte and 47 payload bytes. The CPS PDUs in turn are transported in the payload fields of ATM cells. The start field (STF) comprises an offset value pointing to the start of the next CPS packet to aid CPS packet delineation recovery in the event that an ATM cell has been lost. Figure 10.34 shows a diagram of the structure of AAL2.

AAL2 Error Messages

The following error messages in Figure 10.35 are passed to the layer management in case of transmission errors in AAL2:

Error number	Description
0	The start field in the CPS PDU indicates a parity error. The entire PDU is discarded.
1	The sequence number in the start field is invalid. If the offset is 0, the entire PDU is discarded. Otherwise, processing continues at the byte indicated by the offset.
2	The number of bytes of an overlapping CPS packet does not match the parameters in the start field. If the offset is less than 47, processing continues at the byte indicated by the offset.
3	The offset value is greater than 47. The entire PDU is discarded.
4	The HEC checksum of a CPS packet indicates a transmission error in the header. The corresponding part of the CPS PDU is discarded.
5	The padding bits extend into the following CPS PDU. This extension is ignored and not processed.
6	A CPS packet fragment was received and had to be discarded before it could be reassembled.
7	The HEC checksum of a CPS packet that extends across a CPS PDU boundary indicates a transmission error in the header. If the offset is less than 47, processing continues at the byte indicated by the offset.

Figure 10.35 AAL2 error messages to the layer management

The history of AAL2 is complex. It was originally proposed by the ITU-T in the late 1980s as the adaptation layer to handle Class B services (variable bit-rate, connection-oriented services, specifically video), then withdrawn. It returned through the initially separate then later joint efforts of the ATM Forum and the ITU-T in the latter half of the 1990s to solve the problems inherent in AAL1 for low (and sometimes variable) bandwidth delay sensitive services, such as (compressed) voice, and the inefficiency of AAL5 for handling very short data packets. It was at one point nearly called AAL6 by the ATM Forum before it was concluded that it really did fulfill the original requirements for AAL2 defined a decade earlier. At the time of this writing, therefore, two service types have been defined for AAL2: trunking and SAR. ITU-T Recommendation I.366.2 describes the “AAL type 2 Service Specific Convergence Sublayer for Trunking”—the circuit emulation service used for the transportation of narrow-band, delay-sensitive traffic. AAL2 also provides the means to handle

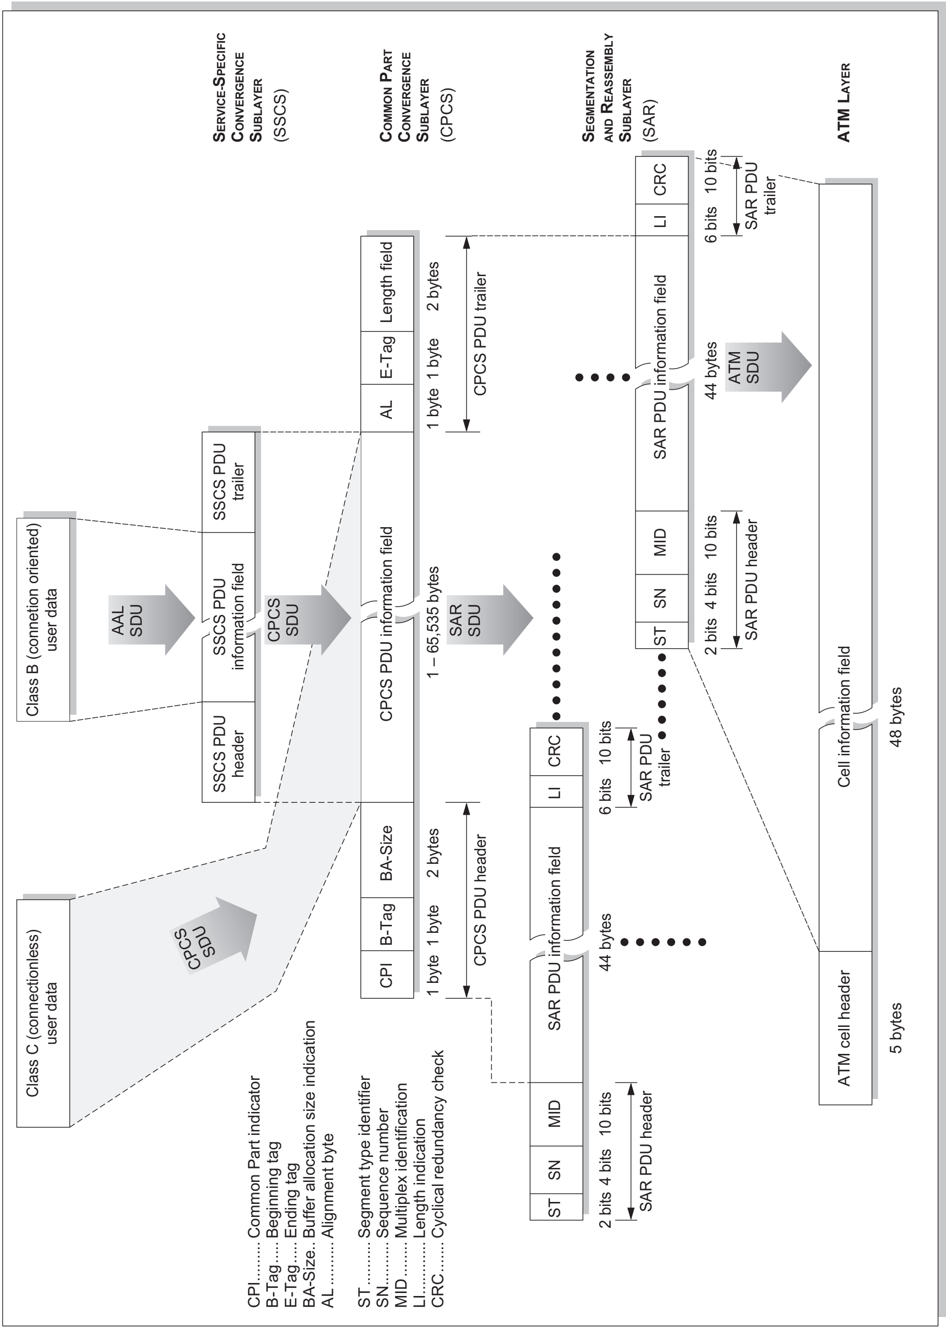

Figure 10.36 Structure of AAL3/4

very short data packets more efficiently, reducing the "cell tax" (inefficient use of ATM cell payloads) inherent in AAL3/4 and AAL5. ITU-T Recommendation I.366.1 describes the "Segmentation and Reassembly Service Specific Convergence Sublayer for AAL type 2". This allows the encapsulation of higher layer variable-length data packets (1 to 65,535 bytes) over AAL2 "mini-cells" in almost exactly the same way that the SAR function of AAL5 allows the encapsulation of data packets over ATM cells (see AAL type 5 in the following discussion). Even the AAL2 SSTED-PDU trailer structure is similar to that of the AAL5 CPCS PDU trailer: it is 8 bytes long and provides a length field and CRC-32 field in the same positions, thus allowing the re-use of hardware designs for processing the SAR function. Third generation wireless services make use of this AAL2 SAR process for carrying data services.

10.1.8.4 AAL Type 3/4

AAL Type 3/4 (ITU-T Recommendation I.363.3) specifies connection-oriented and non-connection-oriented transportation of data packets in ATM networks. Both point-to-point and point-to-multipoint connections can be set up. The AAL3/4 protocol is thus suitable for transportation of non-connection-oriented communications services, such as SMDS/CBDS (Switched Multi Megabit Data Service / Connectionless Broadband Data Service) metropolitan area networks or Frame Relay. Like AAL Type 1, the AAL3/4 protocol consists of two sublayers, the Segmentation and Reassembly (SAR) sublayer and the Convergence Sublayer (CS), although the CS includes both a Common Part Convergence Sublayer (CPCS) and a Service Specific Convergence Sublayer (SSCS). The variable-length data packets (1 to 65,535 bytes) of the application building on AAL3/4 are first padded to an integer multiple of 32 bits to permit an efficient, hardware-based implementation of the AAL processes, and a header and trailer are added. The resulting CS PDU is then split into 44-byte fragments, each of which receives another header and trailer, and is then passed to the ATM layer as a 48-byte SAR PDU. In reality this AAL is now only used for legacy SMDS services, which are themselves obsolescent. AAL5 has overtaken AAL3/4 for interworking with Frame Relay and for the support of almost all other services because it is so much more efficient.

10.1.8.5 AAL Type 5

AAL Type 5 was invented by the ATM Forum in the beginning of the 1990s to simplify the handling of data over ATM. The ITU-T later adopted it and it is now defined in ITU-T Recommendation I.363.5. AAL5 corresponds to a highly simplified implementation of AAL Type 3/4. Like AAL3/4, it is suitable for connection-oriented and non-connection-oriented data communications. And like AAL3/4, AAL5 also consists of a Segmentation and Reassembly (SAR) sublayer and a Convergence Sublayer (CS), with the CS subdivided into a Common Part

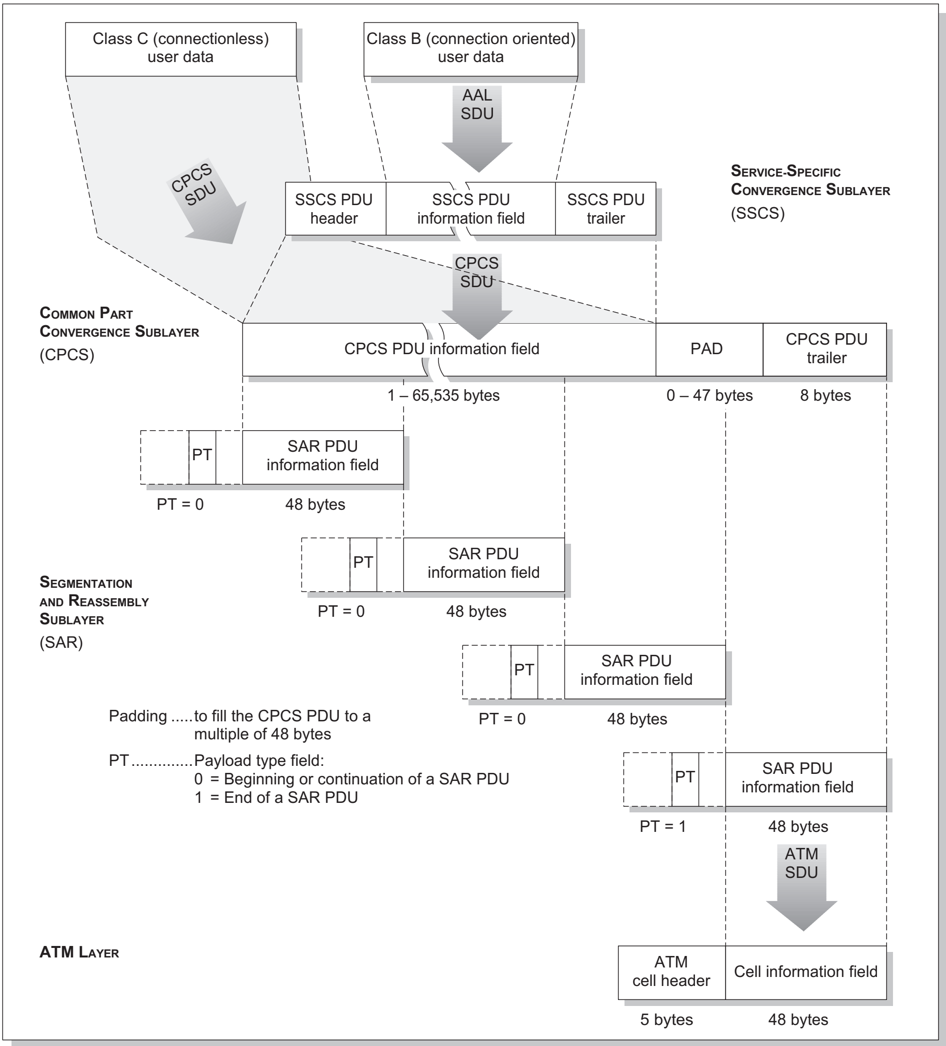

Figure 10.37 Structure of AAL5

Convergence Sublayer (CPCS) and a Service Specific Convergence Sublayer (SSCS).

The processes within the AAL Type 5 sublayers are significantly simpler than in AAL3/4. There is no mechanism for cell multiplexing, for example. All cells that belong to an AAL5 CS PDU are transmitted in one sequential cell stream. The variable-length data packets (1 to 65,535 bytes) of the application building on AAL5 are first padded to an integer multiple of 48 bytes (this ensures that no partially filled cells are transmitted after segmentation) and a trailer is added (no header). The resulting CS PDU is then broken into 48-byte segments that fit directly into the data field of ATM cells. The PT field of the ATM header is used to identify whether further segments follow, or whether the data field contains the end of the CS PDU. Trailer-based PDUs are efficient to process because the large amounts of data often involved can be further processed without having to move them in memory-based implementations. The next section will describe the SAAL, which is also a trailer-based PDU building on AAL5; processing of the higher sublayers of the SAAL is thus simplified. Additionally, AAL5 allows processing on 32-bit boundaries – hence the 8-byte length of the trailer (2 x 32 bits); there is really no need for more than 7 bytes in the trailer, the CPCS-UU and CPI fields (8 bits each) are rarely used and could have shared a single byte.

10.1.8.6 The Signaling ATM Adaptation Layer (SAAL)

As in narrow-band ISDN, B-ISDN signaling is also handled in signaling virtual channels that are separate from the user connections. The Signaling AAL, ITU-T Recommendations Q.2100 – Q.2144 (SAAL) is the AAL used for all ITU-T signaling protocols and ATM Forum UNI signaling protocols versions 3.1 and 4.0, (Version 3.0 used a prenormative version of the SAAL; UNI 3.0 was published about a year before the ITU-T had completed Recommendation Q.2931; UNI 3.0 is now obsolete.). Figure 10.38 shows the protocol layer model for ITU-T UNI and NNI signaling in ATM networks, with the position of the SAAL in the protocol model.

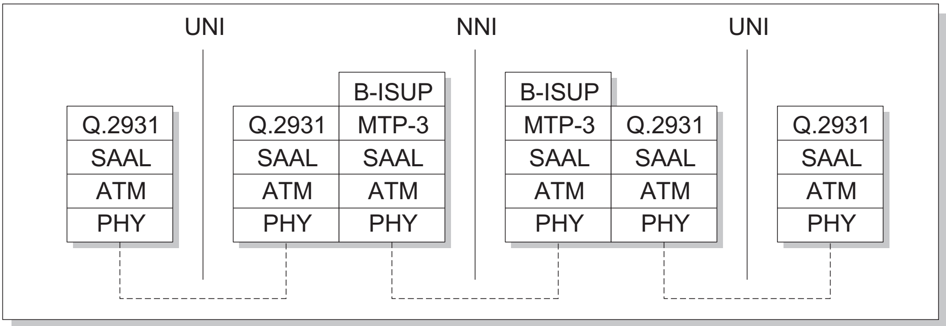

Figure 10.38 Protocol layer model for UNI and NNI signaling

The SAALs for UNI and NNI signaling have several features in common, but the NNI-SAAL has a more complex structure due to the greater number of mechanisms that need to be provided for MTP-3 and B-ISUP. The purpose of the SAAL is to provide the actual signaling layers situated above it in the protocol hierarchy (Q.2931; MTP-3 and B-ISUP) with a reliable transportation service, because these signaling protocols have no error compensation mechanisms themselves. Fault tolerance is provided by SAAL's SSCS sublayer using the Service-Specific

Connection-Oriented Protocol (SSCOP), which builds on the CPCS and SAR sublayers of AAL3/4 or AAL5. Because the CPCS sublayer of AAL3/4 and AAL5 is only able to perform unassured information transfer, a substantial part of the SSCOP protocol is concerned with procedures to guarantee transmission of the SSCOP information field contents. This is analogous to the function of the TCP layer used with IP for guaranteeing reliable transport of data across IP networks. The given signaling layer's requests are translated to SSCOP by an appropriate Service-Specific Coordination Function (SSCF). Although specified as an option for use, in practice AAL3/4 is not used for signaling.

The SSCOP protocol defines 15 distinct PDUs that are used to implement different functions. The PDUs Begin (BGN) and Begin Acknowledge (BGAK) are used to set up the SSCOP connection between two stations and reset the send and receive buffers of the receiving station. Assured data communication can then take place in two ways: the data packets can be sequentially numbered (SD PDUs), and each data packet may contain an individual request for confirmation of receipt (SDP PDUs). Such confirmation is sent by the receiver in a STAT PDU. After the receiver has analyzed the sequential numbers of the PDUs received, it can send USTAT PDUs if necessary to report lost PDUs, indicating the number range missing. Data can also be sent in an unassured mode using Unnumbered

Function	Description	PDU Name	Value
Establishment	Request initialization	BGN	0001
	Request acknowledgement	BGAK	0010
Release	Disconnect command	END	0011
	Disconnect acknowledgement	ENDAK	0100
Resynchronization	Resynchronization command	RS	0101
	Resynchronization acknowledgement	RSAK	0110
Reject	Connection reject	BGREJ	0111
Recovery	Recovery command	ER	1001
	Recovery acknowledgement	ERAK	1111
Assured data transfer	Sequenced connection mode data	SD	1000
	Transmitter state information with request for receiver state information	POLL	1010
	Solicited receiver state information	STAT	1011
	Unsolicited receiver state information	USTAT	1100
Unacknowledged data transfer	Unnumbered user data	UD	1101
Management data transfer	Unnumbered management data	MD	1110

Figure 10.39 SSCOP PDUs and their functions

Data (UD) PDUs. Figure 10.39 lists the 15 types of SSCOP PDUs and their functions.

The SSCOP Timers

The SSCOP protocol defines four timers that control the protocol process. These are the POLL, KEEP-ALIVE, NO-RESPONSE and CONNECTION CONTROL (CC) timers.

The POLL timer monitors the maximum time interval between the transmissions of successive POLL PDUs while SD or SDP PDUs are being transmitted.

Error Type	Error Code	PDU or event that triggered the error
Receipt of unsolicited or inappropriate PDU	A	SD PDU
	B	BGN PDU
	C	BGAK PDU
	D	BGREJ PDU
	E	END PDU
	F	ENDAK PDU
	G	POLL PDU
	H	STAT PDU
	I	USTAT PDU
	J	RS
	K	RSAK PDU
	L	ER
	M	ERAK
Unsuccessful retransmission	O	VT(CC) ≥ MaxCC
	P	Timer_NO_RESPONSE expired
Other list elements error type	Q	SD or POLL, N(S) error
	R	STAT N(PS) error
	S	STAT N(R) or list elements error
	T	USTAT N(R) or list elements error
	U	PDU length violation
SD loss	V	SD PDUs must be retransmitted
Credit condition	W	Lack of credit
	X	Credit obtained

Figure 10.40 SSCOP error codes

When no SD or SDP PDUs are being transmitted, the KEEP-ALIVE Timer controls the interval between successive POLL PDUs. The NO-RESPONSE timer specifies the maximum delay between two PDUs of the type POLL or STAT. The value of the NO-RESPONSE timer must be greater than both that of the KEEP-ALIVE timer and that of the POLL timer. Furthermore, the value of the NO-RESPONSE timer must be more than twice the signal delay over the connection concerned.

The CONNECTION CONTROL timer determines the maximum delay between the transmission of two BGN, END or RS PDUs if no answering PDU is received. The CC timer must also have a value greater than twice the signal delay over the connection concerned.

SSCOP Error Messages to the Management Layer

Error codes are used to forward SSCOP error events to the management layer. The table in Figure 10.40 lists the various SSCOP error types.

SAAL Connection Setup

The connection setup between two SAAL system components is triggered by an AA-ESTABLISH message from the SSCF sublayer. This command contains the parameters SSCOP User to User Parameter (SSCOP-UU) and Buffer Release (BR), which are used in generating the SSCOP sublayer's BGN message. The receiver decodes the BGN message and passes an AA-ESTABLISH.ind to the receiving SSCF. This sublayer responds with an AA-ESTABLISH.res command, which likewise contains the parameters SSCOP-UU and BR. The resulting BGAK message is sent back to the originating SSCOP, which passes an AA-ESTABLISH.conf to the initiating SSCF (see Figure 10.41).

