# Chapter 2

# B-ISDN and ATM

In order to support the diverse requirements of uncertain applications from various traffic sources, the broadband integrated services digital network (B-ISDN) was developed and standardized to provide flexible transport and switching services in an efficient and cost-effective way. There are three primary driving forces for the development of B-ISDN. One is the need for high-speed communications by means of optical fiber transmission systems. The advent of fiber optics has made it possible to achieve drastic improvements of transmission capacity at low error rates and low cost. It is an effective way to service the emerging user applications, such as image and video services and bulk data transfer, which require high-speed communication. Processor technologies to support the services in an integrated manner is another driving force. Due to the advances in processor technologies, it is possible to put more functions on a single chip operating at a higher quality, a higher processing speed, and a higher transmission speed. Because of this, B-ISDN has become available economically. The final driving force is the establishment of digital transmission hierarchy. The synchronous digital hierarchy (or SDH) has been standardized as a broadband physical interface which enables flexible multiplexing of various rates of information streams. As a developed form of narrowband ISDN, B-ISDN should include broadband and multimedia communication capability, intelligent network capability to provide enhanced network services, and sophisticated functional capability to provide reliable services.

## 2.1 B-ISDN architecture

Although the B-ISDN is from a narrowband ISDN, it is quite different from a narrowband ISDN. The requirements in terms of the data rates of the user application are the major discrepancy. In order to support widespread ranges of data rates, optical fiber has been thought to be the most appropriate transmission medium. Hence, the introduction of B-ISDN keeps pace with the introduction of fiber subscriber loops.

Internal to the network, the switching technique should be capable of handling a wide spectrum of bit rates and traffic parameters. Despite the advancement of processor and optical fiber technology, it is exceedingly difficult to handle the diverse requirements of B-ISDN with circuit-switching technology. Thus, *fast packet switching* has been paid much attention to as the basic switching technique for B-ISDN [PRYC93, p.58]. This switching technique, also known as ATM (or **asynchronous transfer mode**), supports the user-network interface protocol, which is detailed later.

Table 2.1: Principles of B-ISDN

**ATM** Asynchronous transfer mode (ATM) is the transfer mode for implementing B-ISDN and is independent of the means of transport at the physical layer.

**Type of connection** B-ISDN supports switched, semi-permanent, and permanent point-to-point and point-to-multipoint connections, and provides on demand reserved and permanent services. Connections in B-ISDN support both circuit-mode and packet-mode services of a mono- and/or multimedia type and of a connectionless or connection-oriented nature and in a bidirectional or undirectional configuration.

**B-ISDN architecture** The B-ISDN architecture is detailed in functional terms and is therefore, technology- and implementation-independent.

**Intelligent capabilities** A B-ISDN will contain intelligent capabilities for the purpose of providing advanced service characteristics and supporting powerful operation and maintenance tools, network control, and management. Further inclusion of additional intelligent features has to be considered in an overall context and may be allocated to different network/terminal elements.

**Conceptual basis** Since the B-ISDN is based on overall ISDN concepts, the ISDN access reference configuration is also the basis for the B-ISDN access reference configuration.

**A layered structure** A layered structure approach, as used in established ISDN protocols, is also appropriate for similar studies in B-ISDN, including information transfer, control, intelligence, and management.

**Expression** Any expression of network capabilities or change in network performance parameters will not degrade the quality of service of existing services.

**Support** The evolution of B-ISDN should ensure the continued support of existing interfaces and services.

**Evolution** New network capabilities will be incorporated into B-ISDN in evolutionary steps to meet new user requirements and accommodate advances in network developments and progress in technology.

**Nation specific implementation** It is recognized that B-ISDN may be implemented in a variety of ways according to specific national situations.

Table 2.1 is an excerpt from [STAL95, p.430] describing ITU-T Recommendation I.121, which states the basic principles of broadband aspects of integrated services digital network

(B-ISDN) and indicates further evolution from TDM (time division multiplexing) and ISDN (integrated services digital network) capabilities in order to support more advanced services and applications. A summary of the main consensus of the *B-ISDN Recommendations* is as follows [STAL95, p.410][HAND94, p.15]:

- There is an emerging demand for broadband services.
- The availability of high-speed transmission, switching, and signal processing technologies is increasing.
- The data and image processing capabilities are improved to be available to the user.
- The software application processing in the computer and telecommunication industries is advancing.
- There is a growing need to integrate interactive and distribution services.
- There is a need to integrate circuit and packet transfer mode into a universal broadband network.
- It is needed to provide flexibility in satisfying the requirements of both user and operator.
- There is a need to cover broadband aspects of ISDN in ITU-T Recommendations.

Figure 2.1 depicts the functional architecture of B-ISDN[HAND94, p.22][MCDY94, p.192]. It is the vision of how B-ISDN should interconnect with ISDN, signaling system number 7 (SS7) and OSI (open systems interconnection). SS7, ISDN and B-ISDN are all lower level capabilities that serve to interconnect terminal equipment (TE) or service providers through local functional capabilities which represent the point of physical interface. Within the network, SS7 provides enhanced out-of-band inter-exchange signaling capabilities for telephony and ISDN, while ISDN provides signaling capabilities for TDM-based services, X.25, and frame relay. Similarly, the user-network control signaling protocol is an enhanced version of Q.931[1]. B-ISDN must support all of the 64 Kbit/s transmission services which are supported by narrowband ISDN. In addition, broadband capabilities are provided for higher-data-rate transmission services. At the user-network interface, these capabilities will be provided with the connection-oriented asynchronous transfer mode (ATM) facility. These all act in support of the higher layer capabilities above layer three in the OSI model.

A reference configuration is a practical tool to define clear interfaces and functions between different entities of the network. Even though ITU-T defines reference configurations for the user-network interface for ISDN, the reference configurations are considered to be applicable to all aspects of the B-ISDN accesses. Figure 2.2 shows the reference configuration for B-ISDN. The

[1] Q.931 is a standardized signaling protocol by ITU-T.

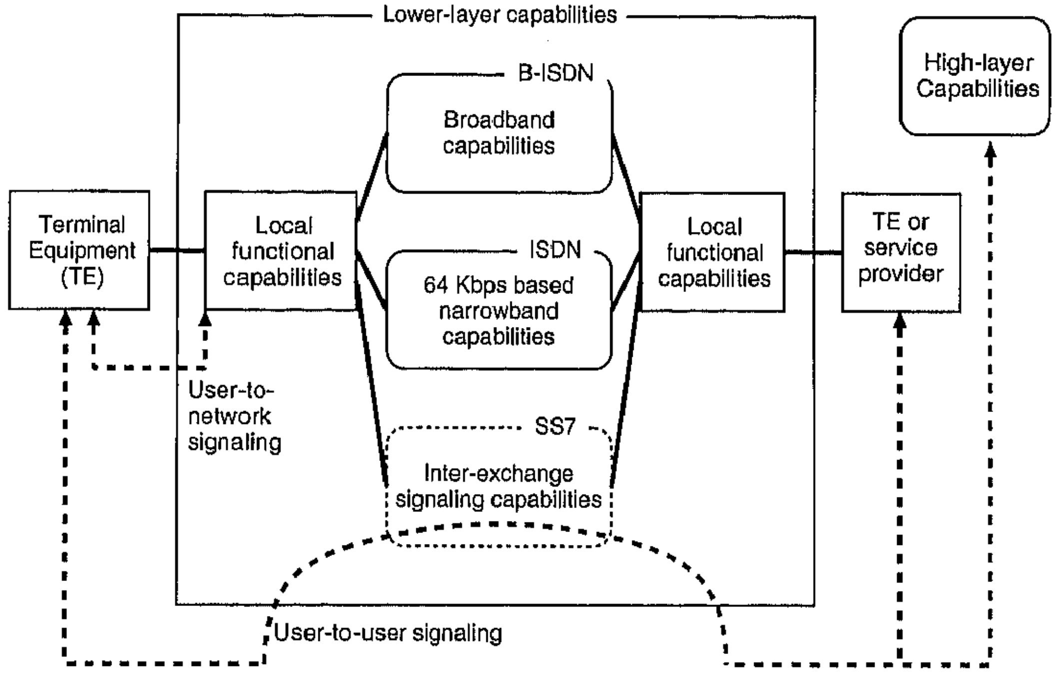


B-ISDN : Broadband ISDN
ISDN : Integrated services digital network
SS7 : Signaling system number 7

Figure 2.1: B-ISDN architecture

reference points $R$, $S$, $T$ and $U$ as defined for narrowband ISDN are also valid in the B-ISDN. Interfaces at the $R$ reference point may or may not have broadband capabilities. In order to clearly illustrate the broadband aspects, the notations for reference points and functional groupings are appended with the letter $B$ (e.g., B-NT1, $T_B$). The B-NT1 mainly performs low layer functions such as line transmission termination, transmission interface handling and OAM (operations, administrations and maintenance) functions. The B-NT2 functional group performs adaptation functions for different media and topologies. Moreover, it handles other higher layer functions, such as cell delineation, concentration, buffering, multiplexing/demultiplexing, resource allocation, usage parameter control (policing), signaling adaptation layer functions, signaling protocol handling, switching of local connections, and OAM functions. The B-TE1 terminates the user interface ($S_B$ or $T_B$), and performs the termination of all end protocols from the low layers up to the higher layers.

Physically, the reference configuration as described can be implemented in several ways. Figure 2.3 shows a few examples for the physical interfaces configurations. In the first example (A), both the $S_B$ and $T_B$ interface are physically connected to interface a physical B-NT2 entity. Physical interfaces, however, are not always needed to occur at any reference point. As in (B), only the $S_B$ is physically presented, in which the B-NT2 and B-NT1 functions might be physically co-located in a single entity. It is also possible for the terminal to include B-NT2

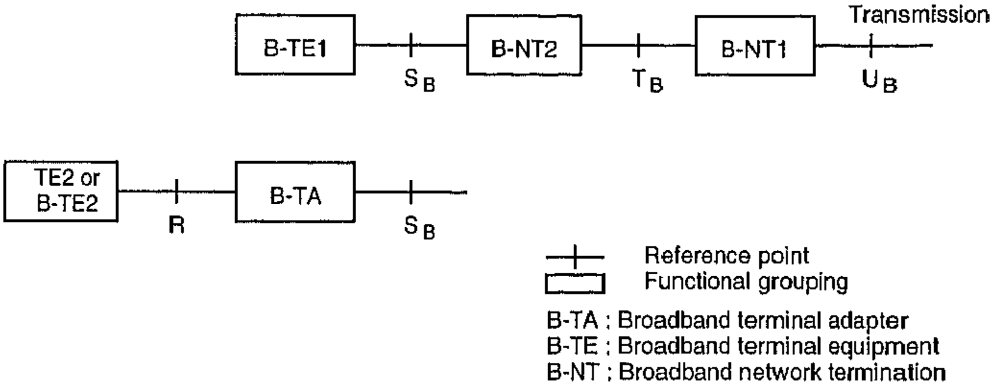


Figure 2.2: B-ISDN reference configuration

functionality. In (C), only $T_B$ is physically implemented resulting in a combined B-TE and B-NT2 entity. If the same interface standard applies to both $S_B$ and $T_B$, these reference points are identical. Thus as in (D), no physical B-NT2 is present at the customer's premises and a terminal is directly connected to the B-NT1. B-ISDN will also offer 64 Kbit/s ISDN services and interfaces. In (E), the physical interfaces between terminals and B-NT2 are both $S_B$ and $S$ (for narrowband ISDN), based on a concentrated B-NT2 grouping. The provision of multiple terminal interfaces by the B-NT2 is not restricted to a specific topology. Other configurations and topologies are also possible (Figure 2.4). B-ISDN has to support various requirements to those different configurations shown in Figures 2.3 and 2.4.