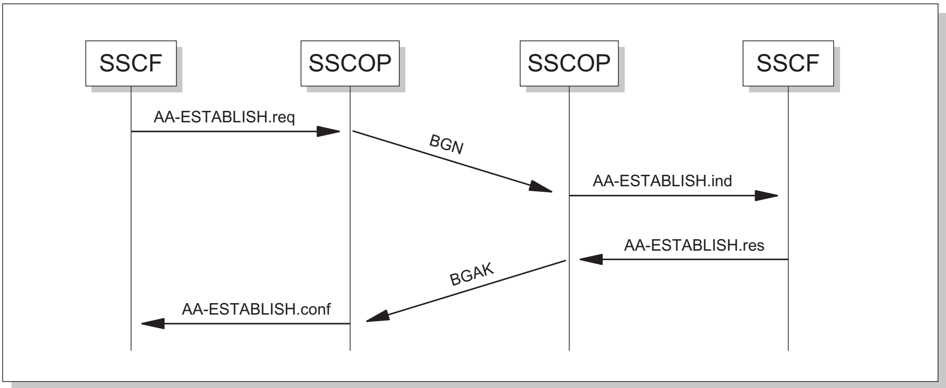

Figure 10.41 SAAL connection setup

10.1.9 ATM Signaling

Signaling in ATM refers to all processes necessary to set up a connection between two or more stations. In contrast to conventional signaling mechanisms, signaling in ATM networks is an extremely complex procedure. All kinds of traffic parameters—such as the AAL type, streaming or message mode, assured or unassured transfer, cell rates (peak cell rate and sustainable cell rate), cell loss ratio, cell delay, cell delay variation tolerance and maximum burst size—must be negotiated and guaranteed over all segments of the connection path. Moreover, novel connection-oriented topologies, such as point-to-multipoint or broadcast connections, must also be handled. For this reason the existing UNI and NNI signaling protocols, Q.931 and Q.933 or ISUP, on which B-ISDN signaling is based and which were originally developed for narrow-band ISDN, have undergone major extension for ATM networks to produce Q.2931 or B-ISUP. Signaling at the user interface takes place either based on ITU-T Recommendation Q.2931 (B-DSS2) or by means of one of the ATM Forum specifications, UNI 3.0, 3.1 or 4.0. At the network interface, public networks use B-ISUP (ITU-T Recommendation Q.2761 – Q.2764) or the corresponding ATM Forum protocols, B-ICI or PNNI.

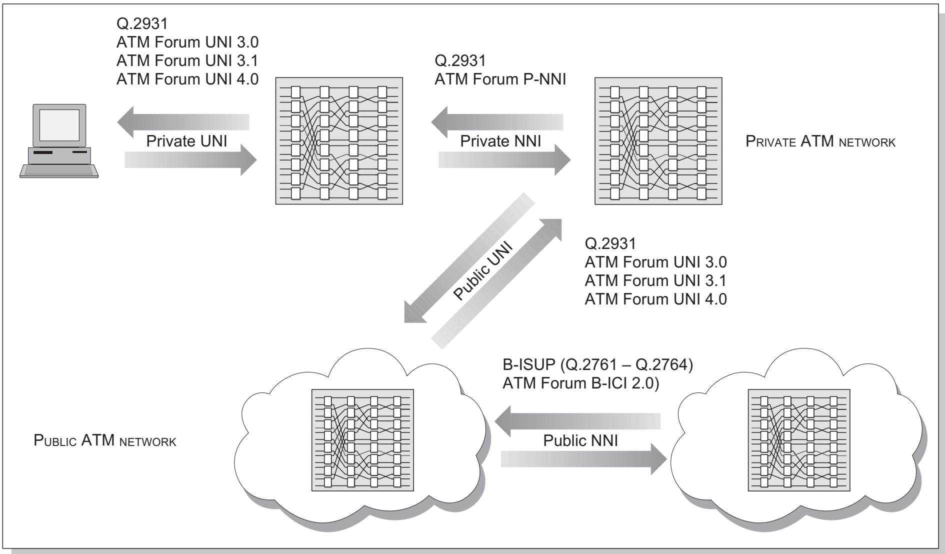

Figure 10.42 ATM signaling in public and private ATM networks

Connection setup can take place either on the reserved signaling virtual channel, VPI-VCI = 0-5, or over VCs selected by means of meta-signaling. Any signaling VCs can be set up without an existing AAL connection, because the simple meta-signaling commands can be transferred in one cell. Once the signaling virtual channel has been determined, the AAL for the signaling virtual channel (SAAL) is established. At this point the signaling protocol becomes active. In practice, however, the signaling virtual channel VPI-VCI = 0-5 is generally used, and meta-signaling is not necessary.

10.1.9.1 The Q.2931 Message Format

Q.2931 defines a total of 15 UNI signaling messages, which can be classified in groups for connection setup, connection clear-down and other messages. For all ordinary signaling processes, the following ten message types are available:

Connection Setup:

- ALERTING
- CALL PROCEEDING
- CONNECT
- CONNECT ACKNOWLEDGE
- SETUP

Connection Clear-Down:

- RELEASE
- RELEASE COMPLETE

Other:

- NOTIFY
- STATUS
- STATUS ENQUIRY

For interoperability with narrow-band ISDN, three more messages have been defined in addition to those listed previously:

Connection Setup (Additional Messages for N-ISDN Interworking):

- SETUP ACKNOWLEDGE
- PROGRESSING

Other (Additional Messages for N-ISDN Interworking):

- INFORMATION

Furthermore, the messages RESTART and RESTART ACKNOWLEDGE are defined for the purpose of requesting a new connection setup attempt (in the event that a connection enters an undefined state, for example). These messages may be used only with the global call reference 0. Every signaling message consists of the following five sections, called information elements:

- Protocol discriminator
- Call reference
- Message type
- Message length
- Message-specific information elements

The first four information elements are mandatory and must be present in every Q.2931 message. The usage of other information elements is dependent on the message type. The diagram in Figure 10.43 illustrates the structure of a UNI signaling message.

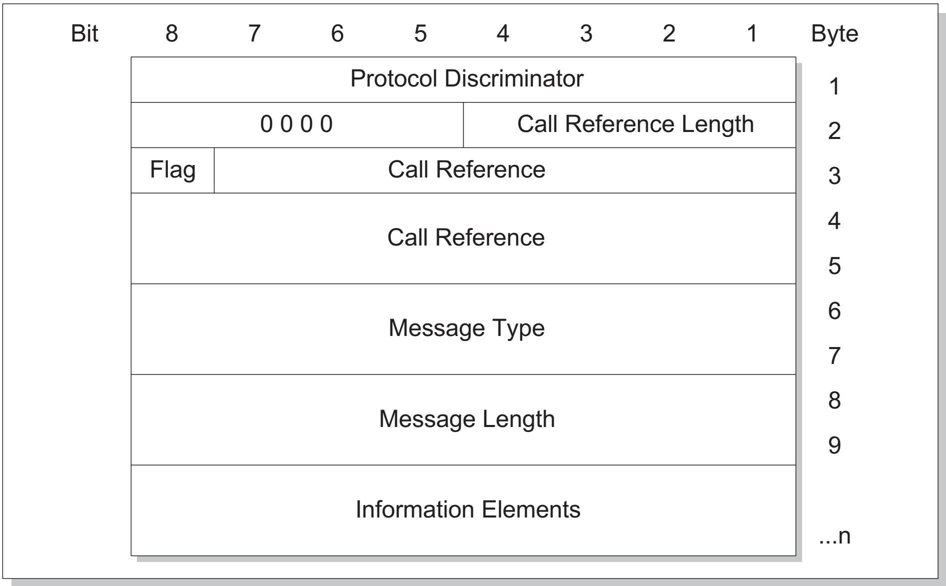

Figure 10.43 Q.2931 message format

Protocol Discriminator

The protocol discriminator is the first element in every Q.2931 message. It identifies the protocol used and is coded as 00001001 (decimal 9) for Q.2931. The ATM Forum signaling variants (UNI 3.x, 4.0) share the same protocol discriminator value.

Bit	8	7	6	5	4	3	2	1
	0	0	0	0	1	0	0	1

Protocol discriminator value for Q.2931 messages

Bit 8765 4321	
0000 0000 to 0000 0111	Reserved
0000 1000	Q.931/ I.415 user network call control
0000 1001	Q.2931 user network call/connection control
0001 0000 to 0011 1111	Other Layer 3 protocols (X.25 etc.)
0100 0000 to 0100 1111	National use
0101 0000 to 1111 1110	Other Layer 3 protocols (X.25 etc.)

Figure 10.44 The protocol discriminator field

Call Reference

The call reference serves to associate Q.2931 messages with a given connection. When a new connection is established, all messages concerning that connection have the same call reference value. When the connection has been cleared down, the call reference is released and can be used again. The same call reference can be used by two connections within an ATM virtual channel only if the respective connection setups take place in opposite directions. The length of the call reference field is measured in bytes; the default length is 3 bytes. The call reference flag identifies the sending and receiving stations. The station that indicates the connection always sets this flag to 0 in its messages, while messages originating from the receiving station have the flag set to 1. The call reference value 0 is known as the global call reference and refers to all connections within a signaling virtual channel.

Message Type

This field indicates the message type. All message types except SETUP ACKNOWLEDGE and INFORMATION are also supported by the corresponding ATM Forum specification UNI 4.0 (see Figure 10.46).

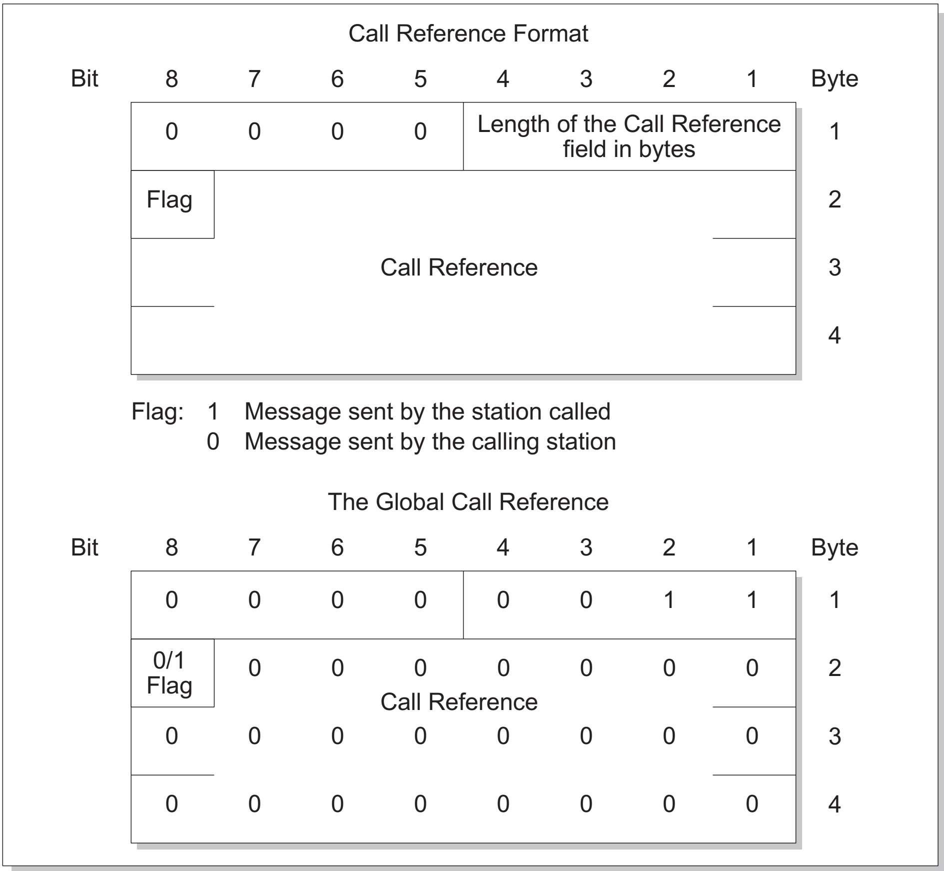

Figure 10.45 Call reference and global call reference

The Message Length Field

The message length field indicates the length of the signaling message in bytes, not counting the protocol discriminator, call reference, message type, and message length fields. The length field itself can be 1 or 2 bytes long. In the first byte, only the first 7 bits can be used. Bit 8 indicates whether the length field includes a second byte. If bit 8 is set to 0, the length field is continued in the following byte; otherwise the length field is a single byte. If the message contains no other information elements after the length field, the message length is coded as all 0s.

UNI Information Elements

After the four mandatory information elements, the various message types use specific information elements of varying lengths to fulfill their respective func-

Bit	8	7	6	5	4	3	2	1	Byte
	Message Type								1
	1 Ext.	0	0	Flag	0	1	Message Action Indicator		2

Bits 8765 4321	Q.2931 Message
0000 0000	Escape sequence for national message types
Connection set-up	
0000 0001	ALERTING
0000 0111	CONNECT
0000 1111	CONNECT ACKNOWLEDGE
0000 0010	CALL PROCEEDING
0000 0011	PROGRESS
0000 0101	SETUP
0000 1101	SETUP ACKNOWLEDGE
Connection clear-down	
0100 1110	RESTART ACKNOWLEDGE
0100 0110	RESTART
0101 1010	RELEASE COMPLETE
0100 1101	RELEASE
Other Messages	
0110 1110	NOTIFY
0111 1011	INFORMATION
0111 0101	STATUS ENQUIRY
0111 1101	STATUS

Figure 10.46 Q.2931 (B-DSS2) signaling messages

tions. Each of these information elements consists of an information element identifier, a length field, a compatibility indicator, and the actual information element contents. Figure 10.47 illustrates the general format of information elements and lists the specific information elements defined.

10.1.9.2 Connection Setup at the Caller's End (UNI)

The caller initiates the connection setup by transmitting a SETUP message containing the desired ATM virtual path, the ATM virtual channel, and quality-of-service and traffic parameters. If the network is able to provide the service requested, and if the specified ATM virtual path and virtual channel are available, the network answers with a CALL PROCEEDING message and forwards

Bit	8	7	6	5	4	3	2	1	Byte
	Information Element Identifier								1
	1 Ext.	Coding Standard		IE Instruction Field: Flag	Res	Action Indicator			2
	Information Element Length								3
									4
	Information Element Contents								5

Information Element Identifiers	
Bits 8 7 6 5 4 3 2 1	
0 1 1 1 0 0 0 0	Called Party Number
0 1 1 1 0 0 0 1	Called Party Sub-Address
0 1 1 1 1 0 0 0	Transit Network Selection
0 1 1 1 1 0 0 1	Restart Indicator
0 1 1 1 1 1 0 0	Narrow-Band Low Layer Compatibility
0 1 1 1 1 1 0 1	Narrow-Band High Layer Compatibility
0 1 1 0 0 0 0 0	Broadband Locking Shift
0 1 1 0 0 0 0 1	Broadband Non-Locking Shift
0 1 1 0 0 0 1 0	Broadband Sending Complete
0 1 1 0 0 0 1 1	Broadband Repeat Indicator
0 1 1 0 1 1 0 0	Calling Party Number
0 1 1 0 1 1 0 1	Calling Party Sub-Address
0 1 0 1 1 0 0 0	ATM Adaption Layer Parameters
0 1 0 1 1 0 0 1	ATM Traffic Descriptor
0 1 0 1 1 0 1 0	Connection Identifier
0 1 0 1 1 0 1 1	OAM Traffic Descriptor
0 1 0 1 1 1 0 0	Quality of Service Parameter
0 1 0 1 1 1 1 0	Broadband Bearer Capability
0 1 0 1 1 1 1 1	Broadband Low Layer Information (B-LLI)
0 1 0 1 1 1 0 1	Broadband High Layer Information (B-HLI)
0 1 0 0 0 0 1 0	End-to-End-Transit Delay
0 0 1 0 0 1 1 1	Notification Indicator
0 0 0 1 0 1 0 0	Call State
0 0 0 1 1 1 1 0	Progress Indicator
0 0 0 0 0 1 0 0	Narrowband Bearer Capability
0 0 0 0 1 0 0 0	Cause

Figure 10.47a Format and coding of information elements

Information Field Coding: Byte 2 Compatibility Indicator	
Extension indicator, bit 8	Set to 1; reserved for future use
Coding standard, bits 7 6	
0 0	ITU-T coding
0 1	ISO/IEC standard
1 0	National standard
1 1	Network-specific standard (private, public)
Flag bit 5	
0	Ignore instruction indicator
1	Obey instruction indicator
Reserved field, bit 4	Set to 0; reserved for future use
Action indicator, bits 3 2 1	
0 0 0	Clear Call
0 0 1	Discard info element and proceed
0 1 0	Discard info element, proceed and report status
1 0 1	Discard info element and ignore
1 1 0	Discard info element and report status

Figure 10.47b Format and coding of information elements

the setup request to the receiving station. If the station called is able to receive the request, the network sends the caller an ALERTING message. If the station called accepts the call, the network sends the caller a CONNECT message. The caller may optionally respond with a CONNECT ACKNOWLEDGE. This completes the connection setup, and the call is then in the active state.

10.1.9.3 Connection Setup at the Station Called (UNI)

After the receiving station has been notified of an incoming call by a SETUP message from the network, it performs a compatibility check. This test compares the address information, the QoS parameters, and the traffic parameters of the SETUP request with the local services available. If the SETUP parameters requested are not locally available, the receiving station responds with a RELEASE-COMPLETE message and Cause 88, “Incompatible Destination”. If the SETUP parameters are compatible, the receiving station can respond with a CALL PROCEEDING, ALERTING, or CONNECT message, depending on the type of end system. The network confirms a CONNECT message with a CONNECT ACKNOWLEDGE and the connection enters the active state.

Numbers

The SETUP message identifies the station called by its number in the information elements “Called Party Subaddress” and “Called Party Number”. The net-

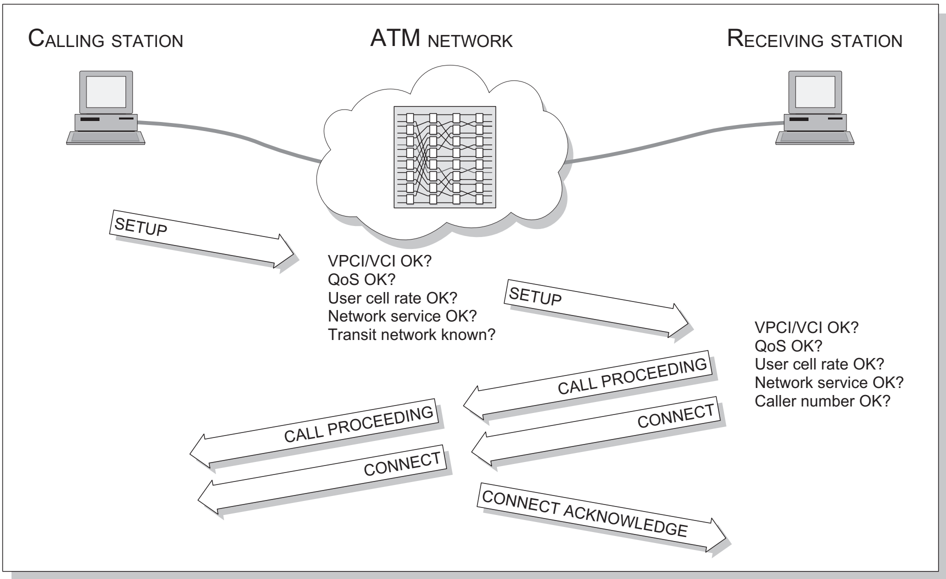

Figure 10.48 UNI connection setup

work verifies that this is a valid number. If not, the network sends a RELEASE COMPLETE with one of the following messages:

1 Unassigned (unallocated) number
2 No route to destination
3 Number changed
28 Invalid number format (incomplete number)

ATM Addressing

Six types of ATM addresses are specified in Q.2931:

- Unknown
- International number
- National number
- Network-specific number
- Subscriber number
- Abbreviated number

The format of the address types "abbreviated number" and "network-specific number" is different from case to case, depending on the network. "Network-specific numbers" can be used for administrative services (such as operator numbers). "Abbreviated number" addresses contain a shortened form of a complete ATM address, and may be defined by the network operator according to the internal structure of the network. All other address types use one of the following address formats:

- Unknown
- ISDN Telephony Numbering Plan (E.164)
- ISO NSAP Addressing (ISO 8348/AD2)
- Private Numbering Plan

The address format “unknown” is used when the user or network does not know any address format. The use of the ISO NSAP address format is optional and takes the ATM address type “unknown”.

10.1.9.4 Connection Clear-Down

Before describing the processes involved in clearing down a connection, we must define the three states “Connected”, “Disconnected”, and “Released”. An ATM virtual channel is in the “Connected” state when it forms part of an active B-ISDN connection. The virtual channel is in the “Disconnected” state when it is not part of such a connection, but is not available for any other connection. When an ATM virtual channel is in the “Released” state, then it is not part of an existing B-ISDN connection and is available for use.

Connection clear-down is normally initiated by a “RELEASE” command, except in the following three cases:

a) If a SETUP message is received without a mandatory information element, the receiver may refuse the connection setup by sending a RELEASE COMPLETE and Cause 96, “Mandatory information element is missing”.
b) If a SETUP message is received with an incorrect mandatory information element, the receiver may refuse the connection setup by sending a RELEASE COMPLETE and message #100, “Invalid information element contents”.
c) After VPI/VCI negotiation has failed, the system that initiated the call may terminate the connection setup with a RELEASE COMPLETE.

Connection Clear-Down by the User

In all cases except the three mentioned previously, sending a RELEASE command begins the connection clear-down. If the user initiates the connection clear-down, the user’s timer T308 is started when the RELEASE message is sent. The first time T308 expires, the user repeats the RELEASE command and starts T308 again. If no RELEASE COMPLETE is received from the network in answer to the second connection clear-down attempt, the user considers the virtual channel concerned as out of order and enters the null state. If the user receives a RELEASE COMPLETE from the network before T308 expires, it clears down the virtual channel, releases the call reference, and enters the null state.