Figure 2.5 is a general depiction of the B-ISDN architecture for user-access. The local exchange to attach subscribers must be able to handle both B-ISDN and ISDN subscribers. ISDN subscribers can be supported with twisted pair (copper) at the basic (144 Kbit/s) and primary access rates (1.544/2.048 Mbit/s), depending on the number of channels and the size of each channel. For B-ISDN subscribers, optical fiber will be used. The data rate from the network to subscribers will be on the order of 600 Mbit/s in order to service multiple video distributions, as might be required in an office environment. The data rate from subscriber to network would normally be much less, since the typical subscriber does not initiate distribution services. A rate of about 150 Mbit/s or less is probably adequate.

In terms of data rates available to B-ISDN subscribers, three transmission services are defined; full-duplex 155.52 Mbit/s service, asymmetric data rate service, and full-duplex 622.08 Mbit/s services. The asymmetrical service provides transmission from the subscriber to the network at 155.52 Mbit/s and in the other direction at 622.08 Mbit/s. A transmission capacity of 155.52 Mbit/s can certainly support one or more basic or primary-rate interfaces to provide all of the narrowband ISDN services. Moreover, it can also support most of the B-ISDN services. At this data rate, several video channels can be supported depending on the video resolution and the coding technique used. Thus, the full-duplex 155.52 Mbit/s service will probably be

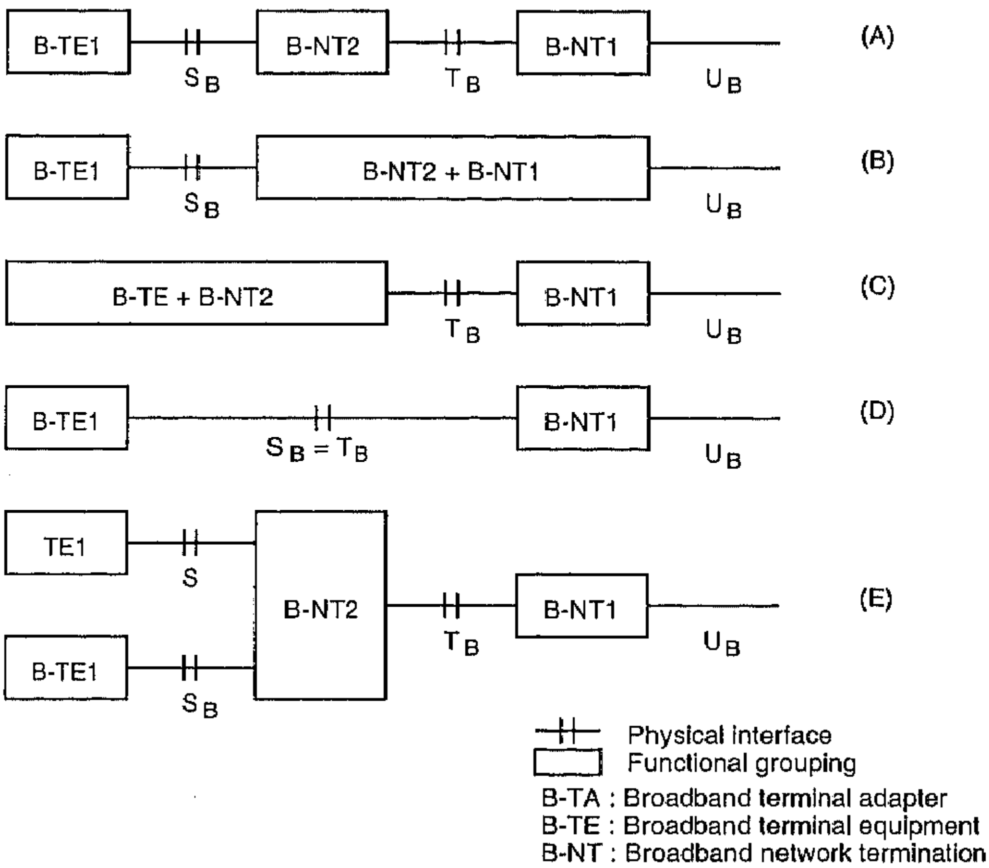


Figure 2.3: Examples of physical configurations

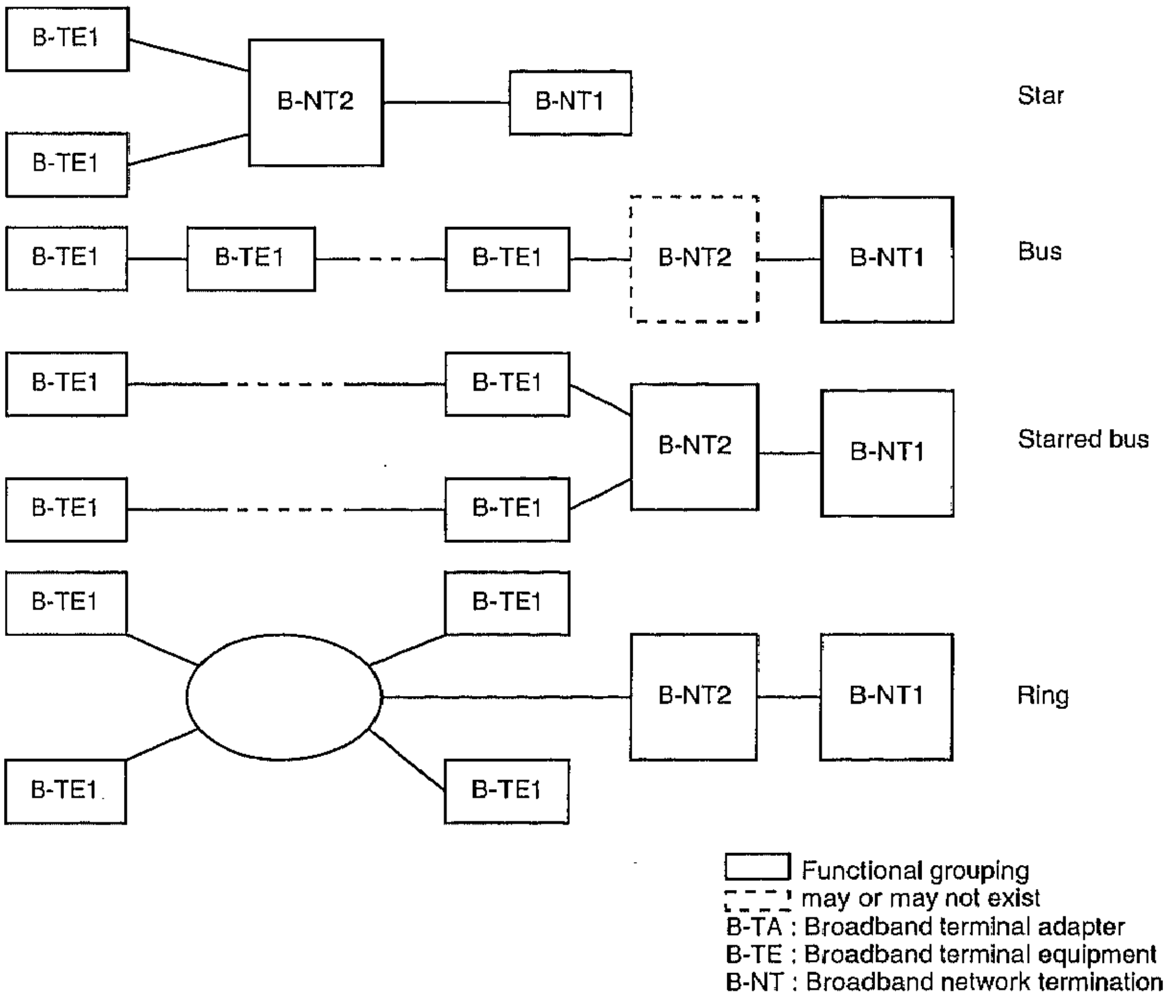


Figure 2.4: Examples of multiple interface arrangements

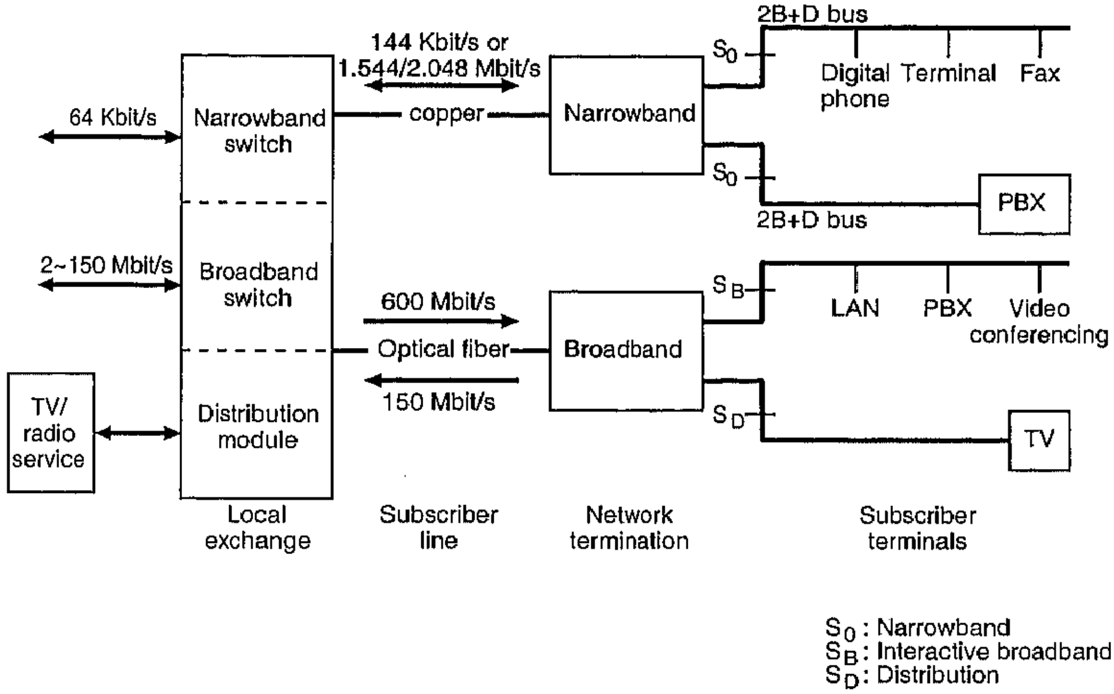


Figure 2.5: B-ISDN user-network interface

the most common B-ISDN service. On the other hand, the higher data rate of 622.08 Mbit/s is needed to handle multiple video distribution, such as multiple simultaneous video-conferences. This data rate makes sense in the network-to-subscriber direction. The typical subscriber will not initiate distribution services and thus it would still be enough for a user to use the lower 155.52 Mbit/s service. The full-duplex 622.08 Mbit/s service would be appropriate for communications between video distribution providers. The ITU-T excludes the specific channel data rates to allow the user and the network to negotiate any channel capacity that can fit in the available capacity provided by the network. Thus, B-ISDN becomes considerably more flexible and can be tailored precisely to a wide variety of applications.

## 2.2 B-ISDN protocols

This section introduces the B-ISDN protocol reference model. For B-ISDN, the transfer of information across the user-network interface uses what is called "asynchronous transfer mode" (ATM). The ATM mechanism is embedded into a protocol reference model that defines the B-ISDN user-network interface. Before describing the B-ISDN protocol, let the author first introduce the basics of ATM.