Connection Clear-Down by the Network

The network can initiate a connection clear-down by sending a RELEASE command and starting its timer T308. The user responds with a RELEASE COMPLETE, whereupon the network stops T308, clears down the virtual channels and releases the call reference. The connection clear-down is then complete. If timer T308 expires before a RELEASE-COMPLETE is received from the user, the RELEASE command is repeated and T308 restarted. If T308 expires a second time before a RELEASE COMPLETE is received, the network marks the ATM virtual channel as out of order and clears the call reference.

10.1.9.5 ATM Forum UNI Signaling: UNI 3.0, 3.1, 4.0

In order to allow for manufacturers to develop ATM components before the definitive adoption of international standards by the ITU-T, the ATM Forum also developed specifications for ATM signaling in its UNI 3.0 document. The signaling specifications in UNI 3.0 were published a year before ITU-T Recommendation Q.2931, and are largely incompatible with it because of the use of a different SSCOP layer, as mentioned earlier. UNI 3.1, however, drafted after work on Q.2931 had been completed, corrected this by specifying the ITU-T Q.2110 SSCOP (the higher layer signaling messages were largely unchanged) and represented a significant step toward harmonization. Next, the ITU-T introduced Q.2971, which added route-initiated join point to multi-point signaling. Finally, the ATM forum introduced UNI 4.0 which, in most respects, is a superset of both Q.2931 and Q.2971 but additionally contains specifications for leaf-initially join point to multi-point signaling; the only signaling messages not contained in UNI 4.0 from the ITU-T recommendations concern interworking with narrow-band ISDN services. The following sections are more detailed but are limited to a description of the major differences that remain between the signaling mechanism described in UNI 3.0, UNI 3.1 and UNI 4.0 on one hand and in Q.2931 and Q.2971 on the other.

Signaling: ATM Forum UNI 3.1 Versus ITU-T Recommendation Q.2931

The AAL for signaling (SAAL) specified in UNI 3.1 is based only on AAL5. The AAL sublayer definitions for this SAAL, CP-AAL and SSCS with the Service-Specific Coordination Function (SSCF) and the Service-Specific Connection-Oriented Peer-to-Peer Protocol (SSCOP), are identical with those in the corresponding ITU-T Recommendation. While the use of VPCIs is supported for identification of the ATM virtual path used for data transmission, these VPCIs are limited to a length of 8 bits (as opposed to 16 bits in Q.2931) in order to be identical with the VPI. Furthermore, no negotiation is possible between user and network regarding the VPCI/VCI values to be used. UNI 3.1 addressing also differs from the Q.2931 specification in that it uses only two number types

Timer Number	Default Value	Start	Stop	On First Expiration	On Second Expiration	Implementation
T301 (Q.2931 only)	≥ 3 min	ALERT received	CONNECT received	Clear call	—	M (if symmetrical connections supported)
T302 (Q.2931 only)	10 – 15 s	SETUP ACK sent	SENDING COMPLETE indication	if information in-complete, clear call Otherwise, CALL PROCEEDING	—	M (if overlap sending and receiving supported)
T303	4 s	SETUP sent	CONNECT, CALL PROCEEDING, ALERT RELEASE COMPLETE received	Repeat SETUP restart T303	Abort call set-up	M
T304 (Q.2931 only)	30 s	SETUP ACK received; restart when INFO received	CALL PROCEEDING; ALERT; CONNECT received	Clear call	—	M
T308	30 s	RELEASE sent	RELEASE COMPLETE or RELEASE received	Repeat RELEASE, restart T308	—	
T309	10 s	SAAL aborted	SAAL active again	Clear call; delete VCIs and call references	—	M
T310	30 – 120 s	CALL PROCEEDING received	ALERT, CONNECT or RELEASE received	send RELEASE	—	M
T313	4 s	CONNECT sent	CONNECT ACK received	send RELEASE	—	M
T316	2 min	RESTART sent	RESTART ACK received	Repeat RESTART several times	Repeat RESTART several times	M
T317	Implemen-tation specific, but < T316	RESTART received	All call references deleted	Report error	—	M
T322	4 s	STATUS ENQUIRY sent	STATUS, RELEASE or RELEASE COMPLETE received	Repeat STATUS ENQUIRY	Repeat STATUS ENQUIRY	M
T398 (ATM Forum only)	4 s	DROP PARTY sent	DROP PARTY ACK or RELEASE received	Send DROP PARTY ACK or RELEASE COMPLETE	Timer is not restarted	M
T399 (ATM Forum only)	14 s (UNI 3.0/3.1) 34 – 124 s (UNI 4.0)	ADD PARTY sent	ADD PARTY ACK, ADD PARTY REJECT, or RELEASE received	Delete party	Timer is not restarted	M
T331 (ATM Forum 4.0 only)	60 s	LEAF SETUP REQUEST sent	SETUP, ADD PARTY, or LEAF SETUP FAILURE received	Repeat LEAF SETUP REQUEST and restart T331	Delete connection	M

Figure 10.49 Timers for UNI signaling processes: Q.2931 and ATM Forum (user system)

Timer Number	Default Value	Start	Stop	On First Expiration	On Second Expiration	Implementation
T301 (Q.2931 only)	≥ 3 min	ALERT received	CONNECT received	Clear call	—	M (if symmetrical connections supported)
T302 (Q.2931 only)	10 – 15 s	SETUP ACK sent; restart when INFO sent	SENDING COMPLETE indication	if information in-complete, clear call Otherwise, CALL PROCEEDING	—	M (if overlap sending and receiving supported)
T303	4 s	SETUP sent	CONNECT, CALL PROCEEDING, ALERT SETUP ACK, RELEASE COMPLETE received	Repeat SETUP restart T 303	Abort call set-up	M
T304 (Q.2931 only)	20 s	SETUP ACK received	INFO sent or CALL PROCEEDING, ALERT, CONNECT received	Clear call	—	
T306 (Q.2931 only)	20 s	RELEASE with Progress Indicator 8 sent	RELEASE COMPLETE received	Stop ringing	—	M (if inband alarms supported)
T308	30 s	RELEASE sent	RELEASE COMPLETE or RELEASE received	Repeat RELEASE, restart T308	Set VC to maintenance state	M
T309	10 s	SAAL aborted	SAAL active again	Clear call; delete VCIs and call references	—	O
T310	10 s	CALL PROCEEDING received	ALERT, CONNECT or RELEASE received	Clear call	—	M
T316	2 min	RESTART sent	RESTART ACK received	Repeat RESTART several times	Repeat RESTART several times	M
T317	Implemen-tation specific, but < T316	RESTART received	All call references deleted	Report error	—	M
T322	4 s	STATUS ENQUIRY sent	STATUS, RELEASE or RELEASE COMPLETE received	Repeat STATUS ENQUIRY	Repeat STATUS ENQUIRY	M
T398 (ATM Forum only)	4 s	DROP PARTY sent	DROP PARTY ACK or RELEASE received	Send DROP PARTY ACK or RELEASE	Timer is not restarted	M
T399 (ATM Forum only)	14 s	ADD PARTY sent	ADD PARTY ACK, ADD PARTY REJECT, or RELEASE received	Delete party	Timer is not restarted	M

Figure 10.50 Timers for UNI signaling processes: Q.2931 and ATM Forum (network system)

("Unknown" and "International number") and the numbering plans E.164 (the ISDN numbering system) and ISO NSAP. When E.164 numbers are used, the number type is "International number". ISO NSAP numbers are designated as "Unknown".

UNI 3.1 does not support the following Q.2931 signaling messages:

- ALERTING
- PROGRESS
- SETUP ACKNOWLEDGE
- INFORMATION
- NOTIFY

Due to the lack of support for the NOTIFY message, the information element "Notification Indicator" is also not supported. This information element is used in the Q.2931 NOTIFY message to obtain information about the connection state, or more specifically the messages CALL PROCEEDING, CONNECT ACKNOWLEDGE, RELEASE, and SETUP. For point-to-multipoint signaling, UNI 3.1 provides following message types not defined in Q.2931:

- ADD PARTY
- ADD PARTY ACKNOWLEDGE
- ADD PARTY REJECT
- DROP PARTY ACKNOWLEDGE

UNI 3.1 supports only the standard code set 0, while Q.2931 also supports code sets 5, 6, and 7. Furthermore, the definitions of certain UNI 3.1 information elements are slightly modified. The maximum length of the "Traffic Descriptor" information element in the SETUP message, for example, is increased to 30 bytes (as opposed to 20 in Q.2931); the length of the "Transit Network Selection" element is limited to 8 bytes. A complete list of the differences between UNI 3.1 and Q.2931 signaling is found in Appendix E of the ATM Forum UNI specification.

The timers defined in UNI 3.1 for B-ISDN signaling are the same as those in Q.2931 with the exception of T301, T302, and T304, which are not supported by UNI 3.1. The ATM Forum defines additional timers for point-to-multipoint processes: T398 and T399 (UNI 3.1), and T331 (UNI 4.0). Figures 10.49 and 10.50 list the UNI protocol timers.

ATM Forum Signaling: UNI 3.0

The most important difference between the signaling specifications in UNI 3.0 and 3.1 is that the adaptation layer for signaling (SAAL) in UNI 3.1 is compatible with the corresponding definition in Q.2931. This is not true of UNI 3.0, which uses the older SAAL specifications Q.SAAL1 and Q.SAAL2. Furthermore, UNI 3.0 signaling supports neither meta-signaling nor broadcast signaling.

ATM Forum Signaling: UNI 4.0

The current ATM Forum signaling specification UNI 4.0 closely approximates ITU-T Recommendation Q.2931 with Q.2971 (point-to-multipoint signaling extensions). With the exception of the messages SETUP ACKNOWLEDGE and INFORMATION, which in any case are only required in the case of transitions between narrow-band and broadband ISDN networks, UNI 4.0 supports all the Q.2931 and Q.2971 signaling messages. In addition to the messages specific to ATM Forum UNI 3.1, 4.0 also provides the messages PARTY ALERTING and LEAF SETUP REQUEST, as well as the corresponding protocol states. It is worth noting that UNI 4.0 was based primarily upon Q.2931 and Q.2971, rather than UNI 3.1. Furthermore, as in much standardization work, many of the same signaling experts were involved in the development of both ITU-T recommendations and ATM Forum specifications. Consequently, with the exception of the SETUP ACKNOWLEDGE and INFORMATION messages previously mentioned, UNI 4.0 can be considered to be a superset of ITU-T UNI recommendations; equipment conforming to the ATM Forum's UNI 4.0 signaling protocols generally are compatible with equipment based on ITU-T recommendations.

10.1.9.6 NNI Signaling

For connection setup between two NNIs, the ITU-T developed the broadband protocol B-ISUP (ITU-T Recommendations Q.2761-Q.2764), modeled after the narrow-band ISDN protocol ISUP. B-ISUP contains many of the mechanisms and functions of the proven ISUP. Only the characteristic parameters and processes of broadband networks with their virtual connections are added. By the same token, ISUP functions that are only practical in conventional connection-oriented networks are omitted. While B-ISUP is used for NNI signaling in wide-area networks, NNIs in private local ATM networks use the ATM Forum PNNI protocol. Unlike B-ISUP, PNNI is able to select routes for the desired connections, a function that has to be implemented completely independently by manufacturers of B-ISUP switching systems. In local ATM networks, the ATM network can be made operational immediately based on PNNI, whereas large, public ATM WAN networks with B-ISUP signaling first require the definition and implementation of route selection and bandwidth monitoring functions. Although designed for private networks, PNNI is also being used in some large public networks, particularly in the United States.

B-ISUP Signaling

The B-ISUP protocol is the NNI counterpart to the UNI protocol in ITU-T Recommendation Q.2931 (or ATM Forum UNI 4.0) and extends virtual UNI connections across the ATM network to the end system called. There the NNI connection is translated back into a UNI connection. However, B-ISUP is unable to fulfill functions such as:

- Bandwidth management
- Route selection
- Routing table maintenance
- OAM process control for existing virtual connections

The implementation of these functions is left to the component manufacturers or the operators of B-ISDN networks. The B-ISUP protocol builds on Message Transfer Protocol Version 3 (MTP-3) specified especially for ATM, whose data packets are transported in turn by means of the NNI SAAL layer.

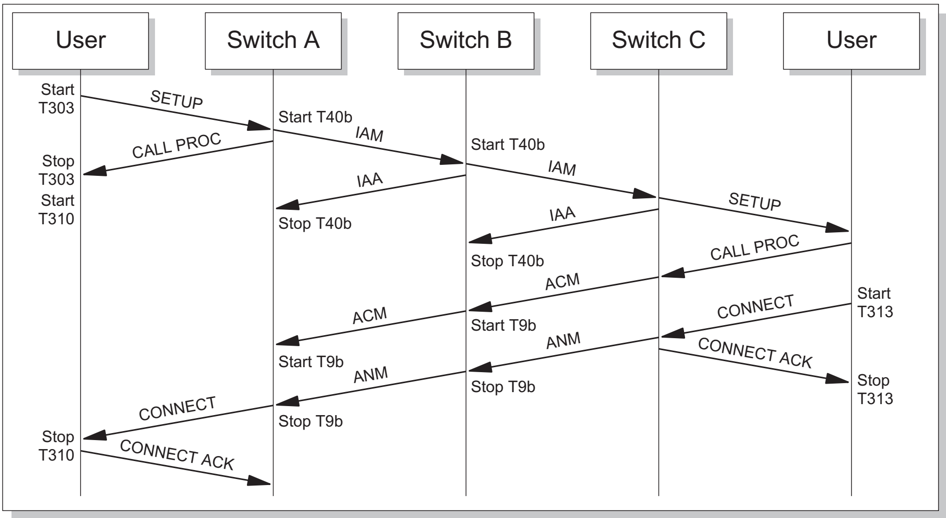

Figure 10.51 Connection setup and tear-down in the UNI/NNI network

PNNI Signaling

PNNI (Private Network to Network Interface) signaling consists of two protocols: a topology protocol, which distributes information about the network topology to the individual network stations, and a signaling protocol, which is necessary for connection setup between PNNI nodes. The signaling protocol is mainly based on the ITU-T Q.2931 signaling recommendation. In addition to the signaling mechanisms drawn from Q.2931, however, PNNI also contains a number of completely new functions. These include searching for alternative routes, which can involve "crankback" (back-tracking) when an attempt to move forward fails for various reasons, maintenance of connection matrices for ascertaining routes, and the operation of a dedicated routing control virtual channel, used only to distribute routing information. Furthermore, all PNNI nodes save updated topology state information at regular intervals. These "topology state parameters" provide information about the availability of connections to adjacent nodes, and are continuously updated by means of a special "hello" procedure carried out between all PNNI nodes.

10.1.10 ATM Interworking

10.1.10.1 LAN-ATM: LLC Encapsulation (RFC 1483)

The first specifications available for connecting LANs with ATM networks were the two Requests for Comments (RFCs) RFC 1483 (superceded in 1999 by RFC 2684) and RFC 1577 (superceded in 1998 by RFC 2225), published in 1993 by the IETF (Internet Engineering Task Force). RFC 2684 specifies two methods for the encapsulation of LAN data packets in ATM:

- LLC encapsulation
- VC-based multiplexing

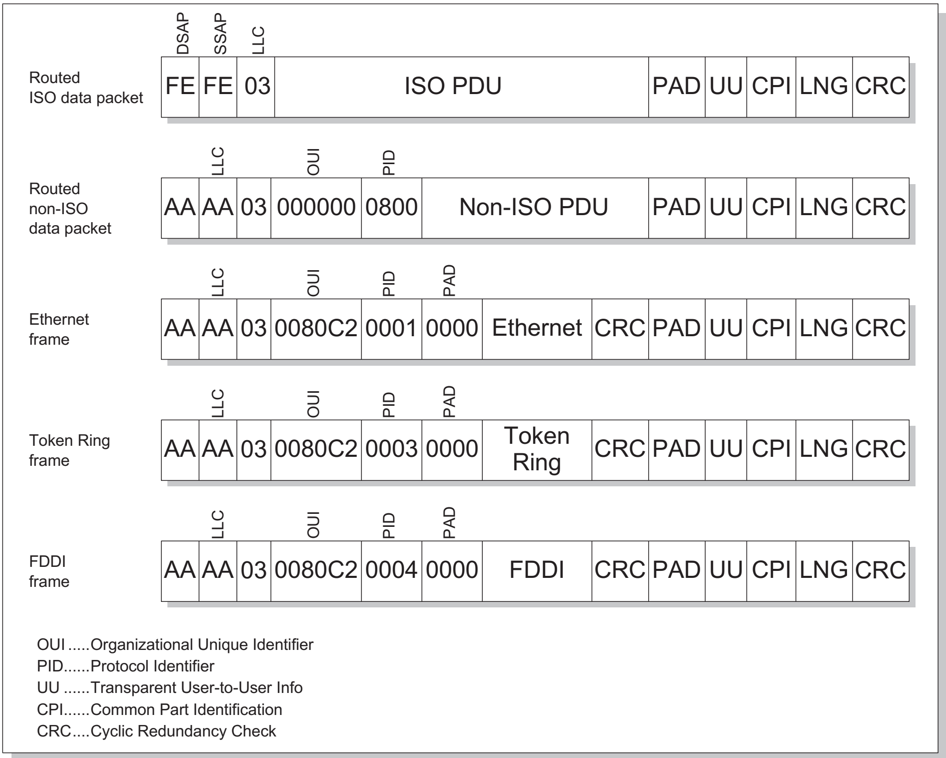

Figure 10.52 ATM-LAN: LLC encapsulation (RFC 2684)

The first method transports all LAN protocols by encapsulating the Logical Link Control (LLC) data packets in AAL5 CPCS PDUs. The entire data stream is transmitted within one VC. All protocols based on Ethernet, Token Ring, FDDI, or Distributed Queue Dual Bus (DQDB) (IEEE 802.6 MAN) can be transported over ATM networks in this way (see Figure 10.52).

LLC encapsulation is used mainly in networks that only support PVCs, and which are unable to manage constantly changing VCs in use. In networks using SVCs, which typically have no trouble dynamically managing a large number of VCs, VC-based multiplexing can be used. This method avoids transporting the LLC header by setting up a separate virtual channel for each protocol. This makes transportation significantly more efficient overall, which is why this method is preferable to LLC encapsulation wherever possible (Figure 10.53).

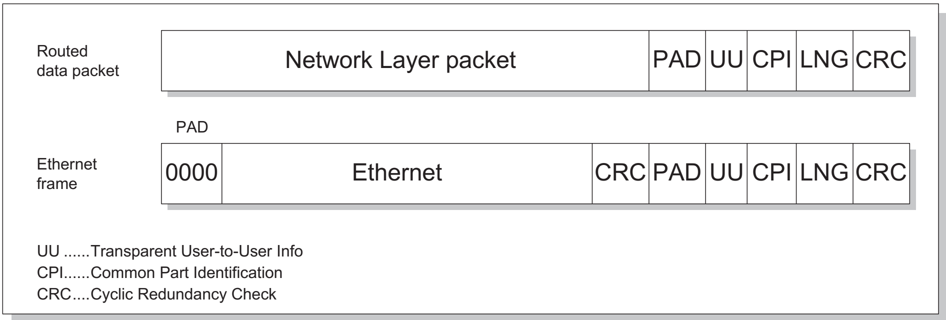

Figure 10.53 VC-based multiplexing (RFC 2684)

10.1.10.2 LAN-ATM: Classical IP and ARP over ATM (RFC 2225)

Classical IP and Address Resolution Protocol (ARP) over ATM go beyond the RFC 2684encapsulation technique to provide a complete implementation of the Internet protocol for ATM. IP address resolution, which is realized in Ethernet by ARP and Reverse Address Resolution Protocol (RARP), is handled by ATMARP and InATMARP functions. The mapping tables for the ATMARP and InATMARP functions are stored in an ATMARP server, which must be present in each logical IP subnet (LIS). The ARP client itself is responsible for registering its own IP/ATM address information with the ATMARP server, and for obtaining the IP/ATM address of the desired destination system from the ARP server. Entries in clients' and servers' ATMARP are subject to an aging process. Client ATMARP entries are valid for a maximum of 15 minutes, server ATMARP entries for at least 20 minutes. ATMARP PDUs, like the IP data packets themselves, are transported in AAL5 CPCS PDUs according to the rules of LLC encapsulation (see Figure 10.54).

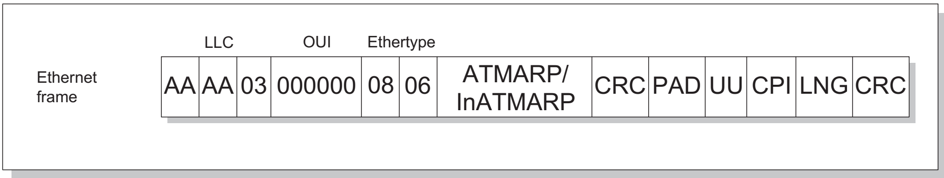

Figure 10.54 ATMARP and InATMARP (RFC 2225)

The following rules govern the design and operation of Classical IP subnetworks (LISs):

- All IP nodes of an LIS must be directly connected to the ATM network.
- All IP nodes of an LIS must have the same IP network or subnet work address.
- Network nodes outside the LIS must be accessible only through routers.
- All IP nodes in an LIS must be able to resolve addresses by ATMARP.
- Every IP node in an LIS must be able to communicate with all other IP nodes.
- Address resolution must function both for PVCs and for SVCs.

The default packet size in CIPs is 9,180 bytes. Adding the 8-byte LLC/SNAP header yields a default AAL5 PDU size of 9,188 bytes in CIP networks. The differences between RFC 2225 and its predecessor RFC 1577 lie mainly in the in the progress made in signaling procedures during the intervening 5 years.

10.1.10.3 LAN Emulation (LANE)

The most universal method of efficiently integrating ATM networks in existing, conventional LAN structures involves a complete emulation of the LAN MAC layer, which allows all existing LAN applications to be extended across ATM networks without modification. From the point of view of traditional local-area networking, the LAN emulation service behaves the same as a conventional MAC LAN driver. The ATM Forum's LAN Emulation (Version 2) is defined in the two documents LANE-LUNI and LANE-LNNI, and is based on five ATM service modules that build on AAL5:

- LAN Emulation Client (LE Client or LEC)
- LAN Emulation Server (LE Server or LES)
- LAN Emulation Configuration Server (LE Configuration Server or LECS)
- Broadcast and Unknown Server (BUS)
- Selected Multicast Server (SMS)

LUNI describes the processes at the interface to the LE Client, while LNNI specifies the protocols for networking LE Client, LE Server, BUS and SMS components. The use of LNNI components permits the coordination of several LE Servers, LE Configuration Servers and BUS systems so that redundant LANE structures can be built. LE Clients can also be grouped in subnetworks regardless of their location or the LE Server or BUS used, which makes it possible to build virtual Emulated LANs (ELANs).

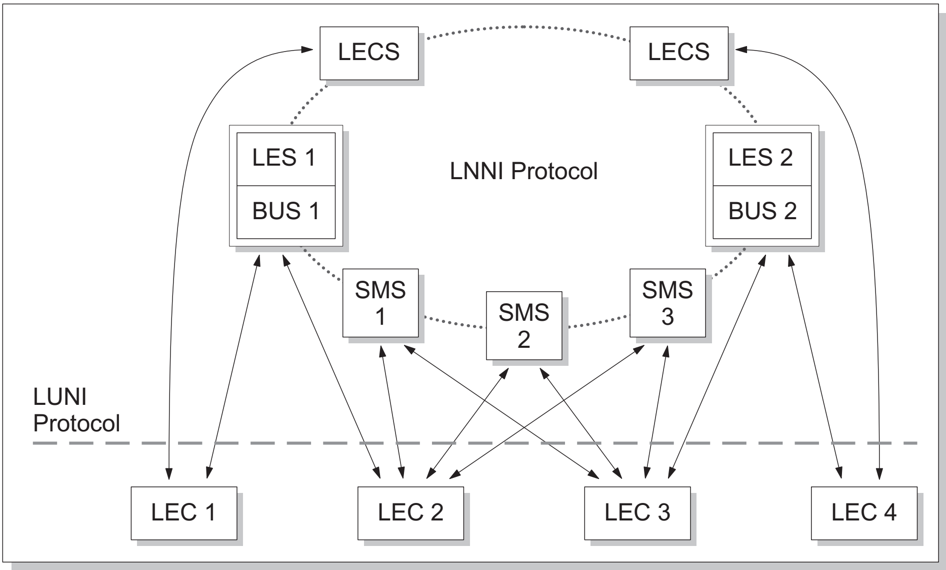

Figure 10.55 LAN emulation: LUNI and LNNI

With regard to the protocol stacks of end systems, LANE emulates the most widely used protocol driver specifications: Network Driver Interface Specification (NDIS), Open Data Link Interface (ODI) and Data Link Provider Interface (DLPI). The LAN data packets themselves (IEEE 802.3 Ethernet or IEEE 803.5 Token Ring) are transported over AAL5 in “LAN emulation frames”. No special LANE frame type is defined for FDDI: either Ethernet or Token-Ring frames can be used. Using the Token-Ring frame format yields better results because the MAC address format is the same as in FDDI. LANE also provides quality-of-service functions for communication between LANE stations connected by ATM.

The LE Client software performs all the necessary control functions as well as the actual data communication over the ATM interface, providing a standard LAN MAC interface (NDIS, ODI or DLPI) to higher-layer applications. LE Clients consist of the following components:

- System hardware (PC, workstation, router, etc.)
- Standard LAN software (MAC address, protocol stack, drivers, etc.)
- LE Client software
- ATM interface with ATM address

The LE Server module controls the emulated LAN (ELAN). It includes functions such as registration of LE Clients and MAC-to-ATM address resolution for all registered stations. Every LE Client participating in an ELAN reports its LAN MAC address, the corresponding ATM address, and any necessary route information to the LE Server. When a LAN data packet needs to be sent, the ATM address of the destination is first sought in the address table of the LE Server. If it is not found there, address resolution must be performed by BUS broadcasts.

The LE Configuration Server manages the assignments of LE Clients to various ELANs by maintaining the configuration information of the ELANs in a configuration database. An LE Client can belong to several ELANs simultaneously.

The BUS module retransmits all the LE Clients' broadcast and multicast data packets. These include:

- All data packets with broadcast or multicast addresses
- Data packets sent to a MAC address for which the LE Client does not know the corresponding ATM address and which could not be resolved by the LE Server
- The source routing mechanism's explorer data packets used to determine optimum routes

Data packets received by the BUS are retransmitted in sequence to the appropriate group of destination LE Clients. This is necessary in order to avoid overlapping of AAL5 data packets from different senders.

Functional Processes in LAN Emulation

Process control between the individual LE Servers takes place by means of control VCCs and data VCCs. Control virtual channels are used to connect LE Clients with LE Servers and with LE Configuration Servers. Communication between LE Clients, as well as between BUSs and LE Clients, takes place over data virtual channels (data VCCs). In contrast to LANE Specification 1.0, LANE v2 also supports LCC Multiplexing. This means that the data streams of several ELANs, and of several different protocols, can be transported over a single VCC.

When a new LE Client is added to an ELAN, the LE Client first sets up a Configuration Direct VCC to the LE Configuration Server in order to register as a member of a particular ELAN. At the same time it may optionally use the LE configuration protocol to negotiate a variety of parameters (addresses, name of the ELAN, maximum frame size).

Next, the Control Direct VCC to the LE Server is set up. At this point the LE Client should be in possession of all information necessary for participation in the LAN emulation service, including the LE Client Identifier (LECID), the LAN type (802.3 Ethernet, 802.5 Token Ring), etc. If necessary, the LE Server may also set up a one-way, point-to-multipoint Control Distribute VCC to several LE Clients. (Point-to-multipoint connections are supported in LANE Version 2.0, but not in LANE 1.0). Finally, communication with the BUS is opened over a Data Direct VCC. The LE Client is then ready for operation under the LAN emulation service (see Figure 10.56). The various connections are set up and cleared down by means of the ATM Forum UNI signaling protocols (UNI 3.0, UNI 3.1, UNI 4.0).

In LAN emulation, data flows either between the LE Client and the BUS, or between one LE Client and another. If an LE Client needs to send a packet to a station whose ATM address is unknown, it first sends an LE-ARP message to the LE Server. The LE Server either forwards the LE-ARP request directly to the desired station, or provides the sender with the desired ATM address in the form of an LE-ARP response. Several data VCCs may be set up between two LE Clients, if both of them support this. The application layer may specify QoS parameters for the connections, such as ATM Traffic Descriptor, Alternative ATM Traffic Descriptor (only in UNI 4.0), Minimum Acceptable Traffic Descrip-

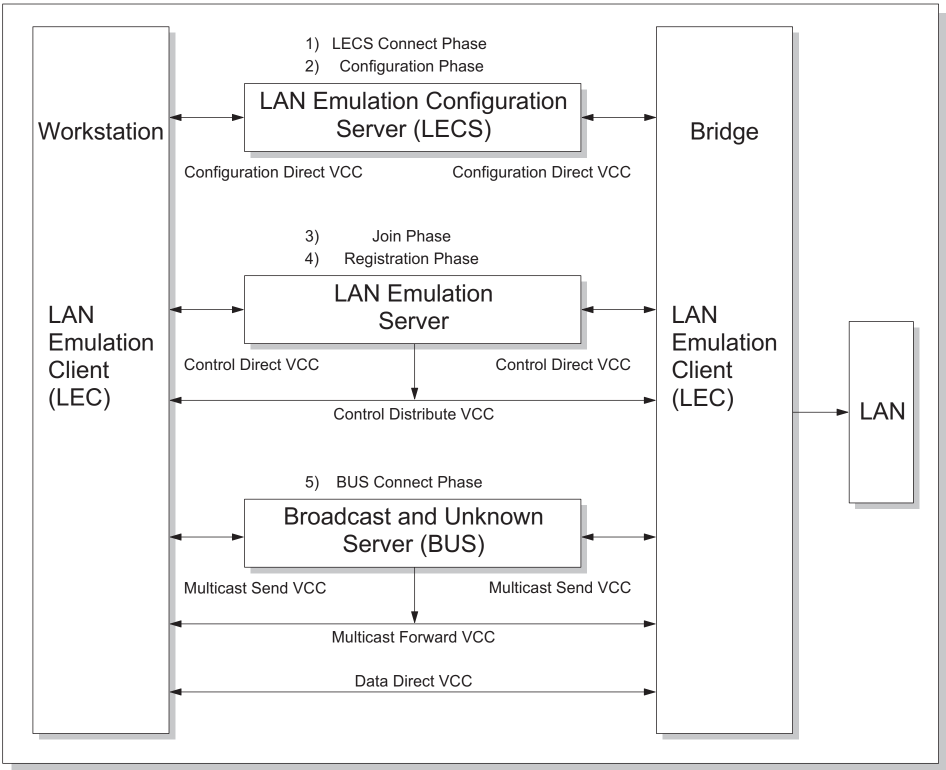

Figure 10.56 LANE functional processes

tor (only in UNI 4.0), Broadband Bearer Capability, QoS parameters, End-to-End Transmit Delay (only in UNI 4.0) etc. If no specific QoS parameters are requested, the default QoS parameters for LANE data VCCs must be used (UBR or ABR, QoS class 0).

Because LE Clients can send data packets to the same destination by different paths—via the BUS or over different data VCCs, for example—the Flush Message protocol is used to preserve the order of the LAN data packets during transmission over the ELAN. The Flush Message protocol in effect "deletes" the old transmission path when communication is moved to a new path. The LE Client sends an LE Flush message over the given VCC to ensure that all data packets sent over that virtual channel have arrived at the receiving station, and that no more packets are forthcoming in the opposite direction. On receiving the Flush response, the LE Client knows that no more data will be sent to it over the given VCC, and can begin using the new VCC.

All LANE components must support one of the following maximum sizes for AAL5 SDUs: 1,516; 1,580; 4,544; 9,234; or 18,190 bytes for non-multiplexed data packets, and 1,528; 1,592; 4,556; 9,246 or 18,202 bytes for multiplex LLC data.

The SDU length of 1,516 is due to the fact that LANE data packets contain the 2-byte LAN emulation header (LEH), but not the 4-byte checksum. For Token Ring, maximum packet sizes of 4,450 bytes (for 4 Mbit/s) and 18,200 bytes (for

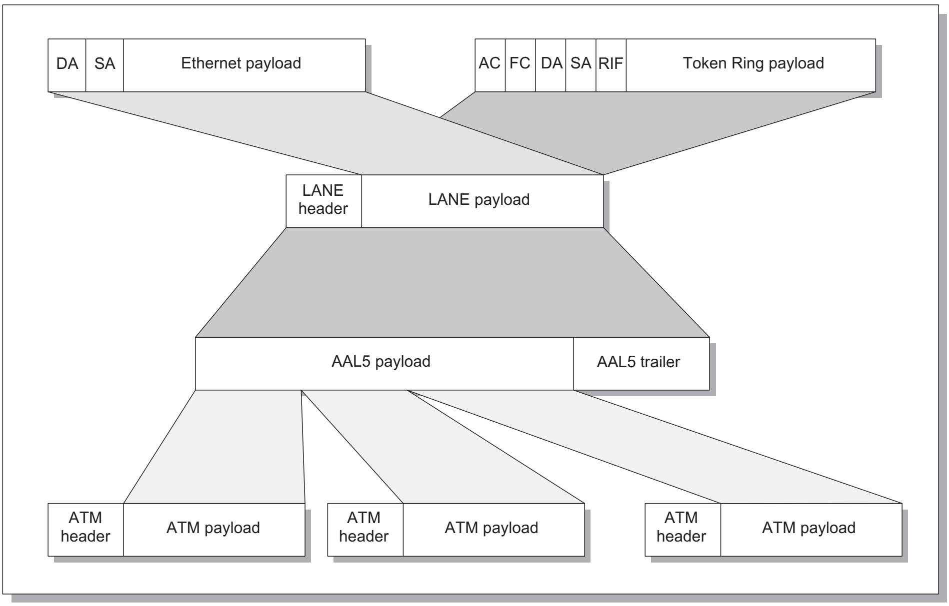

Figure 10.57 Format of LANE data packets

16 Mbit/s) are calculated based on a Token Holding Timer value of 9.1 ms. However, LANE data packets contain no SD, FCS, ED, or FS fields and no interframe spacing gap, so that the resulting lengths are 4,544 and 18,190 bytes.

10.1.10.4 Multiprotocol over ATM (MPOA)

ATM networks can use LAN Emulation to emulate Ethernet and Token Ring topologies. Such emulated ELANs can then form the basis of subnetworks (virtual LANs) that include network nodes of all three network topologies: Ethernet, Token Ring and ATM. Any ATM network can contain several ELANs, but LAN Emulation alone does not make it possible to connect ELANs to one another, except through the use of conventional routers. The MPOA specification does away with this restriction. MPOA, or Multiprotocol over ATM, integrates LANE v2 with the Next Hop Resolution Protocol (NHRP) and the Multicast Address Resolution Protocol (MARS), and permits inter-ELAN communication without the use of additional routers. By separating the route selection and data forwarding functions, MPOA is able to provide efficient routing even in large, complex networks.

Every MPOA network consists of MPOA servers (MPS) and MPOA clients (MPC). MPOA servers are responsible for route selection, while MPOA clients perform the actual forwarding of data from their LE Client interfaces, requesting the appropriate routes from the MPS as necessary.

Prerequisites for the use of MPOA in the LAN include support for:

- ATM UNI signaling (UNI 3.0, 3.1, 4.0)
- LANE v2
- Next Hop Resolution Protocol

Basic MPOA Processes

As in LANE, all MPOA processes are managed through control virtual channels, while the actual data transport takes place through separate data virtual channels. All control and data flows are transported over VCCs with LLC encapsulation (RFC 1483). Channels are set up and cleared down in accordance with one of the UNI signaling specifications, UNI 3.0, 3.1, or 4.0. Four kinds of control virtual channels are defined:

- MPS - MPC configuration virtual channels
- MPC–MPS control virtual channels
- MPS–MPS control virtual channels
- MPC–MPC control virtual channels

MPOA components obtain configuration information from the LE Configuration Server. The MPOA clients obtain route information from the MPOA server over the MPC–MPS control virtual channel. The MPS–MPS control virtual channels are used by standard routing protocols and NHRP to exchange route information. MPCs exchange control information with one another only if one MPC receives misrouted data packets from another. In this case, the sender MPC is notified so that it can delete the incorrect routing information from its cache. MPOA networks are characterized by the following five basic operating processes:

- Configuration
- Discovery
- Target resolution
- Connection management
- Data communication

MPOA obtains all its configuration information from the ELANs' LE Configuration Servers. Additionally, MPOA components can also be directly configured by means of the MPOA MIB. All MPCs and MPSs automatically detect each other's existence by using an extended LE-ARP protocol. This protocol transports not only the ATM addresses of the individual stations, but also information about the MPOA type (MPC or MPS). The target resolution process determines the route to the destination, and is carried out using modified NHRP Resolution Request messages. The connection management process is responsible for setting up and operating the virtual control and data connections, and the data transport process consists of transmitting the user data over the selected routes. The various MPOA processes are controlled by means of eight MPOA and two NHRP control messages:

- MPOA Resolution Request
- MPOA Resolution Reply
- MPOA Cache Imposition Request
- MPOA Cache Imposition Reply
- MPOA Egress Cache Purge Request
- MPOA Egress Cache Purge Reply
- MPOA Keep Alive
- MPOA Trigger
- NHRP Purge Request
- NHRP Purge Reply

Route Selection in MPOA Networks

In selecting routes for data packet transport, MPOA makes a distinction between default routes and shortcuts. Data packets initially enter the MPOA network through an MPC. Because the MPC generally does not know the ATM address that corresponds to the Layer 3 destination address of the packet, it does not attempt to set up a direct VC to the destination. By default, the data is first encapsulated in a LANE frame and forwarded to the default MPS router by the MPC's LE Client unit. This router forwards the data to the destination MPC according to the information available in its routing tables. Then the MPC attempts to resolve the network address of the packets by sending an MPOA Resolution Request message to the MPS. The MPS obtains the corresponding ATM address, along with the information whether a direct connection (shortcut) is available for the given connection, from its assigned Next Hop Server (NHS). If a shortcut is available, the MPS sends an Imposition Request message to the destination MPC to ask whether it is able to accept the shortcut connection. The destination MPC responds with an Imposition Reply message, which the MPS forwards to the originating MPC in the form of an MPOA Resolution Reply. The originating MPC can then set up the shortcut connection. The shortcut route information is entered in the MPC's routing cache, and all data packets for the given destination address are then sent directly over the shortcut. Figure 10.58 shows a schematic illustration of the MPOA route selection process. The default

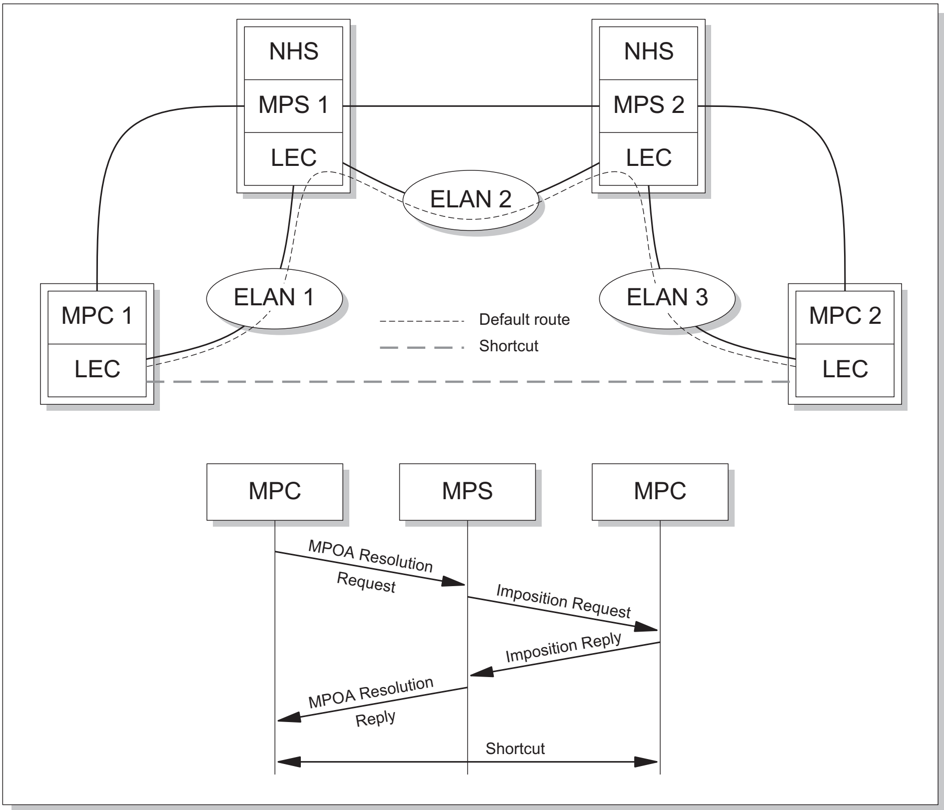

Figure 10.58 Route selection and shortcuts in MPOA networks

route between the MPOA clients MPC1 and MPC2 leads first to the LE Client interfaces of the three MPOA components connected to one another over the three ELANs. Once the shortcut has been obtained from the Next Hop Server, a direct connection can be set up between the LE Client interfaces of the two MPOA clients.

The Next Hop Resolution Protocol (NHRP)

Route optimization in MPOA networks is performed by means of the Non-Broadcast Multiple Access (NBMA) and NHRP because broadcast-based address resolution mechanisms are completely unsuitable for ATM networks. The purpose of NHRP is to find the ATM address that corresponds to a given network address (such as an IP address) so that a direct connection (shortcut) can be set up between two communicating stations. NHRP is based on a client-server model in which Next Hop Clients (NHC) send address resolution requests to a NHS. In MPOA, the MPOA servers play the part of NHCs and initiate Next Hop Resolution requests to the NHS as required by MPCs. Every NHR request contains the following information:

- Network address (for example, IP address) of the destination
- Network address of the sender
- ATM address of the sender

The NHS analyzes the incoming address resolution request and verifies whether it is itself competent to resolve the given destination. If not, it forwards the request to another NHS. Next Hop Servers may forward address resolution requests only to other NHSs that are not more than one hop away—thus the name “next hop” server. If the NHS is competent for the destination indicated in the request, it resolves the address and sends a reply packet containing the Next Hop Network address and the ATM address of the destination to the MPS, which forwards the information to the MPC. The MPC can then set up a direct VCC, or shortcut, to the destination.