### 2.2.1 Basics of ATM

ATM is a specific technology of switching and multiplexing information in a network. In order to understand what ATM is, a brief introduction of **synchronous transfer mode** (STM) is in order. STM is a circuit switched networking mechanism used by telecommunication backbone networks to transfer packetized voice and data across long distances. In the STM, a connection is set up between two end points before data transfer begins, and torn down when the communication is over. Thus the bandwidth is allocated and reserved by two end points during the connection setup duration even though they have no data to transmit. In an STM network, the bandwidth is divided into fundamental units of transmission called "time-slots" or "buckets". These transmission units are organized into a frame containing a fixed number of time-slots labeled from 1 to $N$. The frame repeats periodically every $T$ time period, with the buckets in the train always in the same position with the same label. In the STM, a data unit associated with a given channel is identified by its position in the transmission frame as in Figure 2.6. Thus, if a

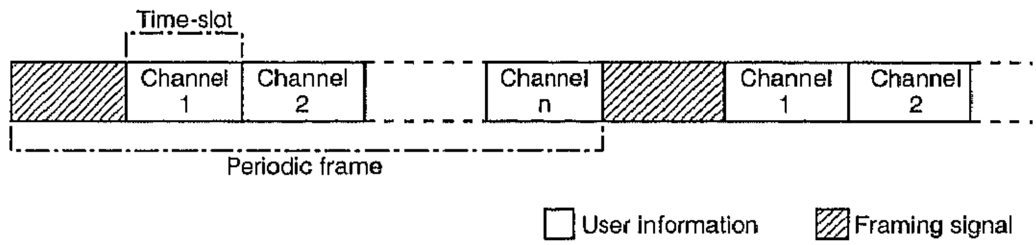


Figure 2.6: STM principle

channel is not transmitting data, the time slot remains reserved and is still transmitted, without any useful payload. In this case, if the other channels have more data to transmit, they have to wait until their reserved, assigned time slot occurs again. Frequent empty time slots result in low line utilization. Even worse, the flexibility of bit rate allocation to a connection in STM is limited by its use of predefined channel bit rates[2]. The conventional transmission frames have a rigid structure, which will not permit individual structuring of the payload or will only permit a quite limited selection of channel mixes. Otherwise the network provider would have to manage a host of different interface types, which should be avoided.

In contrast, ATM uses a completely different approach. The main idea is to carry the connection identifier along with the data in small-sized time-slot, called a "cell" to enable flexible bit rate allocation and easy recovery of dropped buckets. Thus two end points in an ATM network communicate with each other by means of a cell over a connection called a "**virtual channel**" (VC). The identifier of a VC, a **virtual channel identifier** (VCI), is carried in the header of a cell, whose format is shown in Figure 2.7. In ATM, a primary unit of transferred data is the cell. An ATM cell is defined as a fixed-size data unit with a length of 53 octets comprised

[2]For example, the bit rate of B channel is 64 Kbit/s, and those of H channels are multiples of 384 Kbit/s

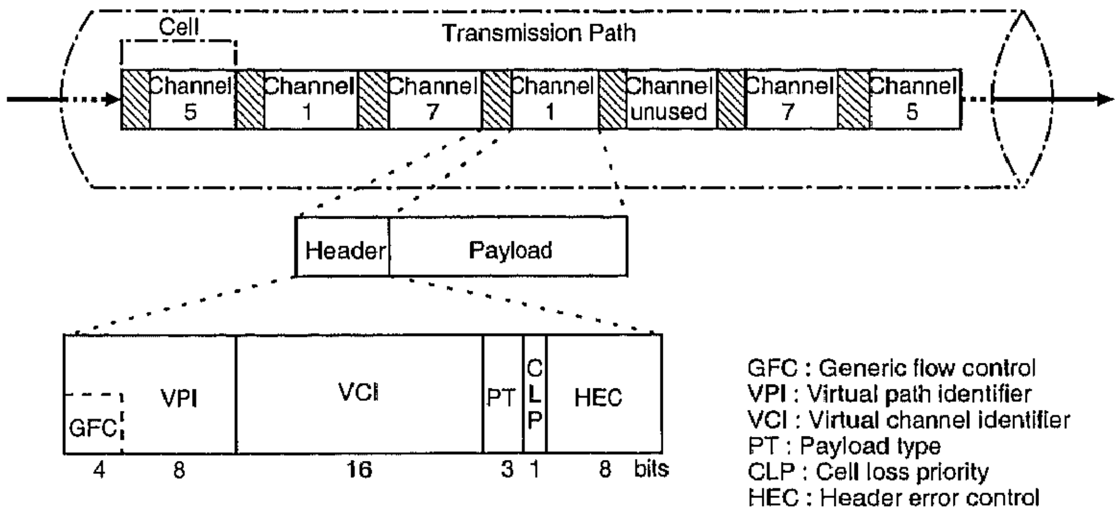


Figure 2.7: ATM principle and cell format

of a five-octet header and a 48-octet payload. Whereas the information field is available for the user, the header field carries information that pertains to the ATM layer functionality, such as destination, cell type and priority. Each entry of the header is briefly described here. The destination is identified by the **virtual path identifier** (VPI) and **virtual channel identifier** (VCI). The **generic flow control** (GFC) field allows a multiplexer to control the rate of an ATM terminal. The **payload type indicator** (PTI) field indicates whether the cell contains user data, signaling data, or maintenance information. Two priority levels, indicated by the **cell loss priority** (CLP) bit, are implemented in ATM layer. The lower priority cells are discarded before the higher priority cells during congested intervals. The cell **header error control** (HEC) detects and corrects errors in the header. It is part of the cell header but it is not used by the ATM layer. Instead, it is processed by the physical layer. The term *asynchronous* in the ATM refers to the fact that, in the context of multiplexed transmission, cells allocated to the same connection may exhibit an irregular recurrence pattern as they are filled according to the actual demand. This is shown in upper part of Figure 2.7.

In ATM-based networks, ATM switches take data, voice, and video from users and chop them into cells and multiplex and/or switch them into a single bit stream which is transmitted across a physical medium (Figure 2.8). Thus, the same piece of equipment can, in principle, handle a low bit rate connection (data) as well as a high bit rate connection (video), be it of stream (voice and video) or bursty nature (data). The flexibility of the ATM network supports the idea of a unique interface which can be employed by a variety of customers with quite different service needs. In addition, it is possible to allocate bandwidth dynamically with a fine granularity in the ATM network.

With regard to a standard cell size, there was a raging discussion between a 32-octet versus a 64-octet payload size. The decision of the 48-octet payload size was a compromise between

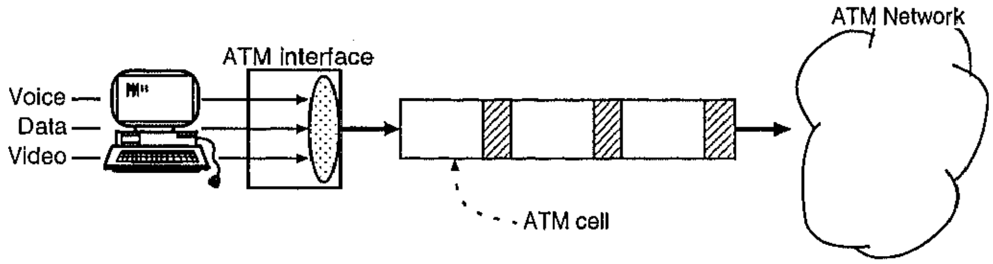


Figure 2.8: Multimedia communication over ATM

the two. The choice of the five-octet header size was a separate tradeoff between a three-octet header and an eight-octet header. There is a basic tradeoff between efficiency and packetization

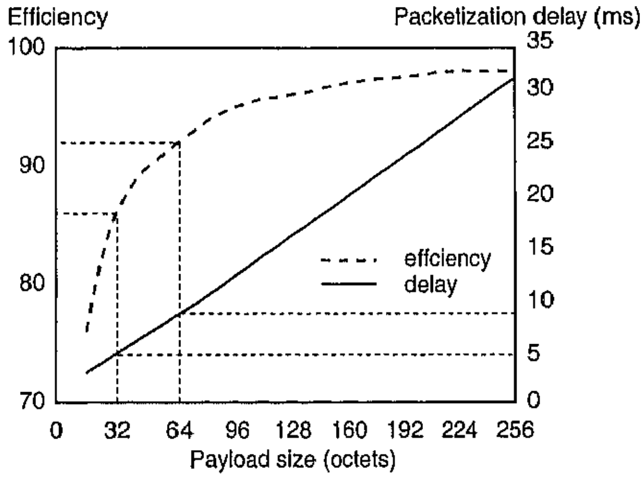


Figure 2.9: Processing delay versus cell size

delay versus cell size as illustrated in Figure 2.9. Efficiency is determined by the ratio between the payload size and a whole packet size. For example, when the payload size is $P$ and the header size is $H$, the efficiency is $P/(P+H)$. Here, the header size $H$ is five octets. Packetization delay is the amount of time required to fill the cell at a rate of 64 kbit/s; that is, the rate to fill the cell with digitized voice samples. For example, when the payload size is 128 octets, the packetization delay is $128 octets \times 8/64[Kbit/s] = 16\ ms$. Ideally, high efficiency and low delay are both desirable, but in reality cannot be achieved simultaneously. As in the figure, efficiency is high at large cell sizes but packetization delay is also increased. In order to carry voice over ATM and interwork with two-wire analog telephone sets, echo cancellation is generally necessary if the the total delay is large. ITU-T Recommendation Q.161 states that a delay of 24 ms is still acceptable without echo cancellers. In case of the longer delays, echo cancellers are required. For a typical national connection, the overall end-to-end delay can be kept within limits if the packetization delay remains around 4 ms (or the cell size around 32 octets). At 64 octets, and a few conversions between ATM and non-ATM networks, this 24 ms is quite short time, requiring frequent echo cancellation. Thus, the ITU-T adopted the fixed-length 48-octet cell payload as a compromise between long cell sizes for time-insensitive traffic (64 octets) and smaller cell sizes for time-sensitive traffic (32 octets).

To sum up, whereas today's networks are characterized by the coexistence of circuit switching and packet switching, B-ISDN will rely on a single switching or multiplexing method, ATM, which combines the advantages of both circuit-oriented and packet-oriented techniques. The former provides low overhead and processing. In addition, once a connection is established, the transfer delay of the information will be low and constant. The latter is much more flexible than the former in terms of the bit rate assigned to each connection. ATM is a simplified, circuit-oriented, hardware-controlled and low-overhead concept technology manipulating cell level transfer over a virtual channel, which has no flow control or error recovery (unlike X.25). The virtual channel communication is implemented by fixed-size small cells and provides the basis for both switched and multiplexed transmission. The use of short cells in ATM and high transfer rates (e.g., 155.52 Mbit/s) result in small transfer delays and delay variations which enable it to be applied to a wide range of services, including real-time services, such as voice and video.

### 2.2.2 B-ISDN protocol reference model

Figure 2.10 depicts the B-ISDN protocol reference model from ITU-T Recommendation I.321, which is used to structure the remaining recommendations [MCDY94]. A significant portion of

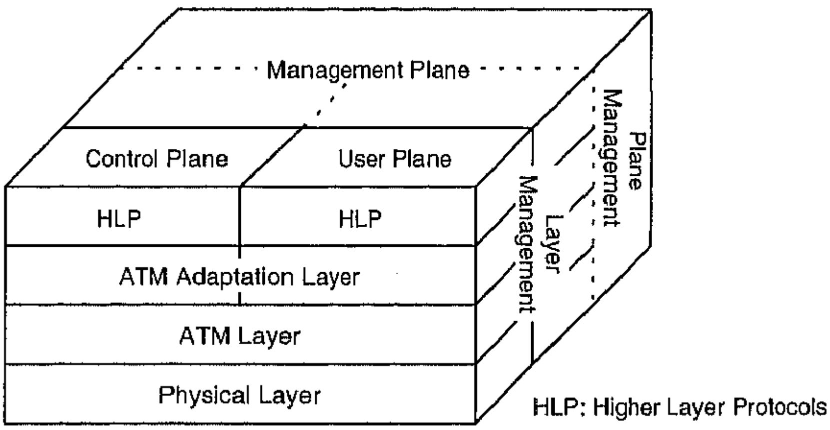


Figure 2.10: B-ISDN protocol reference model

the architecture exists in the protocol reference cube.

It is assumed that the transfer of information across the user-network interface uses asynchronous transfer mode (ATM). ATM is, in essence, a form of packet transmission across the user-network interface with a fixed sized packet, referred to as a cell. In addition, as is characteristic of ATM, ATM makes use of common-channel signaling. The decision to use ATM for

B-ISDN is a remarkable one. This implies that B-ISDN is a packet-based network, certainly at the interface and almost certainly in terms of its internal switching. Although the ITU-T Recommendation also states that B-ISDN will support circuit-mode applications, this is done over a packet-based transport mechanism [STAL95]. Thus, ISDN, which began as an evolution from the circuit-switching telephone networks, has transformed itself into a packet-switching network as it takes on broadband services.