10.1.11 Design Rules for ATM Networks

The design of ATM networks is subject to far fewer restrictions than classical LAN technologies such as Ethernet, FDDI or Token Ring. This is due to the wide range of possible transmission media, and to the fact that ATM was originally designed for use in wide-area networks. Distance limitations are therefore not a fundamental issue. So long as the specifications for the various segments of an ATM network are respected, distances of several thousand km can be bridged. The transmission delay can be calculated using a value of 6 µs per km and 2 µs per switch. The transmission delay for a path of 1,000 km and five switches is 6.01 ms.

The following tables list the limits for the maximum distances that can be spanned using various transmission media as defined in the appropriate ATM standards:

SDH/SONET-based ATM over twisted-pair cabling

Cable type	Transmission speed	Maximum distance
UTP-5	12.96 Mbit/s	400 m
UTP-5	25.92 Mbit/s	320 m
UTP-5	51.85 Mbit/s	160 m
UTP-5	155 Mbit/s	150 m
UTP-3	155 Mbit/s	100 m

Cell-based ATM over twisted-pair cabling

Cable type	Transmission speed	Maximum distance
UTP-3 (100W)	25 Mbit/s	100 m
UTP-4 (120W)	25 Mbit/s	100 m
STP (150W)	25 Mbit/s	100 m

SDH/SONET-based ATM over plastic fiber

Cable type	Transmission speed	Maximum distance
Plastic Optical Fiber (POF)	155 Mbit/s	50 m
Hard Polymer Clad Fiber (HPCF)	155 Mbit/s	100 m

SDH/SONET-based ATM over single-mode fiber

Fiber type	Wavelength	Transmission speed	Laser type	Maximum attenuation	Maximum distance
SM	1,310 nm	- 9.9 Gbit/s	SLM, MLM	0-7 dB	2 km
SM	1,310, 1,550 nm	- 9.9 Gbit/s	SLM, MLM	0-12 dB	15 km
SM	1,310, 1,550 nm	- 9.9 Gbit/s	SLM, MLM	10-14 dB	40 km

SDH/SONET-based ATM over multimode fiber

Fiber type	Wavelength	Transmission speed	Laser type	Maximum distance
MM	1,300 nm	155 Mbit/s	SW, LED-MMF	2 km
MM	1,300 nm	622 Mbit/s	LED-MMF	500 m
MM	1,300 nm	622 Mbit/s	SW-MMF	300 m

10.1.12 ATM Standards

All the important standards for ATM technology are developed by the ITU-T, the ATM Forum and the IETF. A selection of the most important ITU-T, ATM Forum and IETF Recommendations is listed here. A complete list of ITU-T and ATM Forum standards can be found in the Appendix or at the web sites listed below:

10.1.12.1 ATM Forum Standards

Technical Working Group	Approved Specifications	Specification	Approved Date
B-ICI	B-ICI 2.0 (integrated specification)	af-bici-0013.003	Dec, 1995
Control Signaling	PNNI Version 1.0 Security Signaling Addendum	af-cs-0116.000	May, 1999
	UNI Signaling 4.0 Security Addendum	af-cs-0117.000	May, 1999
Data Exchange Interface	Data Exchange Interface version 1.0	af-dxi-0014.000	Aug, 1993
Frame-Based ATM	Frame Based ATM over Sonet/SDH	af-fbatm-0151.000	July, 2000
LAN Emulation/ MPOA	LANE Servers Management Spec v1.0	af-lane-0057.000	Mar, 1996
	LANE v2.0 LUNI Interface	af-lane-0084.000	July, 1997
	LAN Emulation Client Management Specification Version 2.0	af-lane-0093.000	Oct, 1998
	LAN Emulation over ATM Version 2 - LNNI Specification	af-lane-0112.000	Feb. 1999
	Multi-Protocol Over ATM Specification v1.0	af-mpoa-0087.000	July, 1997
	Multi-Protocol Over ATM Version 1.0 MIB	af-mpoa-0092.000	July, 1998
	Multi-Protocol Over ATM Specification, Version 1.1	af-mpoa-0114.000	May, 1999
	Multi-Protocol Over ATM Version 1.0 MIB	af-mpoa-0092.000	July, 1998

Technical Working Group	Approved Specifications	Specification	Approved Date
Physical Layer	Utopia	af-phy-0017.000	Mar, 1994
	Mid-Range Physical Layer Specification for Category 3 UTP	af-phy-0018.000	Sep, 1994
	E3 UNI	af-phy-0034.000	Aug, 1995
	Physical Interface Specification for 25.6 Mb/s over Twisted Pair	af-phy-0040.000	Nov, 1995
	A Cell-Based Transmission Convergence Sublayer for Clear Channel Interfaces	af-phy-0043.000	Jan, 1996
	622.08 Mbps Physical Layer	af-phy-0046.000	Jan, 1996
	155.52 Mbps Physical Layer Specification for Category 3 UTP (See also UNI 3.1, af-uni-0010.002)	af-phy-0047.000	Nov, 1995
	120 Ohm Addendum to ATM PMD Interface Spec for 155 Mbps over TP	af-phy-0053.000	Jan, 1996
	DS3 Physical Layer Interface Spec	af-phy-0054.000	Mar, 1996
	155 Mbps over MMF Short Wave Length Lasers, Addendum to UNI 3.1	af-phy-0062.000	July, 1996
	WIRE (PMD to TC layers)	af-phy-0063.000	July, 1996
	E-1 Physical Layer Interface Specification	af-phy-0064.000	Sep, 1996
	155 Mbps over Plastic Optical Fiber (POF) Version 1.0	af-phy-0079.000	May, 1997
	155 Mb/s Plastic Optical Fiber and Hard Polymer Clad Fiber PMD Specification Version 1.1	af-phy-0079.001	Jan, 1999
	Inverse ATM Mux Version 1.0	af-phy-0086.000	July, 1997
	Inverse Multiplexing for ATM (IMA) Specification Version 1.1	af-phy-0086.001	Mar, 1999
	Physical Layer High Density Glass Optical Fiber Annex	af-phy-0110.000	Feb, 1999
	622 and 2488 Mbit/s Cell-Based Physical Layer	af-phy-0128.000	July, 1999
	ATM on Fractional E1/T1	af-phy-0130.000	Oct, 1999
	2.4 Gbps Physical Layer Specification	af-phy-0133.000	Oct, 1999
	Physical Layer Control	af-phy-0134.000	Oct, 1999
	Utopia 3 Physical Layer Interface	af-phy-0136.000	Nov, 1999
	Multiplexed Status Mode (MSM3)	af-phy-0142.000	Mar, 2000

Technical Working Group	Approved Specifications	Specification	Approved Date
	Frame-Based ATM Interface (Level 3)	af-phy-0143.000	Mar, 2000
	UTOPIA Level 4	af-phy-0144.001	Mar, 2000
PNNI	PNNI ABR Addendum	af-pnni-0075.000	Jan, 1997
	PNNI v1.0 Errata and PICs	af-pnni-0081.000	July, 1997
	PNNI 1.0 Addendum (soft PVC MIB)	af-pnni-0066.000	Sep, 1996
Service Aspects	Audio/Visual Multimedia Services:	af-saa-0049.000	Jan, 1996
and Applications	Video on Demand v1.0		
	Audio/Visual Multimedia Services: Video on Demand v1.1	af-saa-0049.001	Mar, 1997
	ATM Names Service	af-saa-0069.000	Nov, 1996
	FUNI 2.0	af-saa-0088.000	July, 1997
	Native ATM Services DLPI Addendum Version 1.0	af-saa-api-dlpi-0091.000	Feb, 1998
	H.323 Media Transport over ATM	af-saa-0124.000	July, 1999
Traffic Management	Traffic Management 4.1	af-tm-0121.000	Mar, 1999
Voice & Telephony over ATM	(DBCES) Dynamic Bandwidth Utilization in 64 KBPS Time Slot Trunking Over ATM - Using CES	af-vtoa-0085.000	July, 1997
	ATM Trunking Using AAL1 for Narrow Band Services v1.0	af-vtoa-0089.000	July, 1997
	ATM Trunking Using AAL2 for Narrowband Services	af-vtoa-0113.000	Feb, 1999
	Low Speed Circuit Emulation Service	af-vtoa-0119.000	May, 1999
	ICS for ATM Trunking Using AAL2 for Narrowband Services	af-vtoa-0120.000	May, 1999
	Low Speed Circuit Emulation Service (LSCES) Implementation Conformance Statement Performance	af-vtoa-0132.000	Oct, 1999

10.1.12.2 ITU-T ATM Standards

I.113	Vocabulary of terms for broadband aspects of ISDN
I.121	Broadband Aspects of ISDN
I.150	B-ISDN ATM Functional Characteristics
I.211	B-ISDN: Service Aspects
I.311	B-ISDN: General Network Aspects
I.321	B-ISDN Protocol Reference Model and Its Application
I.327	B-ISDN Functional Architecture
I.361	B-ISDN ATM Layer Specification
I.363.1	B-ISDN ATM Adaptation Layer Specification: Type 1 AAL
I.363.2	B-ISDN ATM Adaptation Layer Specification: Type 2 AAL
I.363.3	B-ISDN ATM Adaptation Layer Specification: Type 3/4 AAL
I.363.5	B-ISDN ATM Adaptation Layer Specification: Type 5 AALI.432.1 B-ISDN User-Network Interface – Physical Layer (General) Specification
I.610	B-ISDN Operation and Maintenance Principles and Functions
I.356	B-ISDN ATM Layer Cell Transfer Performance (formerly I.35B)
I.371	Traffic Control and Congestion Control in B-ISDN
I.350	General Aspects of Quality of Service and Network Performance in Digital Networks, Including ISDNs
I.555	Frame Relaying Bearer Service Interworking
I.365.1	B-ISDN ATM Adaptation Layer Sublayers: Frame Relaying Service Specific Convergence Sublayer (FR-SSCS)
I.370	Congestion Management for the ISDN Frame Relaying Bearer Service
I.372	Frame Relaying Bearer Service Network-to-Network Interface Requirements
I.233.1	ISDN Frame Mode Bearer Services: ISDN Frame Relaying Bearer Service
F.811	Broadband Connection-Oriented Bearer Service

F.812	Broadband Connectionless Data Bearer Service
Q.2010	Broadband Integrated Services Digital Network Overview – Signaling Capability Set 1, Release 1
Q.2100	B-ISDN Signaling ATM Adaptation Layer (SAAL) Overview Description
Q.2110	B-ISDN ATM Adaptation Layer – Service Specific Connection Oriented Protocol (SSCOP)
Q.2120	B-ISDN Meta-Signaling Protocol
Q.2130	B-ISDN Signaling ATM Adaptation Layer – Service Specific Coordination Function for Support of Signaling at the User Network Interface (SSFC at UNI)
Q.2140	B-ISDN ATM Adaptation Layer – Service Specific Coordination Function for Signaling at the Network Node Interface (SSCF at NNI)
Q.2610	Broadband integrated services digital network (B-ISDN) – Usage of Cause and Location in B-ISDN User Part and DSS 2
Q.2650	Broadband-ISDN, Interworking Between Signaling System No. 7 – Broadband ISDN User Part (B-ISUP) and Digital Subscriber Signaling System No. 2 (DSS 2)
Q.2660	Broadband integrated services digital network (B-ISDN) – Interworking Between Signaling System No. 7 – Broadband ISDN User Part (B-ISUP) and Narrow-band ISDN User Part (N-ISUP)
Q.2730	Broadband Integrated Services Digital Network (B-ISDN) Signaling System No. 7 B-ISDN User Part (B-ISUP) – Supplementary Services
Q.2761	Broadband Integrated Services Digital Network (B-ISDN) – Functional Description of the B-ISDN User Part (B-ISUP) of Signaling System No. 7
Q.2762	Broadband Integrated Services Digital Network (B-ISDN) – General Functions of Messages and Signals of the B-ISDN User Part (B-ISUP) of Signaling System No. 7
Q.2763	Broadband Integrated Services Digital Network (B-ISDN) – Signaling System No. 7 B-ISDN User Part (B-ISUP) – Formats and Codes
Q.2764	Broadband Integrated Services Digital Network (B-ISDN) – Signaling System No. 7 B-ISDN User Part (B-ISUP) – Basic Call Procedures

Q.2931 Broadband Integrated Services Digital Network (B-ISDN) – Digital Subscriber Signaling System No. 2 (DSS 2) – User-Network Interface (UNI) - Layer 3 Specification for Basic Call/Connection Control

Q.2951.1 Stage 3 Description for Number Identification Supplementary Services Using B-ISDN Digital Subscriber Signaling System No. 2 (DSS 2) – Basic Call [Clauses 1–6 and 8; DDI, MSN, CLIP, CLIR, COLP, COLR, SUB]

Q.2957.1 Stage 3 Description for Additional Information Transfer Supplementary Services Using B-ISDN Digital Subscriber Signaling System No. 2 (DSS 2) – Basic Call: Clause 1 – User-to-User Signaling (UUS)

Q.2971 Broadband Integrated Services Digital Network B-ISDN - DSS 2 - Digital Subscriber Signaling System No. 2 - User-Network Interface Layer 3 Specification for Point-to-Multipoint Call/Connection Control

10.1.12.3 IETF ATM Standards

RFC2379 RSVP over ATM Implementation Guidelines, L. Berger, August, 1998

RFC2225 Classical IP and ARP over ATM, M. Laubach, J. Halpern, April, 1998 (Obsoletes RFC1626, RFC1577)

RFC2226 IP Broadcast over ATM Networks T. Smith, G. Armitage, October, 1997

RFC2320 Definitions of Managed Objects for Classical IP and ARP Over ATM Using SMIv2 (IPOA-MIB), M. Greene, J. Luciani, K. White, T. Kuo, April, 1998

RFC2331 ATM Signaling Support for IP over ATM - UNI Signaling 4.0 Update, M. Maher, April, 1998

RFC2380 RSVP over ATM Implementation Requirements, L. Berger, August, 1998

RFC2381 Interoperation of Controlled-Load Service and Guaranteed Service with ATM, M. Garrett, M. Borden, August, 1998

RFC2382 A Framework for Integrated Services and RSVP over ATM, E. Crawley, L. Berger, S. Berson, F. Baker, M. Borden, J. Krawczyk, August, 1998

RFC2383 ST2+ over ATM Protocol Specification - UNI 3.1 Version, M. Suzuki, August, 1998

RFC2417 Definitions of Managed Objects for Multicast over UNI 3.0/3.1 based ATM Networks, C. Chung, M. Greene, September, 1998 (Obsoletes RFC2366)

RFC2492 IPv6 over ATM Networks, G. Armitage, P. Schulter, M. Jork, January, 1999

RFC2512 Accounting Information for ATM Networks, K. McCloghrie, J. Heinanen, W. Greene, A. Prasad, February, 1999

RFC2515 Definitions of Managed Objects for ATM Management, K. Tesink, Ed, February, 1999 (Obsoletes RFC1695)

RFC2601 ILMI-Based Server Discovery for ATMARP, M. Davison, June, 1999, ASCII

RFC2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5, D. Grossman, J. Heinanen, September, 1999, ASCII (Obsoletes RFC1483)

RFC2761 Terminology for ATM Benchmarking, J. Dunn, C. Martin, February, 2000

RFC2823 PPP over Simple Data Link (SDL) using SONET/SDH with ATM-like framing, J. Carlson, P. Langner, E. Hernandez-Valencia, J. Manchester, May, 2000, ASCII

RFC2844 OSPF over ATM and Proxy-PAR, T. Przygienda, P. Droz, R. Haas, May, 2000

RFC2955 Definitions of Managed Objects for Monitoring and Controlling the Frame Relay/ATM PVC Service Interworking Function, K. Rehbehn, O. Nicklass, G. Mouradian, October, 2000

ITU-T Recommendations:
http://www.itu.int/itudoc/itu-t/rec/

ATM Forum Specifications:
http://www.atmforum.com/atmforum/specs

The IETF RFCs :
http://www.ietf.org/rfc.html

10.2 Troubleshooting ATM

In addition to protocol analyzers and cable testers for twisted pair and fiber optics, troubleshooting in ATM networks involves the use of ATM switch and node management software that is able to track and display the various ATM Operations and Maintenance (OAM) information flows.

ATM contains a number of powerful OAM functions. Because ATM is based on a switched architecture, these integrated monitoring functions are very important; it is no longer possible to monitor the entire activity in a network from a single point, as in traditional network technologies such as Ethernet, Token Ring or FDDI. Monitoring of a single ATM connection only yields information about the traffic between the two connection endpoints (such as a computer system and an ATM switch port). The first step in diagnosing problems in an ATM network is to monitor and analyze data obtained from operation logs and OAM statistics of the various ATM network nodes. Although many of today's ATM components only support, interpret, or display a small proportion of the ATM diagnostics functions, a protocol analyzer can be used to analyze all five OAM flows (F1–F5) and determine whether they report a problem or not.

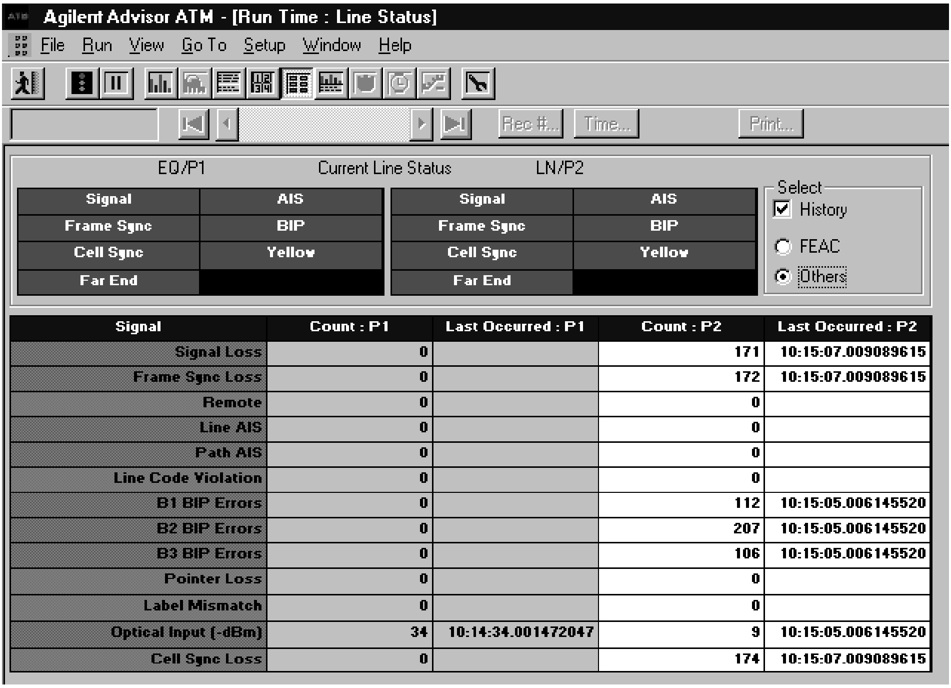

Signal	Count : P1	Last Occurred : P1	Count : P2	Last Occurred : P2
Signal Loss	0		171	10:15:07.009089615
Frame Sync Loss	0		172	10:15:07.009089615
Remote	0		0	
Line AIS	0		0	
Path AIS	0		0	
Line Code Violation	0		0	
B1 BIP Errors	0		112	10:15:05.006145520
B2 BIP Errors	0		207	10:15:05.006145520
B3 BIP Errors	0		106	10:15:05.006145520
Pointer Loss	0		0	
Label Mismatch	0		0	
Optical Input (-dBm)	34	10:14:34.001472047	9	10:15:05.006145520
Cell Sync Loss	0		174	10:15:07.009089615

Figure 10.59 F1–F3 OAM flow statistics with the Agilent Technologies Advisor ATM

Analysis of the integrated network diagnostic functions and the OAM flows is usually sufficient to isolate the systems that are affected by the problem. The troubleshooting process then entails basic functional tests of these components, including loopback tests on the ATM interfaces, firmware tests (is the firmware active?) and hardware self-tests. Cabling and connectors are checked by running bit-error ratio tests and performing OTDR and LED/laser power spectrum measurements. (OTDRs, or optical time domain reflectometers, are physical-layer test instruments for fiber connections. They send defined light pulses over the fiber and measure amplitude and response time of the reflected return signal. The test results include the fiber length and all attenuation components—splices, connectors, fiber loss—along the segment.) If the components pass the basic functional tests, OAM flows F1–F3 in the SDH/SONET layer are examined. If no fault is found here, the ATM layer and the application protocols above it must be analyzed. This begins with checking whether the required PVCs and SVCs are active and working, and whether the ATM addresses are correct. To determine whether the traffic contract parameters for the connections or applications in question are being met, characteristic ATM-layer traffic parameters are measured, including

- User cell rate
- Cell loss
- Cell delay
- Number of cell-sync losses
- Number of cells with corrected headers
- ATM payload bit-error ratio

Finally, if the ATM layer seems to work correctly, the application layer protocols, such as Classical IP over ATM, IP encapsulation, LAN emulation, or PNNI, must be examined.

Each step in the ATM troubleshooting process outlined previously is discussed in detail in the following sections.

10.2.1 Troubleshooting the Physical Layer

Once the error domain has been located, the troubleshooting process can start. If connections are interrupted or network nodes are down, the first step usually consists of basic functional tests of the component systems. Do the activity LEDs of the interface in the problem domain indicate normal working order? Most ATM interfaces indicate normal sending and receiving by a green LED, SDH/SONET-level problems by a yellow LED, and complete signal loss by a red LED. The ATM interface can also be tested using a fiber loop or a UTP loopback connector. If the loopback connection is in order but the loopback test fails, the firmware may not be loaded correctly or may not detect the hardware. If the loopback test is successful, the ATM interface and the firmware are in order. In this case, you must test whether the physical layer connection exists between the network interfaces in the problem domain. This can be done by inserting a protocol analyzer in pass-through mode into the connection between the two nodes. The analyzer may be inserted directly into the ATM connection (active monitoring), or may be connected in passive mode by means of optical power

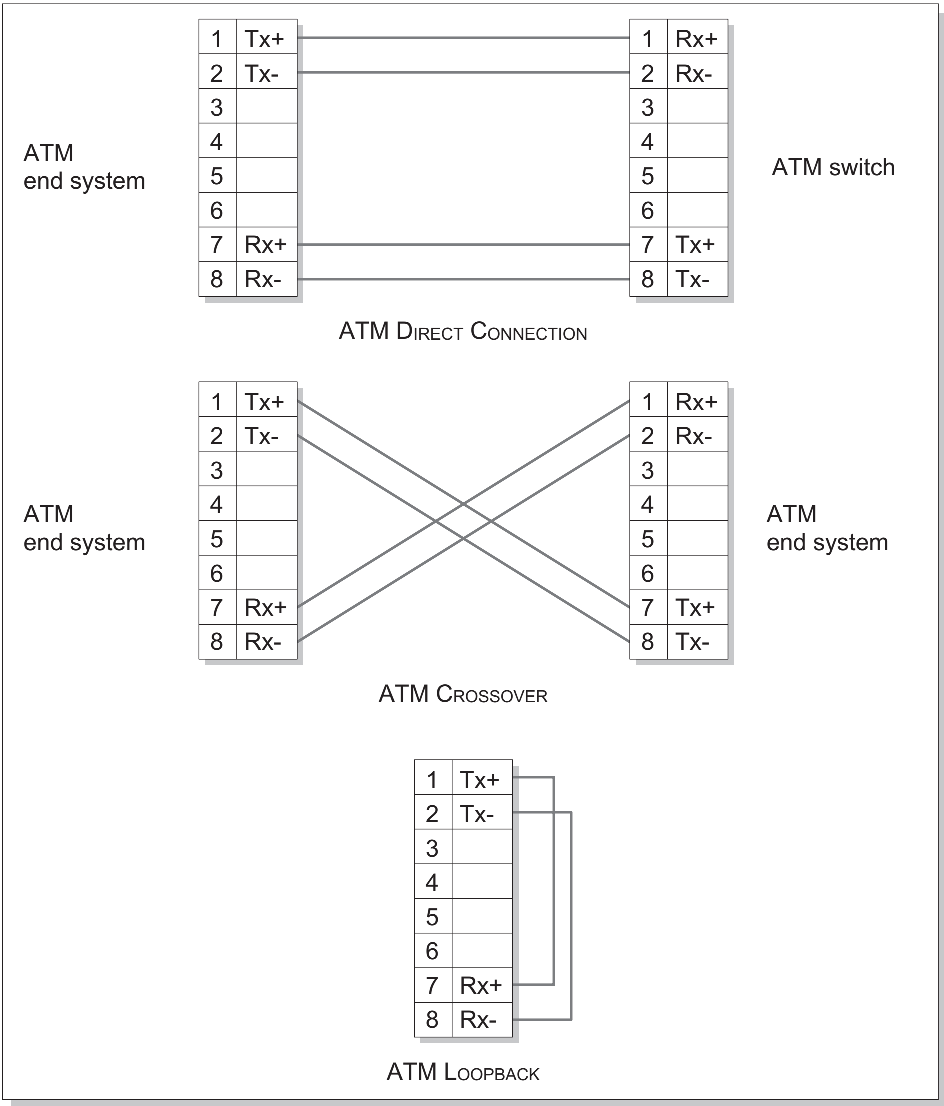

Figure 10.60 UTP Cat. 5 pin assignments for ATM direct connection, crossover, and loopback cables

splitters. Note that the transmitter ports of many analyzers are equipped with single-mode lasers. When actively monitoring multimode fiber components, the transmitters must be fitted with 10 dB attenuators in order to avoid overdriving the receiver electronics. This does not result in damage to the multimode receivers, however. Because lasers still emit a small amount of light when “off” (that is, when sending a “0”): sensitive multimode receivers misinterpret this as a “1” and never see the difference between “0” and “1”. Attenuators can correct this.

10.2.1.1 Analyzing Physical Layer OAM Information Flows

ATM’s integrated error detection mechanisms are contained in the OAM information flows F1 to F5. Flows F1–F3 yield information about the operating state of the SDH/SONET transport structure, while F4 and F5 contain the corresponding ATM layer data. F4 concerns ATM virtual path connections (VPCs), F5 the virtual channel connections (VCCs). The error management function in ATM is based on two types of alarms: Alarm Indication Signal (AIS) and Remote Defect Indicator (RDI). The AIS is sent by the VC or VP node, which recognizes the error condition to all upstream nodes so long as the error condition persists. Immediately after the AIS, an RDI signal is sent upstream to the end nodes of the connections affected. These signals are also sent periodically until the error condition is resolved. VP-AIS and VP-RDI messages are always sent in cells with VCI = 4, while VC-AIS and VC-RDI messages are sent in cells with PT = 101.

Two mechanisms are available to detect error conditions: continuity checks (CC) and loopback tests. Continuity checks continuously monitor the availability of a connection. To this end, CC cells are periodically inserted into the user cell stream. ATM network nodes along the connection path can then check for the presence of these CC cells. When no more cells are received, an AIS alarm for loss of continuity (LOC) is triggered.

If the ATM network components support OAM cell processing, they can often locate the failure domain by analyzing the contents of the OAM flows. If not, the OAM flows must be captured and analyzed using a protocol analyzer.

10.2.1.2 Verifying ATM Cell Transmission on the Physical Layer

The analysis of ATM cell transmission parameters with the help of a protocol analyzer can also provide information about problems in the physical layer. The traffic parameters to examine include corrected header ratio, discarded cell ratio, loss of cell delineation rate, and the demux error ratio.

Corrected Header Ratio

The corrected header ratio is the number of cells with errored but correctable headers divided by the total number of cells received. This parameter is mainly influenced by the bit-error ratio of the transmission path. There is a small probability that cells with errored headers may appear as valid cells, and thus lead to incorrect transmissions (misinserted cells). The probability of such an event can be calculated from the number of errored headers containing more than two incorrect bits—the HEC checksum of such a header no longer indicates

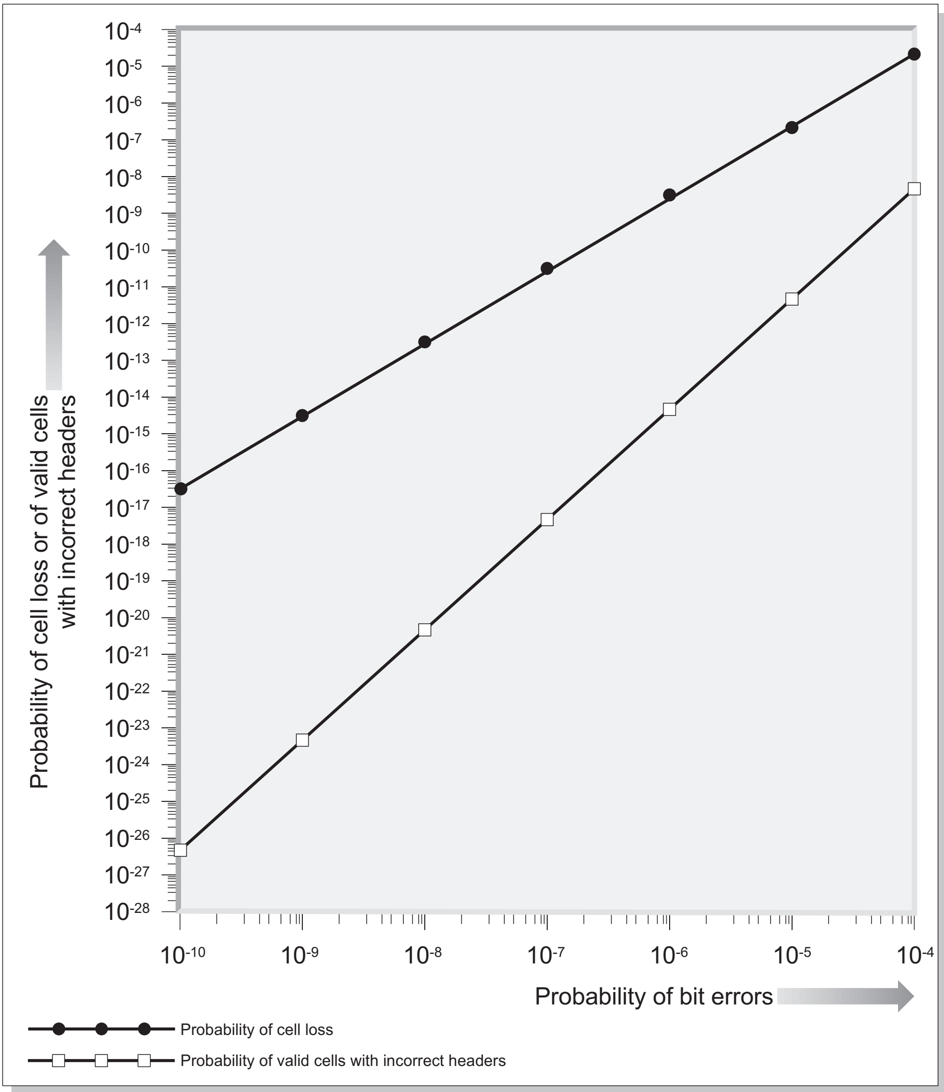

Figure 10.61 Probability of the transmission of errored cell headers as valid cells in relation to the bit-error ratio

whether it is corrupt—and the ratio of valid header values to the number of all possible header values. Figure 10.61 shows the probability of the transmission of errored cell headers as valid cells in relation to the bit-error ratio.

Discarded Cell Ratio

The discarded cell ratio equals the number of cells received with errored headers that cannot be corrected and which are therefore discarded, divided by the total number of cells received (valid or not). This quantity is also influenced by the bit-error ratio of the transmission path.

Loss of Cell Delineation Rate

The loss of cell delineation rate is the number of cell synchronization losses over a certain time interval.

Mean Loss of Delineation Duration

The mean loss of delineation duration is defined as the number of missing Cell Received Events (CRE_2) due to cell synchronization loss within a given time interval divided by the total number of expected CRE_2 events in this time interval (see Figure 10.66).

Demux Error Ratio

The demux (demultiplex) error ratio is the number of all correctly transmitted cells containing an invalid VPI value divided by the total number of correctly transmitted cells.

10.2.1.3 Causes of Problems in the Physical Layer

Physical layer problems in ATM networks may have a variety of causes. Many problems occur on the connections to ATM switches, where a change of the physical transmission media is required (connectors, fiber, switch port). Transmission errors can arise due to aging, humidity, dust, or material flaws. Furthermore, the signal quality of the cabling determines the bit-error ratio in the transmission framework (SDH/SONET, E3, T3, etc.) and consequently the performance of the ATM network.

10.2.2 Troubleshooting the ATM Layer

If no faults can be detected in the physical layer, the ATM layer must be examined. This involves checking whether the required PVCs and SVCs are working, and whether the ATM addresses are correct. The OAM PM cell streams are then monitored using protocol analyzers to determine whether the ATM layer traffic conforms to the traffic contract parameters for the connections or applications in question. Parameters to examine include:

- User cell rate
- Cell loss ratio (CLR)
- Cell transfer delay (CTD)
- Cell delay variation (CDV)
- Number of cell-sync losses
- Number of cells with corrected headers
- Number of cells that violate the traffic contract (non-conforming cells or NCC)

10.2.2.1 Verifying PVCs, SVCs and Addressing

The first step in troubleshooting the ATM layer is to verify whether the ATM connections are working at all. Configuring one ATM node to send a constant stream of pings to another can do this. If the ping does not get through to its destination, use the activity LEDs on the ATM interfaces or a protocol analyzer to determine whether the cells leave the initiating system, reach and leave the switch, and arrive at the interface of the destination node. If no ping packets

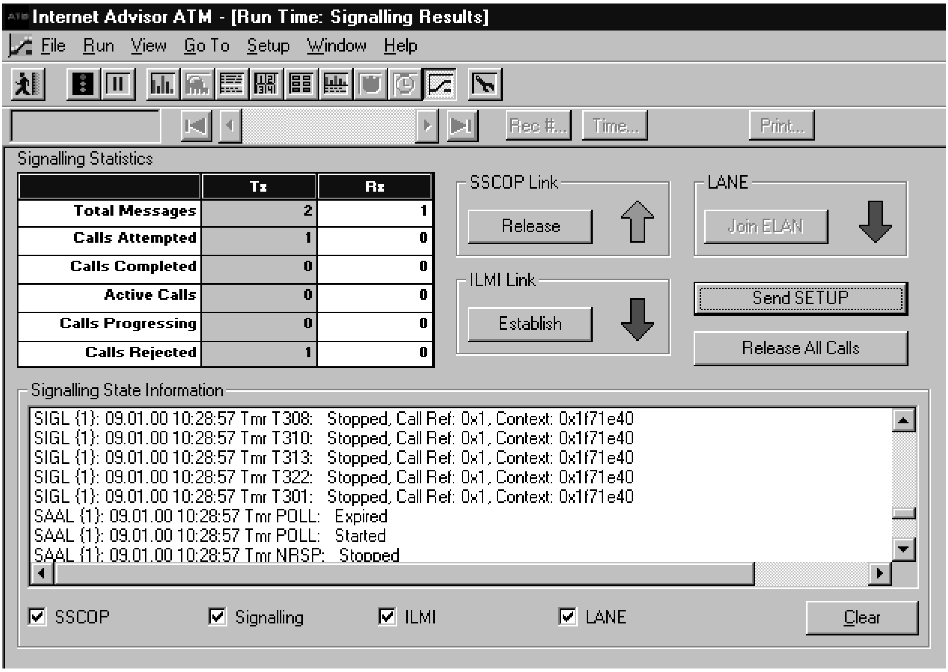

Figure 10.62 Setting up and releasing SVCs with the Agilent Technologies Advisor

arrive, check on the switch whether the PVC is set up at all, the VCI values are correctly configured, and the internal path between the two switch ports is functional.

If you use SVCs you must make sure the signaling process that sets up your connections is working. If pings do not go through between two SVC nodes, the first step once again is to check the configuration of the switch and the end systems. In the case of Classical IP over ATM, are the systems registered on the ATMARP server with their ATM and IP addresses? Is the correct address of the ARP server registered on the clients? If so, at least the SVC setup between the clients and the ARP server should succeed. If the clients are still unable to communicate with one another, make sure the ILMI stack is active on the switch and on the clients. If the problem still persists, you must monitor and analyze the signaling process with a protocol analyzer. Finally, the switch ports and the switch configuration should be examined in detail.

10.2.2.2 Verifying ATM Performance Parameters

If the connections are set up successfully, but problems persist on the application layer and during data transfers, the next step is to examine the ATM performance parameters. Oftentimes functional tests of ATM connections (such as sending pings between two nodes, setting up and clearing down SVCs, etc.) seem to show that everything is working fine, but once application data is transmitted at higher traffic loads problems arise. Reasons for this type of behavior can be excessive cell loss or cell delay values, traffic contracts that provide insufficient bandwidth and therefore cause cells to be discarded, or simply an overloaded switch. Measurements that determine these types of ATM performance parameters can be made either in-service or out-of-service. Out-of-service measurements are performed using special out-of-service test cells defined in ITU-T Recommendation O.191.

5 bytes	4 bytes	4 bytes	37 bytes	1 byte	2 bytes
ATM header	sequence number	time stamp	not used	TCPT	(CRC10)

TCPT ....Test Cell Payload Type field

Figure 10.63 Format of the O.191 test cell

Out-of-service test cells have standard ATM cell headers and can be sent using any VPI/VCI label value (VCI > 31). In the payload they carry a 32-bit sequence number to permit detection of cell loss and cell misinsertion errors, and a 32-bit time stamp to measure cell delay and cell delay variation. This allows CDV and 2-

point CDV measurements up to transmission speeds of 2.4 Gbit/s. The least significant bit of the time stamp has a granularity of 10 ns, though for physical links slower than 2.4 Gbit/s the time stamp is normally incremented from a higher order bit. In the simplest case, a protocol analyzer with one transmit and one receive port is sufficient to perform these out-of-service measurements. Care must be taken with using loopbacks, however, because the traffic contracts of virtual circuits may be asymmetrically specified for the different directions.

A very useful extension to the use of ITU-T Recommendation O.191 test cells is available if the test cells can be shaped to simulate traffic that only just meets the traffic contract in terms of peak cell rate (PCR), cell delay variation tolerance (CDVT), sustainable cell rate (SCR), and maximum burst size (MBS). Shaped test cell traffic can be injected into a network at the UNI and on the far side UNI, and measurements can be made of the delivered QoS in terms of cell loss, CDV, etc. With a pair of analyzers capable of simultaneously sending shaped traffic and analyzing received traffic, bi-directional measurements can be made simultaneously. Asymmetrical traffic contracts can be tested without difficulty in this configuration because each analyzer can independently shape traffic according to the traffic contract for that direction. The fact that it may be impossible to synchronize the clocks of the analyzers with respect to phase as well as frequency need not invalidate the most useful measurements because often the absolute delay is less important than CDV and cell loss.

For in-service measurements, however, more sophisticated test equipment with at least two independent but time-correlated receivers and transmitters is required. Such a tester can be inserted into the lines of transmit and receive ports of the ATM element under test, as shown in Figure 10.64.

If supported by the network components, in-service tests can also be carried out with the help of special performance management OAM cells. These cells are periodically inserted in the user cell streams of the connections to be monitored. The measurement parameters contained in the payload of these cells (sequence number, user cell count, time stamp, cell loss counter) provide information about the operational condition state of the ATM connection.

The ATM layer performance parameters are defined in the following sections.

Cell Error Ratio

The cell error ratio is the number of invalid cells divided by the sum of the number of successfully transmitted cells and the number of invalid cells. Successfully transferred cells, tagged cells, and errored cells contained in severely errored cell blocks are excluded from the calculation of cell error ratio.