The top of the cube illustrates the planes, which are defined on the front and side of the cube. The **user plane** provides for the transfer of user information, along with associated controls, such as flow control and error control. The **control plane**, mainly composed of signaling information, performs call control and connection control functions. These user and control planes span the higher layer, down through the **ATM adaptation layer** (AAL), all the way to the ATM layer and physical layer. Therefore, the physical layer, the ATM layer, and the AAL form the foundation of B-ISDN. Though user and control planes may make use of common ATM and physical layer protocols, the end purpose differs in the AALs and the higher layers. The **management plane** is further broken down into **layer management**, which performs management functions relating to resources and parameters residing in its protocol entities, and **plane management** which performs management functions related to a system as a whole and provides coordination between all the planes. As shown in the Figure 2.10, layer management has a layered structure and interfaces with each layer in the control and user planes. Each of its layers handles the specific operations and maintenance (OAM) information flows for the corresponding layers. In contrast, plane management has no layered structure. It can be viewed as a catch-all for the things that do not fit into the other portions of this model, by having the role of overall system management. Two layers of the B-ISDN protocol architecture relate to ATM functions. There is an ATM layer common to all services, which provides packet transfer capabilities, and an ATM adaptation layer, which is serviced-dependent. The AAL maps higher-layer information into ATM cells to be transported over a B-ISDN, then collects information from ATM cells for delivery to the higher layers.

As the congestion control or flow control schemes are performed in the user plane, this author will concentrate on detailing the user plane in the next few sections.

### 2.2.3 Layer functions

The functions of the physical layer, the ATM layer and the ATM adaptation layer are described in the following sections. In Table 2.2, the lower layers of the B-ISDN protocol reference model and their functions are illustrated.

Table 2.2: The functions of B-ISDN in relation to the B-ISDN protocol reference model

| | Higher-Layer Functions | Higher Layers | |
|---|---|---|---|
| Layer Management | Convergence | CS | AAL |
| | Segmentation and reassembly | SAR | |
| | Generic flow control<br>Cell header generation/extraction<br>Cell VPI/VCI translation<br>Cell multiplexer and demultiplexer | ATM | |
| | Cell rate decoupling<br>HEC header sequence generation/verification<br>Cell delineation<br>Transmission frame adaptation<br>Transmission frame generation/recovery | TC | Physical Layer |
| | Bit timing<br>Physical medium | PM | |

CS: Convergence sublayer
SAR: Segmentation and reassembly sublayer
AAL: ATM adaptation layer
ATM: Asynchronous transfer mode layer
TC: Transmission convergence sublayer
PM: Physical medium sublayer

### Physical layer

The physical layer of the B-ISDN is further subdivided into the **physical medium** (PM) sublayer and the **transmission convergence** (TC) sublayer. The PM layer is the lowest sublayer and includes only the pure bit functions dependent upon physical medium. It provides the bit transmission capability, including bit alignment. Line coding and, if necessary, electrical/optical conversion is also performed by this sublayer. In many cases, the physical medium will be an optical fiber. Other media, such as coaxial and twisted pair cables, are also possible. Therefore, the functions of the *physical medium* are medium specific. *Bit timing* functions in the Table 2.2 are the generation and reception of waveforms which are suitable for the medium, insertion and extraction of bit timing information, and line coding if required.

In the transmission convergence (TC) layer, bits are already recognized as they come from the PM sublayer. This sublayer performs basically five functions as shown in Table 2.2. The lowest function is the *generation and recovery of the transmission frame*. Transmission at the physical layer consists of frames as in basic- and primary-rate interfaces. This function is concerned with generating and maintaining the frame structure appropriate for a given data rate. Information exchange at the ATM layer is a flow of ATM cells. The TC sublayer is responsible for all actions necessary to adapt the cell flow according to the used payload structure of the transmission system in the sending direction. In the opposite direction it extracts the cell flow from the transmission frame. This function is *transmission frame adaptation* function. The transmission frame may be a bare cell stream (i.e., no frame structure is used), a synchronous digital hierarchy (SDH) envelope or an envelope according to other recommendations. In the case of the B-ISDN UNI [3], ITU-T proposes an SDH envelope or no frame structure. These alternatives are described in more detail in later (Section 2.3.1). The functions mentioned so far are specific to the transmission frame. All other TC sublayer functions, which are presented in the following, can be common to all possible transmission frames. *Cell delineation* function is the mechanism that enables the receiver to recover the cell boundaries. This mechanism is described in ITU-T Recommendation [HAND94]. The detailed procedure of the cell delineation mechanism is given in Section 2.3.5. To protect the cell delineation mechanism from malicious attack, the information field of a cell is scrambled before transmission. Then, descrambling is performed at the destination. Each cell header is protected by a *header error control* (HEC) code. HEC sequence generation is done in the transmit direction. It is inserted in its appropriate field within the header. At the receiving side, the HEC value is recalculated and compared with the received value. The header errors may be corrected, otherwise the cell is discarded. The HEC mechanism is detailed in Section 2.3.5. In the sending direction, the *cell rate decoupling* mechanism inserts *idle cells* in order to adapt the rate of valid ATM cells to the payload capacity of the transmission system. In the receiving direction, this mechanism ignores all idle cells and only assigned and unassigned cells are passed to the ATM layer.

ATM layer

The ATM layer is located above the physical layer. Its features are fully independent of the physical medium used to transport the ATM cells and thus of the physical layer. Mainly, the following functions are performed in this layer. In the transmit direction, cells from multiple logical connections are multiplexed into a single cell stream on a physical layer by the *cell multiplexing* function. The composite stream is normally a non-continuous cell flow. At the receiving side, the *cell demultiplexing* function splits the arriving cell stream into individual cell flows appropriate to the VP or VC. The VPI and VCI relate to logical connections and have local significance. These *VPI and VCI translation*, which is required in most cases when switching a cell from one physical link to another, are performed at ATM switching nodes and/or at cross-connect nodes. This translation can be performed either on the VCI or VPI separately, or on both simultaneously. Within a VP switch, the value of the VPI field of each incoming cell is translated into a new VPI value for the outgoing cell. The values of the VPI and VCI are translated into new values at a VC switch. The *cell header generation/extraction* function is applied at the termination points of the ATM layer. In the transmit direction, the

[3]UNI is the interface between the user and the network and will be described in Section 2.4.2.

appropriate cell header (except for the HEC value) is generated and added to the cell information field received from the AAL. VPI/VCI values could be obtained by translation from the SAP (service access point) identifier. In the opposite direction, the cell header is extracted by cell header extraction function. Then, only the cell information field is passed to the AAL. This function could also translate a VPI/VCI into an SAP identifier. The *generic flow control (GFC)* function is only defined at the B-ISDN UNI[4]. GFC supports control of the ATM traffic flow in a customer network. It can be used to alleviate short-term overload conditions (user-to-network direction) at the UNI. GFC information is carried in assigned or unassigned cells.

### ATM adaptation layer

The ATM adaptation layer lies between the ATM layer and the higher layers. Its basic function is to enhance the adaptation of services provided by the ATM layer to a certain level required by the next higher layer. The higher layer PDUs are mapped into the information field of an ATM cell. AAL entities exchange information with their peer AAL entities to support AAL functions. AAL also performs functions for the user, control, and management planes. Specifically, the AAL for the control plane is called the **signaling AAL (SAAL)**. The functions performed in the AAL depend on the higher layer requirements. The AAL is subdivided into the **segmentation and reassembly** (SAR) sublayer and the **convergence sublayer** (CS). The essential functions of the SAR sublayer are, at the transmitting side, segmentation of the higher layer PDUs into a size suitable for the information field of the ATM cells of a virtual connection, and at the receiving side, reassembly of contents of the cells of a virtual connection into data units to be delivered to the higher layer PDUs. The CS is service dependent and provides the AAL service at the AAL service access point (AAL-SAP; Figure 2.11). It performs functions such as message

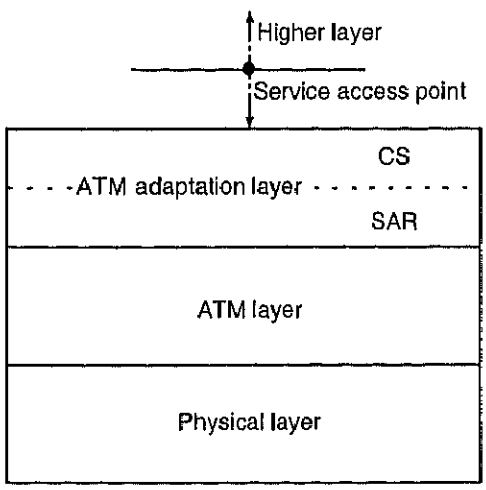


Figure 2.11: Service access point

[4]It is not defined at the NNI (network to network interface)

identification and time/clock recovery. Supporting data transport over the ATM layer, the CS is further subdivided into a **common part convergence sublayer** (CPCS) and a **service specific convergence sublayer** (SSCS) for some AAL types. AAL service data units (SDU) are transported from one AAL service access point (SAP) to one or more other SAPs through the ATM network. The AAL users will have the capability to select a given AAL-SAP associated with the quality-of-service required to transport the AAL-SDU. Several AALs have been defined by now depending on class of services.

The detailed account of each layer will be given later.

### Operation and maintenance

The **operation and maintenance** (OAM) actions are all technical and administrative actions for supervision of the network. With regard to these OAM actions, five levels of connectivity in the ATM transport network have been defined. The ATM transport network comprises the physical layer and the ATM layer. The physical layer consists of three levels: the regenerator section, the digital section, and the transmission path. The ATM layer is composed of the virtual path and the virtual channel. Each of these levels has its own bi-directional operation and maintenance flow, denominated "F1" to "F5", starting with F1 on the regenerator section level.

The transfer mode depends on the nature of the layer. For a physical layer based on SDH or PDH, flows F1 to F3 are carried in synchronous channels in the overhead of the physical layer. In contrast, for a cell based physical layer, the flows are carried by physical layer OAM (PL-OAM) cells. For the ATM layer, the F4 flows are carried by pre-assigned VCIs in the virtual path, and the F5 flows are carried with special PTI codes in the virtual channel.

## 2.3 Physical layer

The B-ISDN physical layer is specified in ITU-T Recommendation. Three options concerning the *physical medium* characteristics of the physical layer at the $T_B$ interface are provided in the standard:

- Full duplex at 155.52 Mbit/s in each direction.
- Subscriber to network at 155.52 Mbit/s and network to subscriber at 622.08 Mbit/s.
- Full duplex at 622.08 Mbit/s.

In addition, three kinds of transmission frame adaptation are specified by ITU-T; the SDH based option, the PDH based option, and the cell based option. The ATM Forum adds a fourth one, the FDDI based option.

### 2.3.1 SONET/SDH

**Synchronous optical network** or **SONET** is the most popular optical transmission interface originally proposed by Bell Core and standardized by ANSI. A compatible version, called **synchronous digital hierarchy (SDH)**, has been published by ITU-T [STAL95]. SONET tries to provide a specification for taking advantage of the high-speed digital transmission capability of optical fiber.

#### Signal hierarchy

The SONET specification defines a hierarchy of standardized digital data rates as shown in Table 2.3. The lowest level digital data rate, referred to as STS-1 (synchronous transport signal level

Table 2.3: SONET/SDH signal hierarchy

| SONET Designation | ITU-T Designation | Data Rate (Mbit/s) |
|---|---|---|
| STS-1 | | 51.84 |
| STS-3 | STM-1 | 155.52 |
| STS-9 | STM-3 | 466.56 |
| STS-12 | STM-4 | 622.08 |
| STS-18 | STM-6 | 933.12 |
| STS-24 | STM-8 | 1244.16 |
| STS-36 | STM-12 | 1866.24 |
| STS-48 | STM-16 | 2488.32 |

STS: Synchronous transport signal
STM: Synchronous transport module

1), is 51.84 Mbit/s. This rate suits to carry a single DS3[5] signal or a group of lower-rate signals, such as DS1[6], DS1C[7], DS2[8], plus ITU-T rates (e.g., 2.048 Mbit/s). Multiple STS-1 signals can be combined to form an STS-N signal. The STS-N signal is generated by interleaving bytes from N mutually synchronized STS-1 signals. The lowest rate for the ITU-T synchronous digital hierarchy is 155.52 Mbit/s, which is designated STM-1. This corresponds to SONET STS-3.