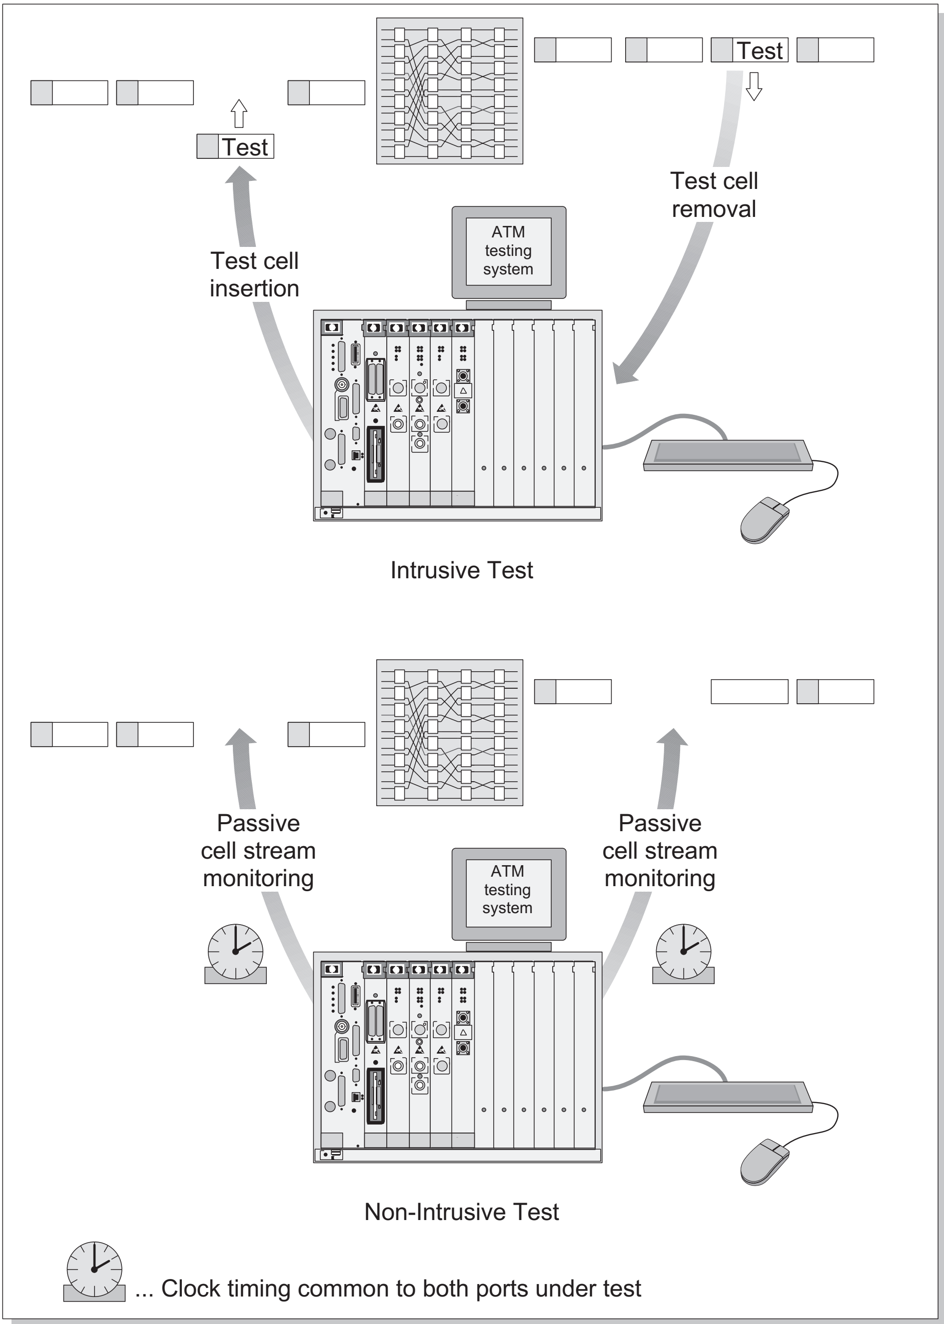

Figure 10.64 Intrusive and non-intrusive tests in ATM networks

Cell Loss Ratio

The cell loss ratio is the number of cells lost divided by the total number of cells transmitted. Lost cells and transmitted cells in severely errored cell blocks are excluded from the calculation of cell loss ratio. There are three different cell loss ratio measurements:

a) Cell loss ratio for cells with high priority (cell loss priority bit = 0): CLR_0
   If $N_t(0)$ is the number of cells with CLP = 0 and $N_l(0)$ is the number of lost cells plus the number of tagged cells, then CLR_0 is defined as $N_l(0) / N_t(0)$.

b) Cell loss ratio for the entire cell stream: CLR_{0+1}
   If $N_t(0 + 1)$ is the number of all cells transmitted and $N_l(0 + 1)$ is the number of lost cells, CLR_{0+1} is defined as $N_l(0 + 1) / N_t(0 + 1)$.

c) Cell loss ratio for cells with low priority: CLR_1
   If $N_t(1)$ is the number of cells with CLP = 1 and $N_l(1)$ the number of lost cells, CLR_1 is defined as $N_l(1) / N_t(1)$.

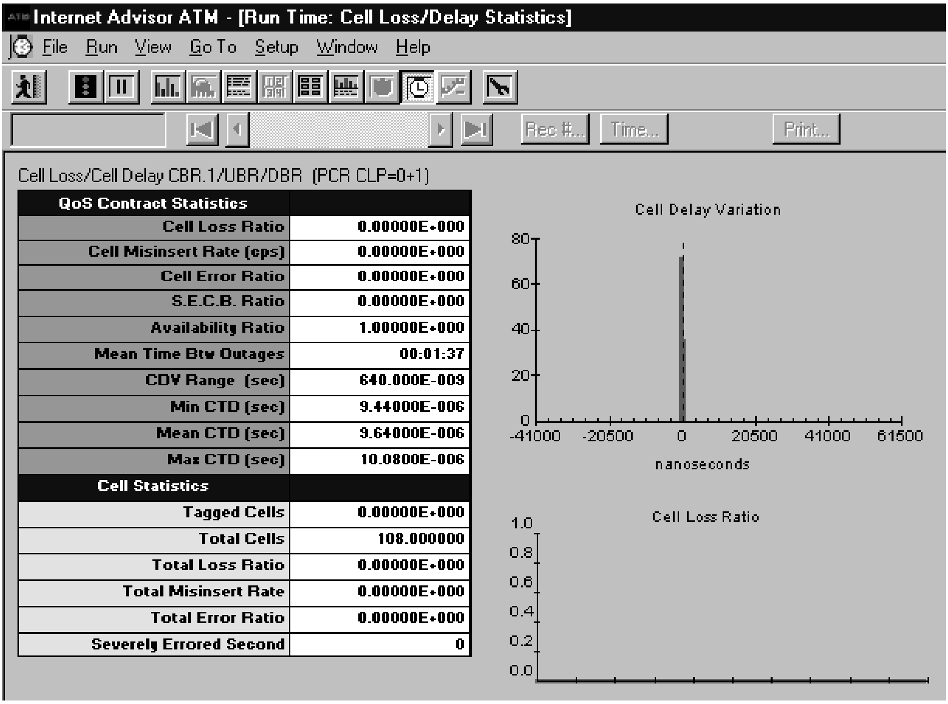

Figure 10.65 Determining CDV using a protocol analyzer

Cell Misinsertion Rate

The cell misinsertion rate is defined as the number of defective cells (cells containing a wrong VPI/VCI due to non-corrected header errors) transmitted within a time interval divided by this interval.

Cell Transfer Delay (CTD)

The cell transfer delay is defined as the time $t_2 – t_1$ between two corresponding cell transmission/reception events CRE_1 (t_1) and CRE_2 (t_2) (where $t_2 > t_1$).

Cell Delay Variation (CDV)

Two types of variations in the cell transfer delay are defined: one-point cell delay variation, which examines cells arriving at one measurement point, and two-point cell delay variation, which examines cell delay variation at measurement point two relative to measurement point one.

One-Point Cell Delay Variation

The one-point cell delay variation (y_k) for cell k at measurement point MP is defined as the difference between the reference arrival time (c_k) of the cell and

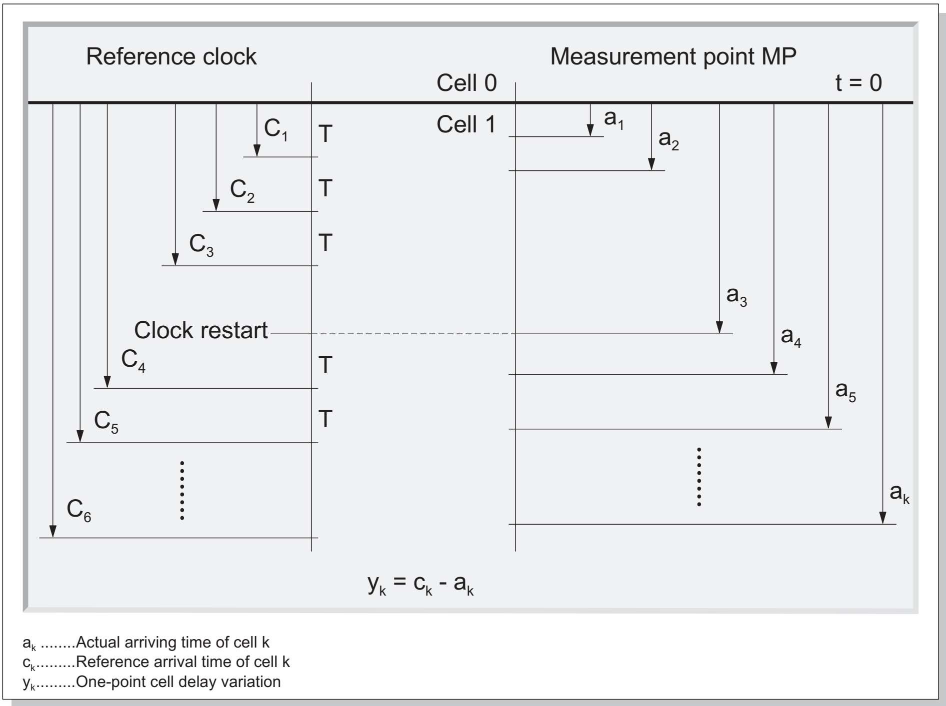

Figure 10.66 One-point CDV

the actual arrival time (a_k), that is, $y_k = c_k - a_k$, where the reference arrival time is defined as follows:

$$c_0 = a_0 = 0$$

$$c_{k+1} = c_k + T \qquad \text{if } c_k > a_k$$

$$c_{k+1} = a_k + T \qquad \text{in all other cases}$$

Two-Point Cell Delay Variation

The two-point cell delay variation (v_k) for cell k between MP1 and MP2 is defined as the difference between the actual cell delay (x_k) and the reference delay ($d_{1,2}$) between the two measurements points $v_k = x_k - d_{1,2}$. The actual cell delay (x_k) is defined as the difference between the actual cell arrival time at MP2 ($a2_k$) and the actual arrival time at MP1 ($a1_k$): that is, $x_k = a2_k - a1_k$. The reference cell delay ($d_{1,2}$) between MP1 and MP2 equals the actual cell delay of cell 0 between the two measurement points.

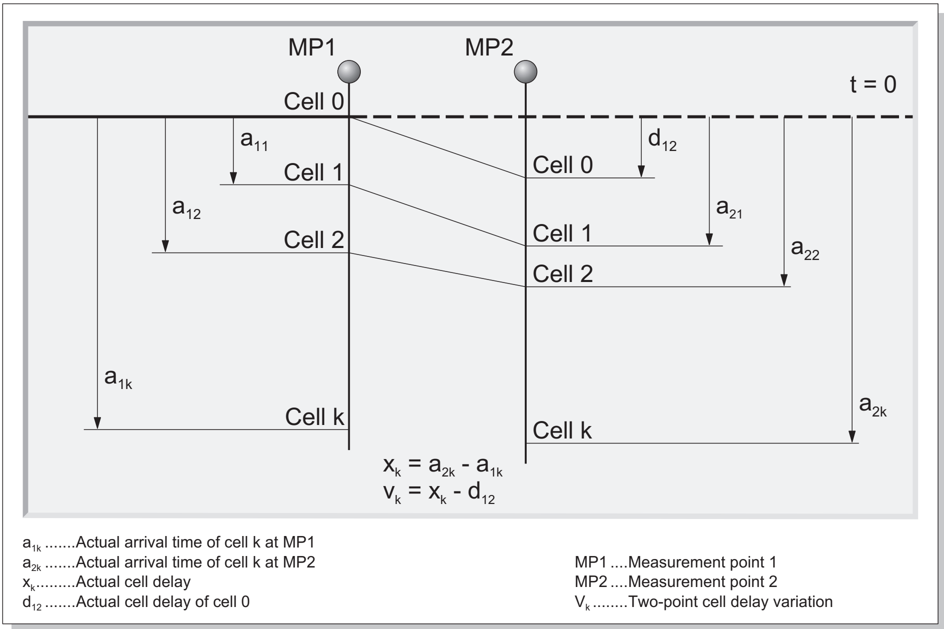

Figure 10.67 Two-point CDV

Figure 10.68 shows an overview of the ATM layer performance parameters for the various QoS classes, that can be attained in ATM wide-area networks with the reference diameter of 27,500 km.

			QoS Class			
	Limit	Default	Stringent	Tolerant	Bi-level	U
CTD	Maximum average CTD	none	400 ms	U	U	U
2-point CDV	Maximum difference between upper and lower 10^{-8} range of CTD	none	3 ms	U	U	U
CLR_{0+1}	Maximum cell loss probability	none	$3 \cdot 10^{-7}$	10^{-5}	U	U
CLR_0	Maximum cell loss probability	none	none	none	10^{-5}	U
CER	Maximum cell error probability	$4 \cdot 10^{-6}$	default	default	default	U
CMR	Maximum cell misinsertion rate	1/day	default	default	default	U
SECBR	Maximum SECB probability	10^{-4}	default	default	default	U
U ....unspecified or unlimited						

Figure 10.68 ATM layer performance parameters

10.2.2.3 Symptoms and Causes

The most frequent causes of cell loss are buffer overflows or faults in the physical layer that lead to non-correctable errors. Cell misinsertion is caused by multiple bit transmission errors in the header, which can no longer be corrected as a result of physical layer problems or malfunctions in the switching fabric. Cell errors usually indicate the occurrence of bit errors in the payload field (bit errors in the header are reflected in cell loss figures). In most cases these bit errors are caused by a higher degree of signal jitter than the ATM interface can tolerate. Cell transfer delay is caused by ordinary electronic switching and signal propagation delays. The cause of cell delay variations usually lies in the varying states of buffers that the cells must pass through on their way to the destination, and in the effects of cell encapsulation in the physical layer transmission framing. Thus two cells within a single SDH container (SONET SPE) will have a smaller CDV relative to one another than two cells transported in different containers /SPEs.

Figure 10.69 lists the typical causes of problems in the ATM layer, grouped by symptoms: Cell Error Rate (CER), Cell Loss Ratio (CLR), Cell Misinsertion Rate (CMR), Mean Cell Transfer Delay (MCTD), Cell Delay Variation (CDV).

Sources of error	CER	CLR	CMR	MCTD	CDV
Signal propagation delay				X	
Fault in communication medium	X	X	X		
Switch architecture		X		X	X
Buffer capacity		X		X	X
Number of nodes along a given VPC/ VCC connection	X	X	X	X	X
Network load		X	X	X	X
Error		X			
Bandwidth allocation to a given VPC/ VCC		X		X	X

Figure 10.69 Symptoms and causes of problems in ATM networks

10.2.3 Troubleshooting Higher Layers

If the ATM layer is working and the problem persists, the higher layer protocols in use must be analyzed. Two of the most common higher layer protocols—besides UNI signaling—are LAN Emulation (LANE) and the Private Network to Network Interface (PNNI).

10.2.3.1 LAN Emulation

If problems occur in LAN emulation environments, the first step is to make sure that the connected traditional LANs (10/100/1,000 Mbit/s Ethernet, FDDI, etc.) and the LAN interfaces of the LAN/ATM interworking devices are functioning correctly. This includes verifying the various LAN configuration settings and measuring basic network statistics with a protocol analyzer. Examination of the LAN emulation components begins only after the LAN part has been proven to be working properly. The first step in LANE troubleshooting is to send pings between two LE Clients and check whether a connection can be set up at all. If the ping does not go through, check the IP interfaces of the LE Clients. Examine whether the IP interfaces are active at all, and whether the IP addresses and subnet masks are correct and in the same subnet. Then determine whether the LANE software on the clients is active, and whether both LE Clients belong to the same ELAN. If no error is found, the following configuration parameters must be systematically checked through the system management interface of the ATM components:

- Are both LE Clients registered on the same LE Server and Broadcast-Unknown Server (BUS)?
- Is the ATM address of the primary (and secondary, if configured) LE Server correct, and are the Configuration Direct VCCs set up?
- Is the ATM address of the primary (and secondary) BUS correct, and are the Multicast Send VCCs set up?

If it is still impossible to set up data VCCs between the LE Clients, the last error cause to check for is a restricted VC capacity on one of the systems due to traffic contracts. This is done by verifying the compatibility of the traffic contracts for the LE Client interfaces. If a protocol analyzer is available, a trace of the unsuccessful connection setups is the fastest way to find the cause of the problem.

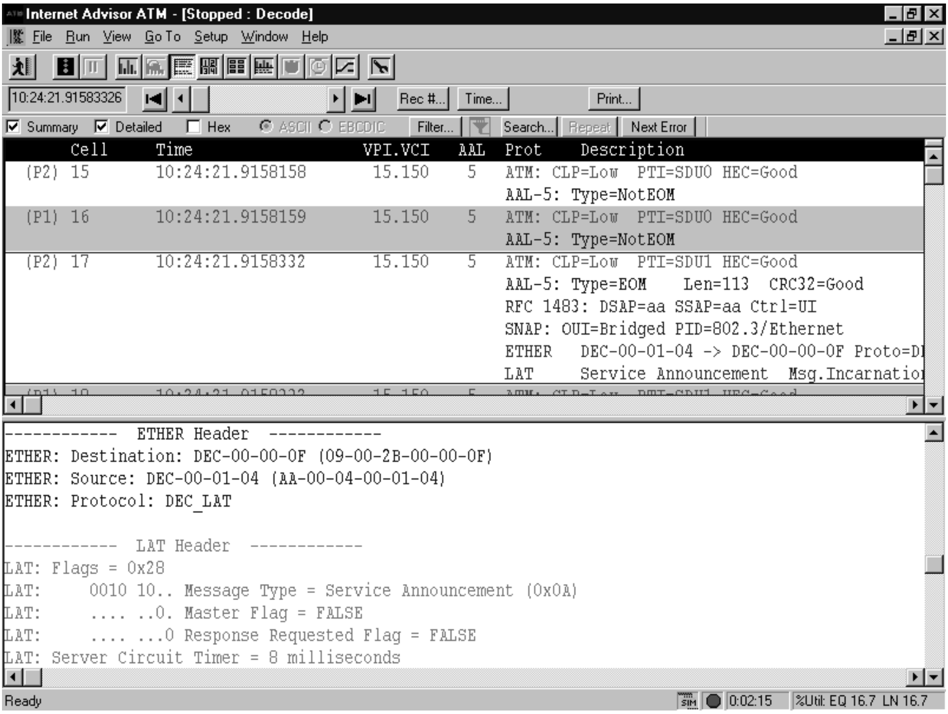

Figure 10.70 LANE analysis by means of a protocol analyzer

10.2.3.2 UNI Signaling

The first steps in troubleshooting ATM UNI signaling problems are the following basic checks:

- Are the interfaces of the affected systems active?
- Are the affected systems using compatible signaling variants (for example, UNI 3.1, 4.0)?
- Are ILMI (if supported) and the SSCOP layer active?

If no error is detected, the next step is to try to set up an SVC connection while using the debug mode of the ATM nodes or a protocol analyzer to trace the signaling process. Such an analysis of the signaling process usually leads to the cause of the problem. Typical causes are invalid called party or calling party addresses; invalid, unknown or disordered mandatory information elements; invalid call reference numbers; or rejection of the connection setup by a RELEASE message from the destination station.

The most important error states in UNI 3.1/4.0/Q.2931 signaling that can occur during the connection setup and clear-down processes are described in the following sections:

Invalid protocol discriminator

Messages with an invalid protocol discriminator are discarded.

Short messages

Messages that are too short to contain a complete information element are discarded.

Invalid call reference format

If bytes 1 and 5 through 8 of the call reference information element are not set to the value 0, or if the call reference length field contains a value other than 3, the message is discarded.

Invalid call reference

a) Whenever any message except SETUP, RELEASE, RELEASE COMPLETE, STATUS ENQUIRY, or STATUS is received with a call reference that does not refer to an active call or to a call in progress, the receiver shall initiate clearing by sending a RELEASE COMPLETE message with cause No. 81, “Invalid call reference value’’, specifying the call reference of the message received, and shall remain in the null state.

b) When a RELEASE COMPLETE message is received with a call reference that does not refer to an active call or to a call in progress, no action should be taken.

c) When a SETUP message is received with a call reference that does not refer to an active call or to a call in progress, and with a call reference flag incorrectly set to "1", the message shall be ignored.

d) When a SETUP message is received with a call reference that *does* refer to an active call or to a call in progress, the SETUP message shall be ignored.

e) When any message except RESTART, RESTART ACKNOWLEDGE, or STATUS is received with the global call reference, no action should be taken on this message, but a STATUS message shall be returned using the global call reference with a call state indicating the current state associated with the global call reference and cause No. 81, "Invalid call reference''.

f) When a STATUS message is received specifying a call reference that is not recognized as relating to an active call or to a call in progress, it shall be cleared with cause 101 "Message not compatible with call state". Alternatively, any other action specific to the implementation that attempts to recover from this mismatch can be taken.

g) If a STATUS or a STATUS-ENQUIRY message is received with a call reference that does not refer to an active call or to a call in progress, a STATUS ENQUIRY message shall be sent to check the correctness of the call state.

h) When a RESTART message is received specifying the global call reference with a call reference flag incorrectly set to "1", or when a RESTART ACKNOWLEDGE message is received specifying the global call reference with a call reference flag incorrectly set to "0", no action should be taken on this message, but a STATUS message shall be returned with a call state indicating the current state associated with the global call reference and cause No. 81, "Invalid call reference''.

Message type or message sequence errors

Whenever an unexpected message is received, except RELEASE, RELEASE COMPLETE, or when an unrecognized message is received, no state change shall occur and a STATUS message shall be returned with one of the following causes:

a) No. 97, Message type non-existent or not implemented
b) No. 101, Message not compatible with call state

Two exceptions to this procedure exist, however. The first is when the network or the user receives an unexpected RELEASE message in response to a SETUP message. In this case, a no STATUS or STATUS ENQUIRY message is sent. Whenever the network receives an unexpected RELEASE message, the network shall release the virtual channel, clear the network connection and the call to the remote user indicating the cause received in the RELEASE message sent by the user or, if no cause was included, cause No. 31, "Normal, unspecified''. Furthermore, the network shall return a RELEASE COMPLETE message to the user, release the call reference, stop all timers, and enter the null state. Whenever the user receives an unexpected RELEASE message, the user shall release the virtual channel, return a RELEASE COMPLETE message to the network, release the call reference, stop all timers, and enter the null state.