#### System hierarchy

SONET capabilities have been mapped into a four-layer logical hierarchy shown in Figure 2.12.

[5] 44.736 Mbit/s
[6] 1.544 Mbit/s
[7] $2 \times DS1$
[8] $2 \times DS1C$

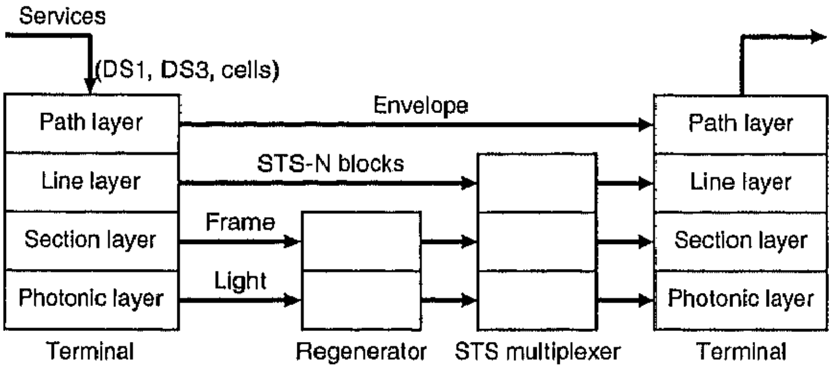


Figure 2.12: SONET system hierarchy

Photonic This is the physical layer. It includes a specification of the type of optical fiber that may be used and details such as the required minimum powers and dispersion characteristics of the transmitting lasers and the required sensitivity of the receivers.

Section This layer creates the basic SONET frames, converts electronic signals to photonic ones, and has some monitoring capabilities.

Line This layer is responsible for synchronization, multiplexing of data onto the SONET frames, protection and maintenance functions, and switching.

Path This layer is responsible for end-to-end transport of data at the appropriate signaling speed.

### Transmission convergence characteristics of SONET/SDH

The basic SONET building block is the STS-1 frame in Figure 2.13(a). It consists of 810 octets and is transmitted every 125 $\mu$sec, thereby an overall data rate of 51.84 Mbit/s ($810octets \times 8bits \times 1/125\mu sec$). The frame can logically be viewed as a matrix of 9 rows of 90 octets each. It is being transmitted one row at a time, from left to right and from top to bottom. The first three columns ($3octets \times 9rows = 27octets$) of the frame are devoted to line overhead. Nine octets from the first three rows are devoted to section-related overhead, and 18 octets from the rest six rows are devoted to line overhead. Figure 2.14(a) shows the fields of section overhead, and Table 2.4 defines these various fields. The remainder of the frame is the payload, which is provided by the path layer. The payload specifically includes one column of path overhead, which is not necessarily in the first column. The line overhead contains a pointer which indicates the starting point of the path overhead. Figure 2.14(b) shows the arrangement of path overhead octets, and Table 2.4 also includes definition of them. Based on the STS-1 frame format shown in Figure 2.13(a), ITU-T designated the general format for higher-rate frames as in Figure 2.13(b). In this option, the bit rate available for user information cells, signaling cells and OAM cells excluding

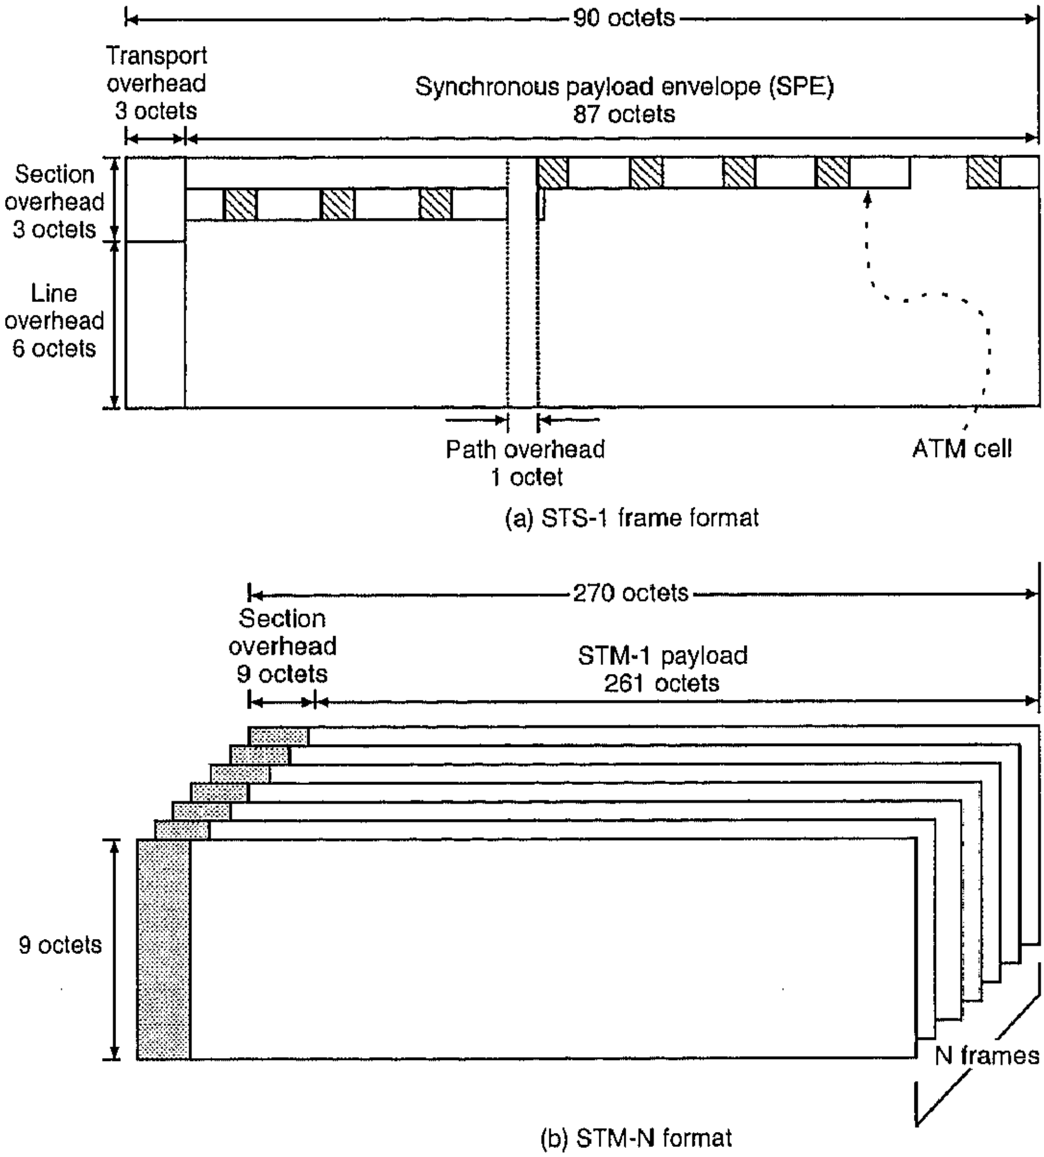


Figure 2.13: SONET/SDH frame formats

Table 2.4: STS-1 overhead bits

| | |
|---|---|
| | **Section overhead** |
| A1, A2 | Frame alignment word, used for synchronization of the start of each $90 \times 9$ block. |
| C1 | STS-1 ID identifies the STS-1 number (1 to N) for each STS-1 within an STS-N multiplexer for interleaving STS frames. |
| B1 | Bit-interleaved parity byte providing even parity over previous STS-N frame after scrambling, which checks errors on the regenerator section overhead. |
| E1 | Section-level 64 Kbit/s PCM orderwire (local orderwire), which provides a 64 Kbit/s voice channel for maintenance personnel in each of the regenerator and multiplexer sections. |
| F1 | 64 Kbit/s alarm channel set aside for user purposes |
| D1-D3 | 192 Kbit/s data communications channel for alarms, maintenance, control and administration between sections. |
| | **Line overhead** |
| H1-H3 | Pointer bytes used in frame alignment and frequency adjustment of payload data |
| B2 | Bit-interleaved parity for line-level error monitoring, which error-checks on the multiplexer section overhead. |
| K1,K2 | Two bytes allocated for signaling between line-level automatic protection switching equipment |
| D4-D12 | 576 Kbit/s data communications channel for alarms, maintenance, control, monitoring and administration at the line level |
| Z1, Z2 | Reserved for future use |
| E2 | 64 Kbit/s PCM voice channel for line-level orderwire |
| | **Path overhead** |
| J1 | 64 Kbit/s channel used to repetitively send a 64-byte fixed-length string so a receiving terminal can continuously verify the integrity of a path; the contents of the message are user-programmable |
| B3 | Bit-interleaved parity at the path level |
| C2 | STS path signal label to designate equipped versus unequipped STS signals and, for equipped signals, the specific STS payload mapping that might be needed in receiving terminals to interpret the payloads |
| G1 | Status byte sent from path-terminating equipment back to path-originating equipment to convey status of terminating equipment and path error performance |
| F2 | 64 Kbit/s channel for path user |
| H4 | Multiframe indicator for payloads needing frames that are longer than a single STS frame; multiframe indicators are used when packing lower-rate channels (virtual tributaries) into the SPE |
| Z3-Z5 | Reserved for future use |

| | 1 octet | 1 octet | 1 octet |
|---|---|---|---|
| Section overhead | Framing A1 | Framing A2 | STS-ID C1 |
| | BIP-8 B1 | Orderwire E1 | User F1 |
| | Data Com D1 | Data Com D2 | Data Com D3 |
| Line overhead | Pointer H1 | Pointer H2 | Pointer action H3 |
| | BIP-8 B2 | APS K1 | APS K2 |
| | Data Com D4 | Data Com D5 | Data Com D6 |
| | Data Com D7 | Data Com D8 | Data Com D9 |
| | Data Com D10 | Data Com D11 | Data Com D12 |
| | Growth Z1 | Growth Z2 | Orderwire E2 |

(a)Section overhead

| 1 octet |
|---|
| Trace J1 |
| BIP-8 B3 |
| Signal label C2 |
| Path status G1 |
| User F2 |
| Multiframe H4 |
| Growth Z3 |
| Growth Z4 |
| Growth Z5 |

(b)Path overhead

Figure 2.14: SONET STS-1 overhead octets

physical layer frame structure octets is 149.76 Mbit/s ($8KHz \times 260 octets \times 8 bits \times 9 columns$) on a 155.52 Mbit/s ($8KHz \times 270 octets \times 8 bits \times 9 columns$) transmission system, and 599.04 Mbit/s on a 622.08 Mbit/s transmission system.

In case of conventional circuit-switched networks, a piece of information that is addressed to a node is needed to be extracted from the entire signal while the remaining part must be left unchanged. Thus, multiplexers and telephone company channel banks in between the circuits require the demultiplexing and remultiplexing of the signal. As an example, consider the case

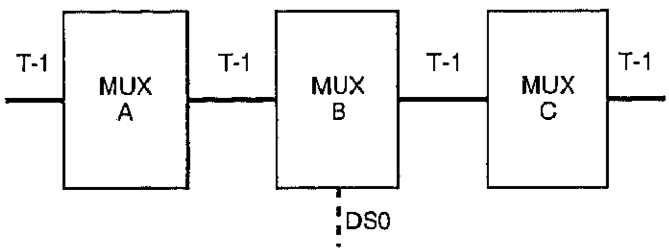


Figure 2.15: Example of demultiplex and remultiplex

shown in Figure 2.15. A T-1[9] multiplexer B (MUX B) receives data on a single T-1 circuit from multiplexer A (MUX A) and passes the data on to multiplexer C (MUX C). In the stream

[9] 1.544 Mbit/s

from MUX A, a single DS0 channel[10] is addressed to node B. The rest will pass on to MUX C and further on into the network. MUX B first demultiplexes every bit of the 1.544 Mbit/s signal, then removes the DS0 channel data, and re-multiplexes every bit. A few proprietary T-1 multiplexers allow this drop-and-insert capability, which means only part of the signal is to be demultiplexed and remultiplexed. However, this equipment will not communicate with that of other vendors.

SONET offers a standard drop-and-insert capability, which applies not only to 64 Kbit/s channels but to the higher data rates as well. SONET uses a set of pointers to locate channels within a payload and a payload within a frame. Thus information can be accessed, inserted, and removed with a simple adjustment of the pointer. Pointer information in the path overhead refers to the multiplex structure of the channels contained within the payload. A pointer contained in the line overhead serves a similar function for the entire payload. The use of this latter pointer is describe here. The position of a synchronous payload envelope (SPE) of an STS-1 frame can float within the frame. The payload including path overhead ($87 columns \times 9 rows$) can straddle two frames as in Figure 2.16. The H1 and H2 octets in the line overhead is the pointer of the

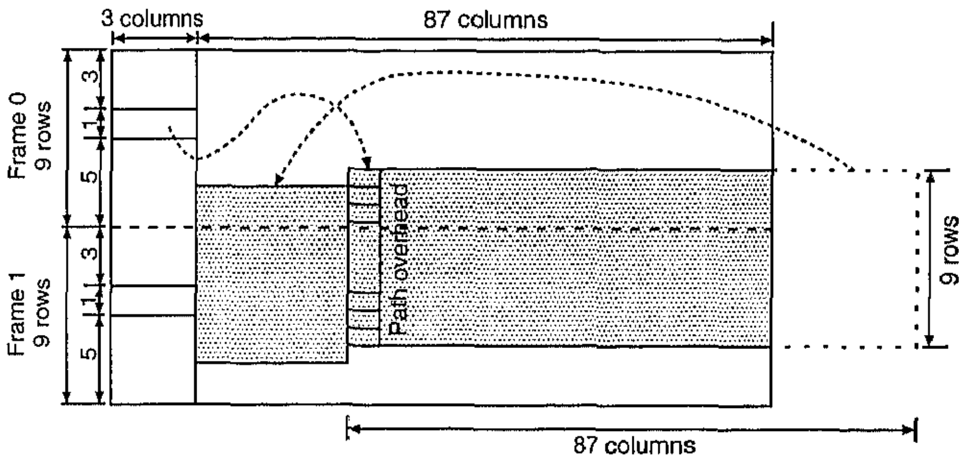


Figure 2.16: Representative location of SPE in STS-1 Frame

start of the payload. Each node must recalculate the pointer to alert the next receiving node of the exact location of the start of the payload. In order to cope with the timing differences between timing sources of the nodes, the payload is allowed to change the pointer value by one byte in an STS-1 frame. If the payload rate is higher than the local STS frame rate, the pointer is decreased by one octet position so that the next payload will begin one octet earlier. To prevent the loss of an octet on the payload, the H3 field is used to hold an extra octet for the frame Figure 2.17. Similarly, if the payload rate lags behind the frame rate, the insertion of the next payload is delayed by one octet. In this case, the octet which follows the H3 octet is left empty to allow for the movement of the payload (Figure 2.18).

[10] 64 Kbit/s

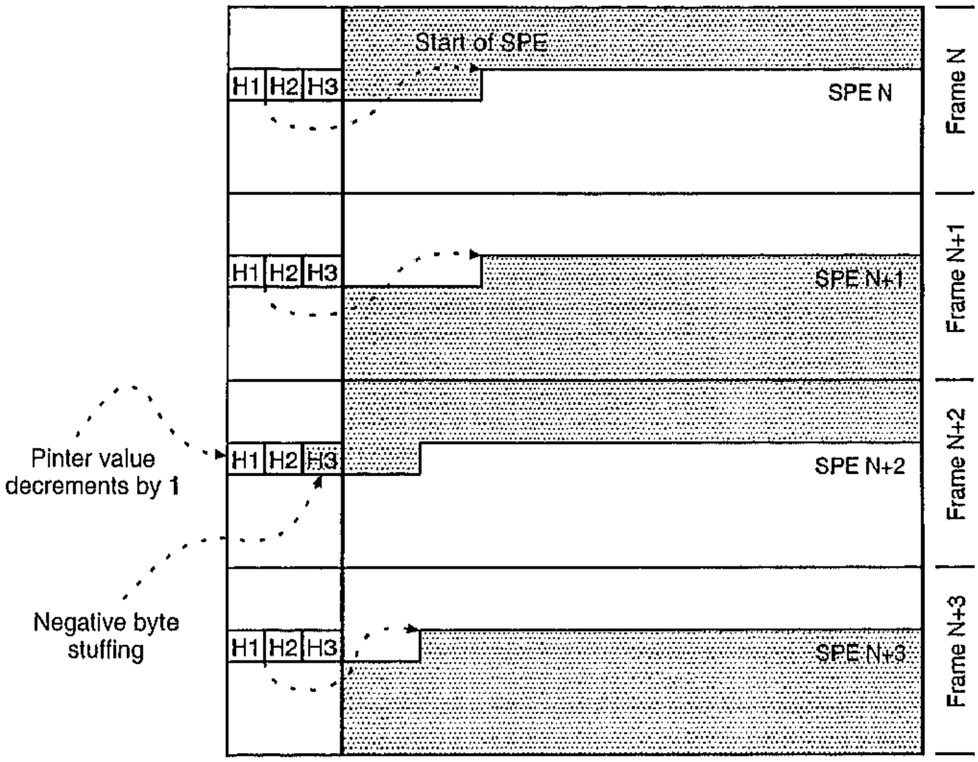


Figure 2.17: STS-1 pointer adjustment: negative pointer adjustment

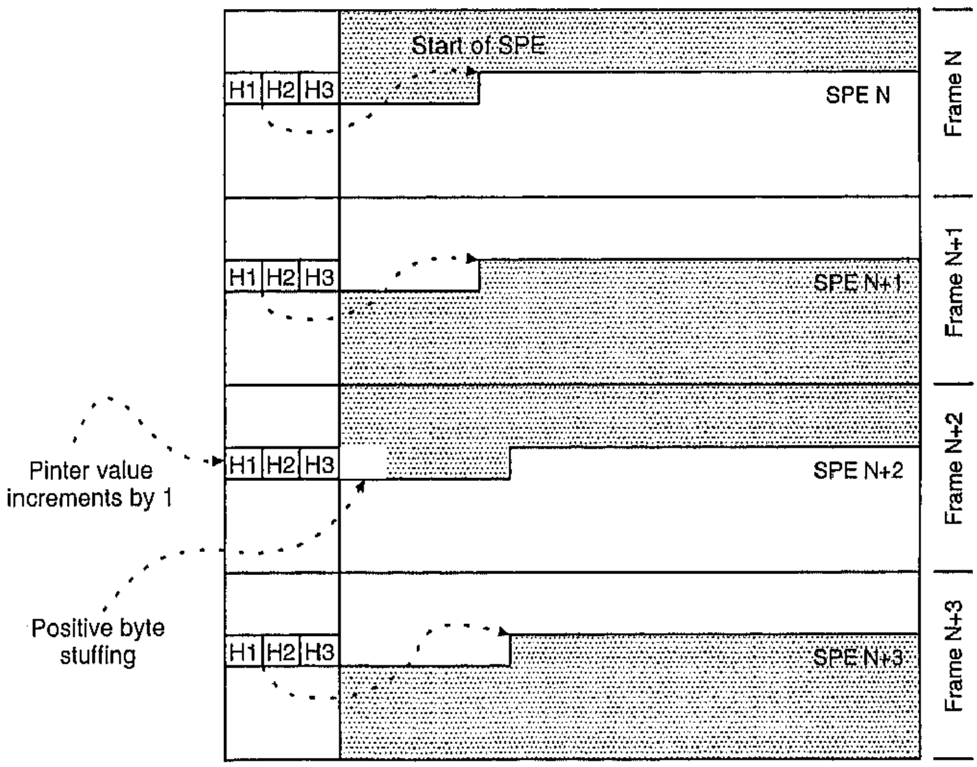


Figure 2.18: STS-1 pointer adjustment: positive pointer adjustment

### 2.3.2 Cell based interface

In cell based interface, a continuous stream of cells each containing 53 octets is transported without any regular framing related to a time frame (Figure 2.19). The physical layer cells

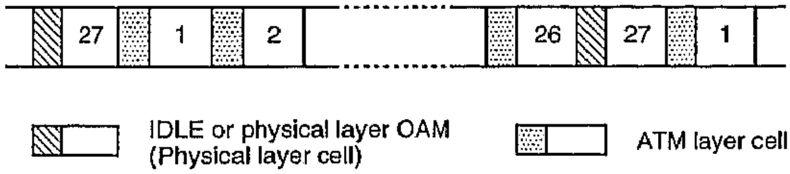


Figure 2.19: The cell based interface

are generated and interpreted in the physical cell based layer. The maximum spacing between successive physical layer cells is 26 ATM layer cells. After 26 continuous ATM layer cells, the insertion of physical layer cells is enforced to adapt the transfer capability to the interface rate. The physical layer cells can be either idle cells or physical layer OAM cells (PL-OAM), depending on operation and maintenance requirements. Physical layer cells are identified by a pre-defined header. Idle cells merely perform cell rate adaptation and PL-OAM cells convey OAM information concerning the physical layer itself. The pre-assigned header values for PL cell type are shown in Table 2.5. PL-OAM cells carry regenerator level (F1) and transmission path

Table 2.5: Pre-assigned values of the cell header at the physical layer

| Cell type | Octet 1 | Octet 2 | Octet 3 | Octet 4 |
|---|---|---|---|---|
| Idle cell | 00000000 | 00000000 | 00000000 | 00000001 |
| Physical layer OAM (F1) | 00000000 | 00000000 | 00000000 | 00000011 |
| Physical layer OAM (F3) | 00000000 | 00000000 | 00000000 | 00001001 |
| Reserved for use by PHY | pppp0000 | 00000000 | 00000000 | 0000ppp1 |

p: the bit is available for use by the PHY layer

(F3) level information. They are inserted in the ATM layer cell flow at least every 513 cells. Note that at the cell-based interfaces only F1 and F3 OAM flows. The digital section level flow (F2) is not used, since level F2 and F3 coincide for the cell-based interfaces, and the corresponding functions are supported by the F3 flow. Functions such as the performance monitoring, and detection and reporting of transmission errors are supported. Performance checking includes counting and calculating an error code over the ATM layer cells and IDLE cells between two subsequent PL-OAM cells. The results are conveyed in the information field of the PL-OAM cell, together with maintenance signaling, and a CRC on the PL-OAM cell information field itself. Additionally, the cell based interface specified by the ATM Forum supports 125$\mu$sec clock delivery across a transmission link by means of a special symbol at the private UNI only. The bit rate available for user information cells, signaling cells and OAM cells is 149.76 Mbit/s on a

155.52 Mbit/s transmission system, and 599.04 Mbit/s on a 622.08 Mbit/s transmission system, which is the same as the ratio 26:27. These values are identical to the payload of the SDH frames. When no ATM layer cells are available, physical layer cells are inserted.

### 2.3.3 Plesiochronous digital hierarchy based interface

Until the middle of 1970s, all switching was analog. Thus, analog signals received by the digital transmission systems are once converted to a digital stream using an internal crystal clock and reconverted to analog. A few years later, as faster digital transmission was realized, it made it possible to multiplex tributaries from a certain number of DS-1 or E-1 signals. However, the multiplexer had to take into consideration the fact that the clocks of the tributaries were all slightly different. In order to wipe out the difference of the clocks, the method called the **plesiochronous digital hierarchy** (PDH) was developed. "Plesio", from Greek, means "almost" and "plesiochronous" means "almost or nearly synchronous." In PDH, each clock speed is allowed within a certain range. The multiplexer takes bits from each tributary at the highest allowed clock speed and puts them into a higher-order stream. When there are no bits available in the input buffer because of a slower clock, it stuffs an additional bit to adapt the signal up to the higher clock speed. It also has a mechanism to notify the demultiplexer of stuffing, and the demultiplexer must know which bit to unstuff or negative stuff. Carrying ATM cells in PDH frames has the advantage of using the existing transmission network infrastructure, instead of having to rely on the deployment of new SDH transmission equipment. This PDH method was once the basis for all digital transmission systems [OMID93].