The second exception is when the network or the user receives an unexpected RELEASE COMPLETE message. Whenever the network receives an unexpected RELEASE COMPLETE message, the network shall disconnect and release the virtual channel, clear the network connection and the call to the remote user indicating the cause given by the user or, if no cause was included, cause No. 111, "Protocol error, unspecified''. Furthermore, the network shall release the call reference, stop all timers, and enter the null state. Whenever the user receives an unexpected RELEASE COMPLETE message, the user shall disconnect and release the virtual channel, release the call reference, stop all timers, and enter the null state.

Information Element Sequence

Information elements must be sent in the following order:

- Protocol discriminator
- Call reference
- Message type
- Message length
- Other information elements

Information elements of variable length can be sent in any order.

Duplicate Information Elements

If an information element is repeated in a message in which repetition of the information element is not permitted, only the contents of the information element appearing first shall be handled. All subsequent repetitions of the information element are ignored.

Mandatory Information Element Missing

When a message other than SETUP, RELEASE, or RELEASE COMPLETE is received that lacks one or more mandatory information elements, no action shall be taken on the message and no state change should occur. A STATUS message is then returned with cause No. 96, "Mandatory information element is missing''.

When a SETUP message is received that lacks one or more mandatory information elements, a RELEASE COMPLETE message is returned with cause No. 96, "Mandatory information element is missing''.

Mandatory Information Element Content Error

When a message other than SETUP, RELEASE, or RELEASE COMPLETE is received in which one or more mandatory information elements have invalid contents, no state change occurs. A STATUS message is returned with cause No. 100, "Invalid information element contents''. When a SETUP message is received in which one or more mandatory information elements has invalid contents, a RELEASE COMPLETE message is returned with cause No. 100, "Invalid information element contents''.

Unrecognized Information Element

If a message is received that contains one or more unknown information elements, action is taken on the message and those information elements that are recognized and have valid contents. If the received message is not a RELEASE or RELEASE COMPLETE, a STATUS message is returned containing one cause information element. The information element contains cause No. 99, "Information element non-existent or not implemented", and the diagnostic field, if present, contains the information element identifier of each unrecognized information element.

When a RELEASE message is received that has one or more unrecognized information elements, a RELEASE COMPLETE message with cause No. 99, "Information element non-existent or not implemented'' is returned. The cause information element diagnostic field, if present, contains the information element identifier of each unrecognized information element. A RELEASE COMPLETE message with unknown information elements is ignored completely.

If a message contains one or more information elements with contents that are in part invalid, then action is taken on those information elements that appear correctly. A STATUS message is also sent with cause No. 100, "Invalid information element contents", and the information element identifier of each invalid information element in the diagnostic field. If address information fields are also corrupt, then cause 43, "Access information discarded", is sent in place of cause 100. If an information element is recognized but should not be present in the given message, it is treated as an unrecognized information element.

AAL Signaling Error

If an AAL signaling error occurs, all connections not yet started are initialized and a T309 timer is started for each active connection. Then a restart of the AAL signaling layer is initiated. If a connection's T309 expires before the signaling layer can be restarted, that connection is deactivated with cause 27, "Destination out of order", and its call reference is deleted.

Status Enquiry

A STATUS ENQUIRY message can be sent to check the call state at a peer entity. Furthermore, whenever the SAAL indicates that a disruption has occurred at the data link layer, a STATUS ENQUIRY message is sent to check for a correct call state at the peer entity. When the STATUS ENQUIRY message is sent, timer T322 is started in anticipation of an incoming STATUS message. Only one unanswered STATUS ENQUIRY may be outstanding at any given time. The receiver of a STATUS ENQUIRY message responds with a STATUS message indicating cause 30, "Response to STATUS ENQUIRY", and reporting the current call state.

If no STATUS response is received before T322 expires, the STATUS ENQUIRY can be repeated one or more times, depending on the implementation. If the timer expires after the last attempt, the connection is cleared down with cause 41, "Temporary failure".

Procedure on Receipt of a STATUS Message

When a STATUS message is received that indicates that the peer station is in an incompatible state for call handling, the connection can be cleared down with cause 101, "Message not compatible with call state", or—if so implemented—an attempt may be made to correct the fault. The decision as to whether the two stations' call states are incompatible with one another is left to the given implementation, except in the following three cases:

a) If a STATUS message is received signaling that the peer station is in a state other than null, and the station receiving the STATUS message is in the null state itself, then the receiver responds with a RELEASE message and cause 101, "Message not compatible with call state".

b) If a STATUS message is received signaling that the peer station is in a state other than null, and the station receiving the STATUS message is in the "Release Request" state, the receiver shall not respond.

c) If a STATUS message is received signaling that the peer station is in the null state, and the station receiving the STATUS message is not in the null state, then the receiver of the STATUS message shall change to the null state.

If a STATUS message is received that signals a compatible call state but contains cause 96, 97, 99, 100 or 101, the response is left to the given implementation. If no particular reaction is specified, the connection in question should be cleared down with the cause indicated in the STATUS message received.

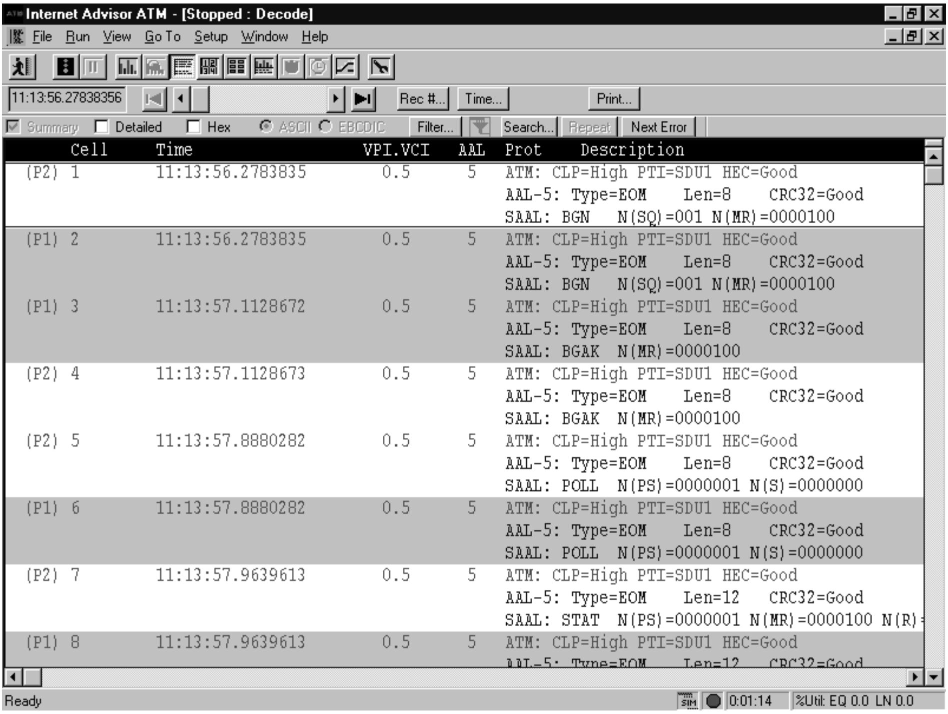

Figure 10.71 UNI signaling decoded by a protocol analyzer for LANE analysis

10.2.3.3 PNNI Signaling

PNNI is the signaling protocol used to set up SVC connections between two NNIs (that is, switch-to-switch signaling). It actually consists of two protocols: the topology protocol, which distributes information about the network topology to the network nodes, and the signaling protocol, which is basically an extension of the ITU-T Recommendation Q.2931 UNI protocol.

Diagnosis of PNNI problems begins with an examination of the UNI signaling functions at the end systems in the problem domain:

- Are the interfaces of the ATM nodes and switches active?
- Are the signaling versions of all systems compatible (UNI 3.1, 4.0)?
- Are ILMI and the SSCOP layer active?

If the end systems begin signaling processes that cannot be finished successfully, the next step is to analyze the PNNI SVC routing. First, read out the PNNI routes determined by the UNI send port of the switch to the destination node. Then check whether an active route to the destination was found. This is done using the vendor's system management interface to read out the PNNI operations log and examine the PNNI routes calculated by the sending node. If the route leads to a wrong node, a wrong ATM address is probably involved. If the switch cannot find a route to the destination at all, its topology information may be incorrect. If the topology information does not contain the destination node, this may be due either to physical problems in the network, or to the switch's inability to update its topology database. If the sending node does have an uninterrupted route to the destination system, but the connectivity problem remains, the trouble may be caused by traffic parameters that cannot be provided by one of the ATM switches.

If the topology information is wrong or incomplete, the first troubleshooting step is to list all peer neighbors that can be reached directly from the UNI port and their PNNI operation states (up, down, changing, loading). Then the topology information (PNNI Topology State Elements or PTSE) of the peer nodes at the lowest hierarchical level are analyzed to determine whether they reach the peer group leader. Conversely, the PTSEs of the peer group leaders must also be examined to verify that they reach the peer nodes at the lowest level. If all the PTSEs are in order and no configuration errors can be detected (peer group

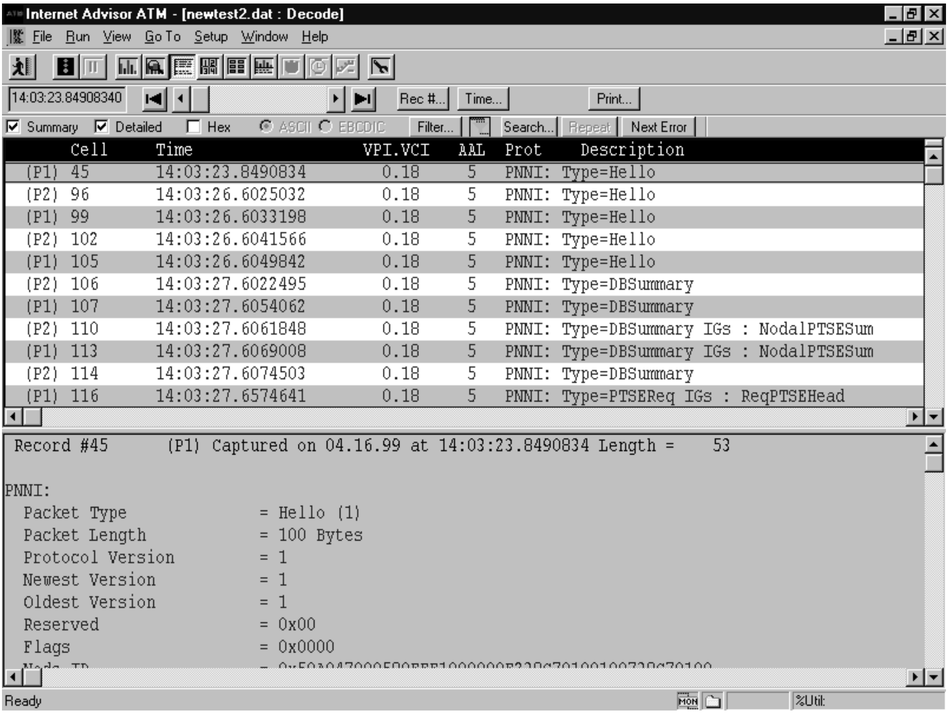

Figure 10.72 Decoding PNNI signaling with a protocol analyzer

leader status is active; parent logical group node (LGN) is configured), then the following checks can be performed successively:

- Is the Hello protocol active on the PNNI?
- Is the PNNI Routing Control Channel (SVCC-RCC) active?
- Are the PNNI port parameters set correctly (cell rate, cell delay, cell rate)?
- Are the peer group uplinks configured and active?
- Are the PNNI addresses (prefixes) and the PNNI short form addresses (summary addresses) correct?

Once again, the fastest way to diagnose the problem is to monitor the PNNI messages during connection setup with a protocol analyzer.

10.2.4 Cabling Problems

Cabling problems are also very common in ATM networks. Typical causes include bad splices, low quality cables, wiring faults, excessive segment length and, for UTP, incorrect characteristic impedances or noise due to electromagnetic interference caused by air condition systems, photocopiers, pagers, elevators or production environments. These types of problems are discussed in detail in Chapter 6.

10.2.5 Problems with ATM Interface Cards

The typical symptoms of defective interface cards in ATM networks are high rates of cell errors or complete loss of cell synchronization. Because of ATM's connection-oriented architecture, it is easy to determine whether the problem is caused by a network interface card. Starting in the middle of the affected connection, the number of cells received is compared with the number of cells transmitted at each ATM interface. If the numbers do not match, the QoS parameters of the given interface must be examined. If no restrictions can be detected, a loopback test shows whether or not cells are being lost due to NIC problems. When changing ATM interface cards, care must be taken due to the high temperatures at which these cards normally operate. Either wait until the card cools, or avoid touching the chips on the card.

10.2.6 Problems with Routers

Routers are internetwork components that connect network segments at OSI Layer 3, and are therefore able to link networks of different topologies. For this reason, there are no troubleshooting issues for routers that apply specifically to ATM networks. Please refer to the troubleshooting section on routers in Chapter 7.

10.2.7 Symptoms and Causes: ATM

Symptom: No Connection over a PVC

Cause (1): Problems with ATM interface card or driver.
Cause (2): No PVC set, selected VCI is incorrect.
Cause (3): Hardware or software problems on the switch.
Cause (4): Misconfigured ATM port (bit rate, scrambling, interface type, frame type (PLCP, SDH/SONET, SONET)).

Symptom: No Connection over SVC (UNI Signaling Problems)

Cause (1): Problems with ATM interface card or driver.
Cause (2): ATMARP server misconfigured; clients are not set up with their correct ATM and IP address on the ATMARP server.
Cause (3): The address of the ATMARP server is not configured correctly on the client system.
Cause (4): ILMI is not active on the client or the server.
Cause (5): The ILMI software versions on client and server are incompatible.
Cause (6): SSCOP layer not established.
Cause (7): Incompatible UNI signaling variants (UNI 3.0, 3.1, 4.0).
Cause (8): Wrong Called Party or Calling Party number.
Cause (9): Unknown or invalid information elements, or mandatory information elements in wrong order.
Cause (10): Incorrect call reference numbers.
Cause (11): Called party is not ready to accept call, call setup attempt is rejected with RELEASE message.
Cause (12): Misconfigured ATM port (bit rate, scrambling, interface type, frame type (PLCP, G.804, SDH, SONET)).
Cause (13): Hardware or software problems on the switch.

Symptom: High Cell Error Rate (CER)

Cause (1): Problems on the physical layer (cabling, connectors, ATM port).
Cause (2): Too many nodes along the transmission path of a VP or VC connection.

Symptom: High Cell Loss Ratio (CLR)

Cause (1): Problems on the physical layer (cabling, connectors, ATM port).
Cause (2): Too many nodes along the transmission path of a VP or VC connection.
Cause (3): ATM switch overloaded.
Cause (4): Insufficient buffering in the switch.
Cause (5): High network load.
Cause (6): Limits of traffic contract are exceeded.

Symptom: High Cell Misinsertion Rate (CMR)

Cause (1): Problems on the physical layer (cabling, connectors, jitter, ATM port)
Cause (2): Too many nodes along the transmission path of a VP or VC connection
Cause (3): High network load
Cause (4): ATM switch malfunction

Symptom: High Mean Cell Transfer Delay (MCTD)

Cause (1): High signal delay due to long transmission path
Cause (2): Too many nodes along the transmission path of a VP or VC connection
Cause (3): ATM switch overloaded
Cause (4): Insufficient buffering in the switch
Cause (5): High network load
Cause (6): Limits of traffic contract are exceeded

Symptom: High Cell Delay Variation (CDV)

Cause (1): Too many nodes along the transmission path of a VP or VC connection
Cause (2): ATM switch overloaded
Cause (3): Insufficient buffering in the switch
Cause (4): High network load
Cause (5): Limits of traffic contract are exceeded

Symptom: No Connection over Emulated LAN (ELAN)

Cause (1): Problems with the connected traditional LANs (Ethernet, FDDI, Token Ring)
Cause (2): IP interfaces on the LAN emulation clients are not active or not functioning
Cause (3): IP addresses and subnet masks are incorrect; interfaces belong to different subnets
Cause (4): LANE software on the client is not active
Cause (5): The LE Clients trying to communicate do not belong to the same ELAN
Cause (6): The LE Clients are not registered on the same LE Server/BUS
Cause (7): The VCC and ATM address of the LANE server (LE Server) are incorrect
Cause (8): The VCC and ATM address of the BUS are incorrect
Cause (9): LANE-ARP entries are incorrect (MAC-ATM address resolution is not working)
Cause (10): The traffic contracts of the LE Clients are incompatible

Cause (11): The primary LANE service failed and the backup LANE service was not activated

Symptom: No Connection over PNNI Network

Cause (1): Signaling problems (SVC) between the end systems involved
Cause (2): Wrong route selection due to incorrect ATM addressing of the end systems
Cause (3): Topology information on the switch port is incomplete or outdated
Cause (4): Misconfigured peer group leader (PGL not active or no designated parent LGN)
Cause (5): Hello protocol is not active on the PNNI
Cause (6): The PNNI Routing Control Channel (SVCC-RCC) is inactive
Cause (7): Misconfigured PNNI port parameters (cell rate, cell transfer delay, bit rate)
Cause (8): Uplinks to neighboring peer groups are inactive or not defined
Cause (9): PNNI addresses, prefixes or summary addresses are incorrect

Symptom: Loss of ATM Connections

Cause (1): Violation of the traffic contract; traffic shaping activated
Cause (2): Cell streams with different priorities are transmitted at high load, and cells with low priority are discarded
Cause (3): Clocking and synchronization problems due to configuration errors on the ATM port
Cause (4): Problems on the physical layer (cabling, connectors, ATM port)

Gathering Information; Common Errors

The first step in any troubleshooting process is to gather information. In diagnosing ATM problems, comprehensive information about the context of the problem provides a detailed description of the symptoms and clues to possible causes. Questions to ask at this stage include:

- Do the symptoms occur regularly or intermittently?
- Are the symptoms related to certain applications, or do they affect all network operations?
- Do the symptoms correlate to other activities in the network?
- When was the first occurrence of the symptom?
- Was there any change in any hardware or software network component?
- Has anyone connected or disconnected a PC (laptop or desktop) or any other component to or from the network?
- Has anyone installed an interface card in a computer?
- Has anyone stepped on a cable?

- Has any maintenance work been performed in the building recently (by a telephone company or building maintenance personnel, for example)?
- Has anyone (including cleaning personnel) moved any equipment or furniture?

The following table lists the most common causes of problems in ATM networks:

- ATM interface card defective
- ATM interface card incorrectly configured (interrupt, driver, timers)
- Cell streams with different priorities are being transmitted at high load, and cells with low priority are discarded
- Classical IP: ATM ARP server address not configured on the client systems
- Classical IP: Misconfigured ATM ARP server: clients are not registered at all or registered under a wrong address
- Faulty cable infrastructure: see Chapter 6
- Electromagnetic interference (ATM over UTP)
- Hardware or software problems on the switch
- High signal transit delay due to long transmission path
- ILMI not active on the client or on the ATM switch
- Incompatible ILMI software versions on client and server
- Incorrect port configuration: bit rate, scrambling, interface type, frame type (PLCP, G.804, SDH, SONET)
- Incorrect router configuration (port inactive, wrong operating mode, protocol not active)
- Incorrect router filters
- Insufficient buffering in the switch
- LANE: IP addresses and subnet masks are incorrect; interfaces belong to different subnets
- LANE: IP interfaces on the LE clients are not active or not functioning
- LANE: LANE software on the client or switch is not active
- LANE: LANE-ARP entries are incorrect (MAC-ATM address resolution is not working)
- LANE: LE Clients are not registered on the same LE Server/BUS
- LANE: LE Clients trying to communicate do not belong to the same ELAN
- LANE: The primary LANE service failed and the backup LANE service was not activated
- LANE: The traffic contracts of the LE Clients are incompatible

Figure 10.73a The most frequent causes of trouble in ATM networks

- LANE: The LE server (LES) VCC is inactive or the ATM address of the LES is incorrect
- LANE: The BUS VCC is inactive or the ATM address of the BUS is incorrect
- Loose or defective connectors on interface cards, wall jacks, MAUs, hubs, bridges, routers
- Misconfigured ATM interface card (interrupts, drivers, timers)
- Misconfigured routing protocol entries (address tables, mapping tables, subnet masks, default gateways, routing tables, timers)
- PNNI: Hello protocol on the PNNI interface is not active
- PNNI: Misconfigured peer group leader (PGL is not active or no designated parent LGN)
- PNNI: Misconfigured PNNI port parameters (cell rate, cell transfer delay, bit rate)
- PNNI: PNNI addresses, prefixes or summary addresses are incorrect
- PNNI: The PNNI Routing Control Channel (SVCC-RCC) is inactive
- PNNI: Topology information on the switch port is incomplete or outdated
- PNNI: Uplinks to neighboring peer groups are inactive or not defined
- PNNI: Wrong route selection due to incorrect ATM addresses for the destination node
- PVC not set up; invalid VCI
- Faulty physical router or switch installation (cables, connectors, plug-in modules are loose, backplane connections are miswired)
- Problems on the physical layer (cabling, connectors, ATM port)
- Switch is overloaded
- Too many nodes along the transmission path of a VPI/VCI connection
- Traffic contract is exceeded, cells are being discarded

Figure 10.73b The most frequent causes of trouble in ATM networks

Index of chapter 10

A

G

H

I

L

M

N

O

P

Q

R

S

T

U

V