#### Transmission convergence characteristics.

Several ways for mapping ATM cells of different bit rates in PDH frames have been proposed. Some of them implemented cell based transmission on top of the existing PDH frame, using PL-OAM cells for maintenance. These old frame structures of the PDH signals have been partly replaced with new ones which are capable of supporting ATM cell transport and SDH element transport. They use an SDH-like approach, in which, maintenance, performance monitoring and reporting are based on the use of special octets added to the frame. The remaining payload of the PDH frame is filled with ATM cells, which are octet aligned to the PDH frame payload area. The ATM cells are delineated by using the HEC, and their payload is scrambled to avoid false frame and cell synchronization. As an example, the frame structure for the 34.386 Mbit/s $((59 \times 9 + 6) \times 8 \times 8000)$ PDH interface is shown in Figure 2.20. The following path overhead functions are defined.

- FA: Frame alignment signal
- EM: Error monitoring (bit interleaved parity; BIP-8)

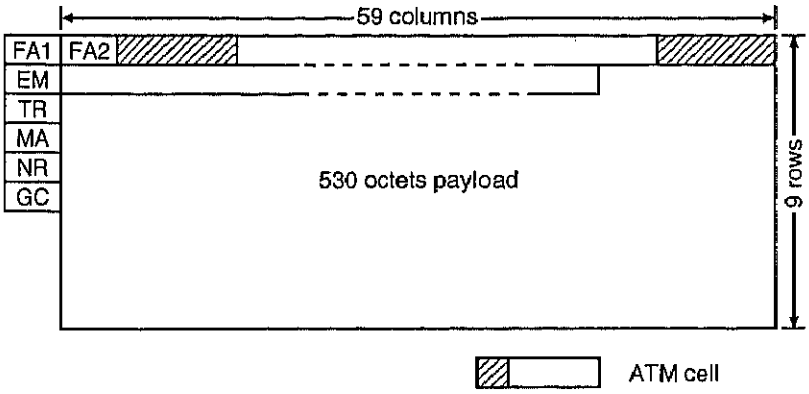


Figure 2.20: The frame structure at 34.368 Mbit/s

- TR: Trail Trace
- MA: Maintenance and adaptation byte; far end receive failure (FERF), far end block error (FEBE), payload type
- NR: Network operator byte
- GC: General purpose communications channel; data or voice for maintenance.

The functions supported by the above overhead are similar to the SDH functions. The ATM cells are mapped into the 530 payload octets ($59columns \times 9rows - 1octet$).

In contrast to the 34.368 Mbit/s mapping, the ATM Forum specified another mapping which makes use of **physical layer convergence protocol** (PLCP), which is defined in IEEE 802.6 (DQDB). Figure 2.21 illustrates the DS3 mapping using the PLCP. The integer number of ATM cells are enclosed in a 125 $\mu$sec PLCP frame, which is defined in the standard DS3 M-frame. In the PLCP frame, each ATM cell is preceded by PLCP framing octets A1 and A2, a path overhead identifier octet (POI), and a path overhead octet (POH). The framing octets allow delineation of the ATM cells without using the HEC. Hence, the HEC is only used for cell header checking and correction. The POI octets identify the contents of the succeeding POH octets, i.e., BIP-8, stuff counter, path status, and so on. The trailer of the PLCP frame is stuffed with a nibble (1100). As the data rate of DS3 is approximately 44.21 Mbit/s, the trailer should be between 13 and 14 nibbles (between 6.5 and 7 octets). Note that the BIP-8 indicator is computed over the POH and associated ATM cells of the previous PLCP frame. The PLCP mapping transfers occurs 8000 times (1/125 $\mu$sec) across the DS3 interface. Hence, the cell transfer rate is only 40.704 Mbit/s ($53 \times 12 \times 8 \times 8000$), which is only about 90 % of the DS3's approximately 44.21 Mbit/s payload rate.

| PLCP Framing 2 octets | | POI 1 octet | POH 1 octet | PLCP payload 53 octets | 13 or 14 nibbles |
|---|---|---|---|---|---|
| A1 | A2 | P11 | Z6 | First ATM cell | |
| A1 | A2 | P10 | Z5 | ATM cell | |
| A1 | A2 | P9 | Z4 | ATM cell | |
| A1 | A2 | P8 | Z3 | ATM cell | |
| A1 | A2 | P7 | Z2 | ATM cell | |
| A1 | A2 | P6 | Z1 | ATM cell | |
| A1 | A2 | P5 | X | ATM cell | |
| A1 | A2 | P4 | B1 | ATM cell | |
| A1 | A2 | P3 | G1 | ATM cell | |
| A1 | A2 | P2 | X | ATM cell | |
| A1 | A2 | P1 | X | ATM cell | |
| A1 | A2 | P0 | C1 | Twelfth ATM cell | Trailer |

Object of BIP-8 calculation

PLCP : physical layer convergence protocol
A1 : 11110110
A2 : 00101000
P0-P11 : path overhead identifier (POI)
POH : path overhead
Z1-Z6 : Growth octets (00000000)
X : unassigned
B1 : PLCP bit interleaved parity-8 (BIP-8)
G1 : PLCP path status
C1 : Cycle stuff counter
Trailer : 1100 (nibble)

Figure 2.21: The DS3 PLCP frame

### 2.3.4 FDDI based interface

The ATM Forum also specifies an FDDI based multimode fiber interface for private UNI. The physical medium dependent sublayer applies a 4B/5B line code and 125 M baud, resulting in a bit rate of 100 Mbit/s.

In contrast to other interfaces which use HEC for cell delineation, ATM cell delineation is done by using special line codes in FDDI based interface. The HEC can be used for error detection only, but not for error correction, because the line code causes bit error multiplication.

### 2.3.5 ATM specific transmission convergence sublayer functions

#### Cell rate decoupling.

Whenever no cell[11] is ready for transmission, special codings of the ATM cell called ***unassigned cells***[12] or *idle cells* will be inserted to adapt the cell stream to the transmission speed. All

[11] In this case, the cells include data cells and unassigned cells from ATM layer, or physical layer OAM cells

[12] The unassigned cells are visible at both the physical and ATM layer, though they are treated differently at each layer.

other cells are *assigned cells* which correspond to the cells generated by ATM layer. Those unassigned or idle cells will be discarded at the receiving side. The insertion and discarding of those unassigned or idle cells is called "cell rate decoupling" or "speed matching". Figure

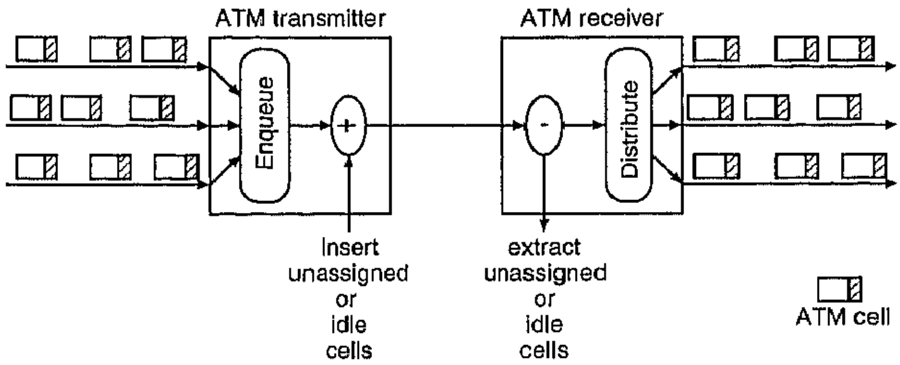


Figure 2.22: Cell rate decoupling

2.22 illustrates this operation between a transmitting device and a receiving ATM device. The transmitter multiplexes multiple cell streams. Cells from the streams are queued if no ATM slot is available. If the queue is empty when the next synchronous time slot arrives, the transmission convergence sublayer inserts an unassigned or idle cell. The receiver extracts those cells and demultiplexes the other assigned cells to the destinations.

ITU-T have placed this cell rate decoupling function in the physical layer and uses idle cells. In contrast, the ATM Forum places it in the ATM layer and uses unassigned cells. Hence, there is a potential low-level incompatibility if different systems follows different rules and cell types for cell rate decoupling. Thus, it is sometimes necessary to examine the ATM systems with regard to the inter-operability of the methods[MCDY94].

Unassigned cells and idle cells are identified by a standardized pattern for the cell header, which is shown in Tables 2.6 and 2.5, respectively. This is used throughout the ATM network

Table 2.6: Header pattern for unassigned cell and idle cell identification

| Cell type | Octet 1 | Octet 2 | Octet 3 | Octet 4 |
|---|---|---|---|---|
| Unassigned cell | 00000000 | 00000000 | 00000000 | 00000000 |
| Idle cell | 00000000 | 00000000 | 00000000 | 00000001 |

to identify unassigned cells and idle cells. Each octet of the information field of an idle cell is filled with 01101010.

Header error control.

According to the B-ISDN protocol reference model, the ATM cell header error control is a physical layer function and is described in the ITU-T Recommendation. Every ATM cell transmitter calculates the one octet length of header error control (HEC) over the first four octets of the cell header and inserts it into the fifth octet position. The transmitter calculates the HEC value using the polynomial $P(x) = x^8 + x^2 + x + 1$. The header bits (excluding the HEC field) multiplied by $x^8$ is divided by the polynomial $P(x)$. The remainder, whose initial value is preset to all zeros, is added to the fixed pattern called a "coset value" (01010101) to improve the cell delineation performance. This final value becomes the HEC code. The HEC code is capable of correcting any single-bit error in the header. It is also capable of detecting many patterns of multiple-bit errors in the ATM header. Since the header contains important information which tells the ATM layer what to do with the cell, it must be flawless; otherwise it might be delivered to the wrong user, or an undesired function in the ATM layer may be inadvertently invoked.

The TC sublayer generates HEC on transmit and uses it to determine if the received header has any errors according to the state-event diagram depicted in Figure 2.23. At first, the receiver

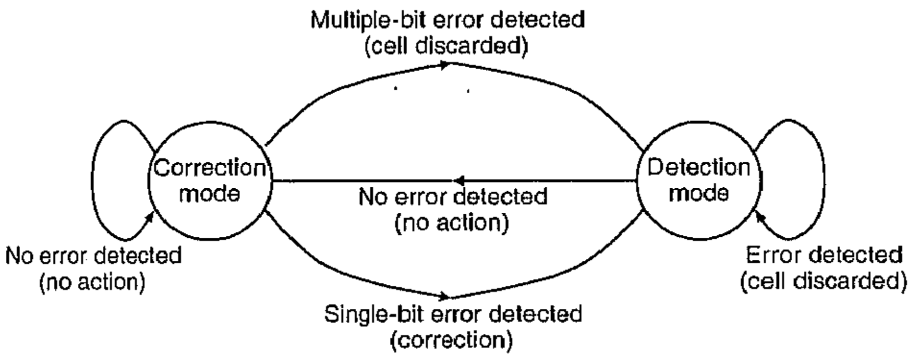


Figure 2.23: HEC receiver actions

is in *correction mode.* As each cell is received, the HEC calculation and comparison is performed. As long as no errors are detected, the receiver remains in correction mode. When a single-bit error is detected, the receiver will correct the error, or when a multiple-bit error is detected, the received cell is discarded. In both cases, the receiver now moves to *detection mode.* In this receiver state, no attempt is made to correct errors and each cell with a detected single-bit or multiple-bit error is discarded. The receiver remains in detection mode as long as erroneous cells are received. When a header is examined and found not to be in error, the receiver switches back to correction mode.

The above receiver operation has been chosen to take into consideration the error characteristics of fiber based transmission systems, which appears to be a mixture of single-bit errors and relatively large error bursts. For some transmission systems, the error-correction capability, which is more time-consuming, might not be invoked. The specified HEC method guarantees recovery from single-bit errors, and a low probability of delivering cells with erroneous headers under bursty error conditions.

Figure 2.24 shows the performance of the HEC mechanism as a function of random bit error probability[HAND94][STAL95]. For a bit error probability of about $10^{-8}$, the probability that

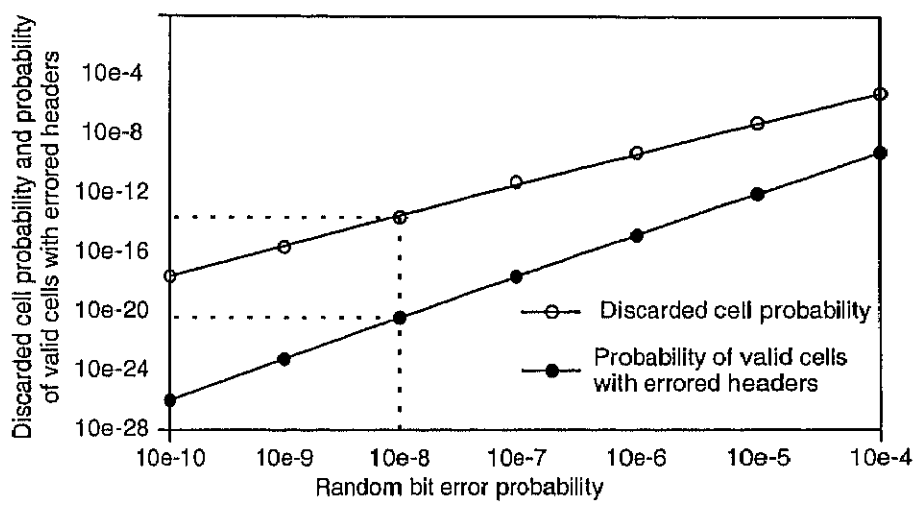


Figure 2.24: Impact of random bit errors on HEC performance

cells will be discarded (as a result of header errors detection but not correction) is approximately $10^{-13}$ and the probability of valid cells with erroneous headers is about $10^{-20}$.

The HEC can also be used to locate cells by the TC. The HEC will not match random data in the cell payloads when the five octets being checked are not part of the header. Thus, it can be used to find cells in a received bit stream. Once several cell headers are found through the use of HEC, TC will expect the next 53 octet cell later on. This process, called "HEC-based cell delineation", is described next.

### Cell delineation.

*Cell delineation* is the process which allows identification of the cell boundaries [HAND94]. The cell delineation algorithm has to be self supporting so that it can be transported on every network interface, independent from the used transmission system (cell based, SDH based, plesiochronous and so on) [PRYC93]. The recommended cell delineation mechanism is based on the correlation between the first four header octets to be protected and the corresponding HEC octet. The state diagram for HEC-based cell delineation is shown in Figure 2.25. In the HUNT state, a bit by bit check of the assumed header field is performed. When octet timing can be obtained in the PHY layer, the cell delineation in the HUNT state may be performed octet by octet. This is applicable for instance, if SDH transmission is applied. When correct, the PRESYNCH state is entered. In the PRESYNCH state, it is assumed that a header has been found and a correct cell delineation has been conducted. Now the HEC correlation check is performed cell

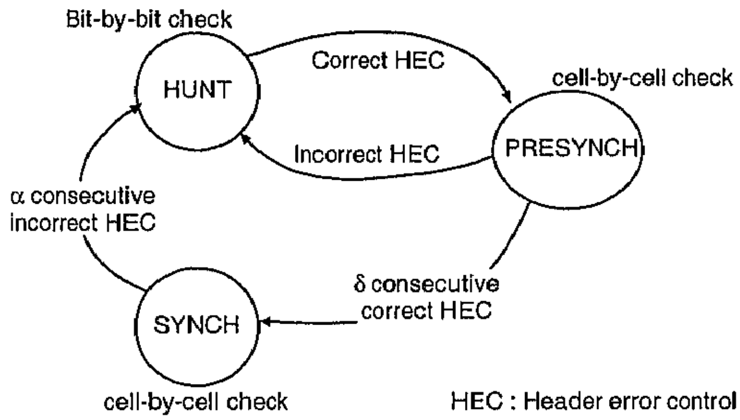


Figure 2.25: Cell delineation state diagram

by cell. However, a further confirmation is required. Therefore, the correctness of the HEC field is checked. An incorrect HEC before the SYNCH state was reached makes the system goes back to the HUNT state. If $\delta$ consecutive correct HECs are found, the SYNCH state is entered. The system leaves the SYNCH state for the HUNT state due to the loss of cell delineation if $\alpha$ consecutive incorrect HECs are encountered. It is obvious that the value of $\alpha$ and $\delta$ influence the performance of the cell delineation process. Robustness against false misalignment caused by bit errors depends on $\alpha$, while robustness against false delineation in the resynchronization phase depends on $\delta$. Values of $\alpha = 7$ and $\delta = 6$ are suggested by ITU-T for SDH based physical layers, and $\alpha = 7$ and $\delta = 8$ are suggested for cell based physical layers.

The next two figures (Figure 2.26 and 2.27) show the impact of random bit errors on cell-delineation performance for various values of $\alpha$ and $\delta$ [STAL95]. Figure 2.26 shows the average amount of time that the receiver will maintain synchronization in the face of errors with respect to a parameter $\alpha$. With $\alpha = 7$, once a system state is SYNCH, a 155.52 Mbit/s ATM system will remain in a SYNCH state for more than a year even when the bit error probability is about $10^{-4}$. Figure 2.27 shows the average amount of time to acquire synchronization as a function of bit error probability, with $\delta$ as a parameter. With $\delta = 6$, the same system with the same bit error probability will need about 10 cells or 28 $\mu$s to re-enter SYNCH after a loss of cell synchronization.

This cell delineation could fail if the header HEC correlation were imitated in the information field of ATM cells. This might be effected by a malicious user or might happen inadvertently by an application that accidentally uses the same generator polynomial. To increase the security and robustness of the delineation process against them, the bits of the information field are randomized. For the SDH-based interfaces, the information field contents will be scrambled using a self-synchronizing scrambler with a polynomial of $x^{43} + 1$ [PRYC93]. The scrambler is effective only in PRESYNCH and SYNCH and is disabled in HUNT. For cell-based interfaces, a distributed sample scrambler is recommended. This randomizes the information field contents by

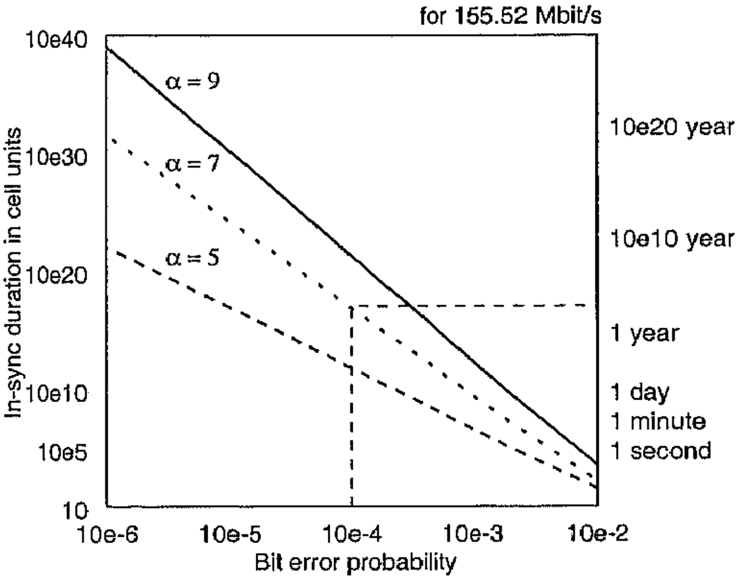


Figure 2.26: Impact of $\alpha$ on performance of cell delineation

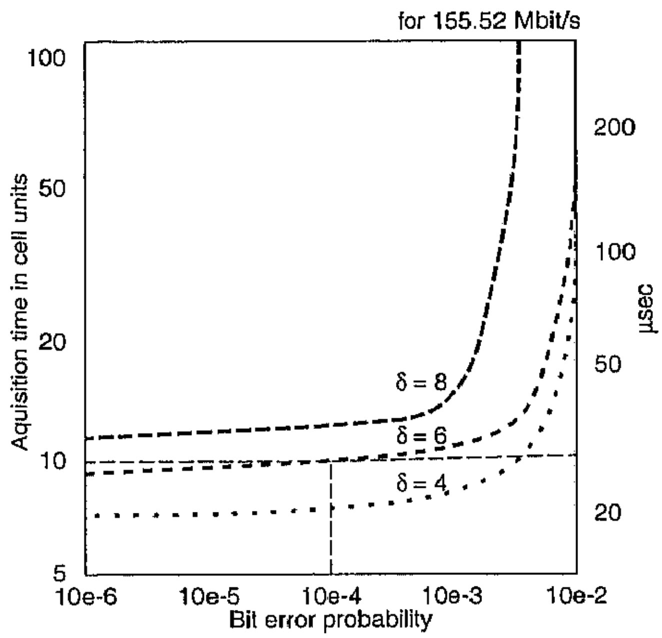


Figure 2.27: Impact of $\delta$ on performance of cell delineation

modulo addition of a pseudo-random sequence. The receiver descrambles it by modulo addition of an identical locally generated pseudo-random sequence, which has synchronous phase to the transmitted cells.

## 2.4 ATM layer

This section describes the ATM layer. The relationship of the ATM layer to the physical layer and its division into a **virtual path** (VP) and **virtual channel** (VC) level are covered in detail. A transfer unit of ATM, called ATM cell, is described next.

### 2.4.1 Physical links and ATM virtual paths and channels

A key concept in constructing the logical ATM layer connection is the **virtual channel** (VC) and the **virtual path** (VP). A virtual channel is analogous to a virtual circuit in X.25 or a frame relay logical connection. It is the basic unit of switching in B-ISDN. A virtual path is a bundle of virtual channels that have the same endpoints (Figure 2.28). Thus, all of the cells

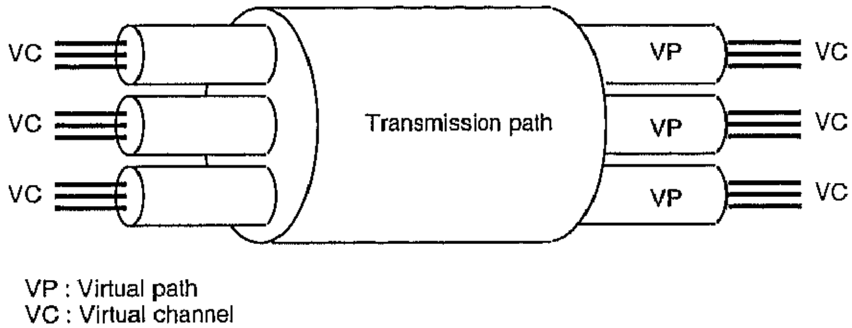


Figure 2.28: ATM connection relationships

flowing over all of the virtual channels in a single virtual path are switched together.

Figure 2.29 shows the ATM transport hierarchy derivation. The physical layer can be subdivided into three functional levels; transmission path level, digital section level, and regenerator section level. Transmission path extends between network elements that assemble and disassemble the payload of a transmission system. For end-to-end communication, the payload is end-user information. For user-to-network communication, the payload may be signaling information. Cell delineation and header error control functions are required at the endpoints of each transmission path. The digital section extends between network elements which assemble and disassemble continuous bit or byte streams. These are the exchanges or signal transfer points in a network that are involved in switching data streams. The regenerator section is a portion of a digital section extending between two adjacent regenerators. An example of this level is a repeater that is used to simply regenerate the digital signal along a transmission path. No switching is involved in the regenerator section.

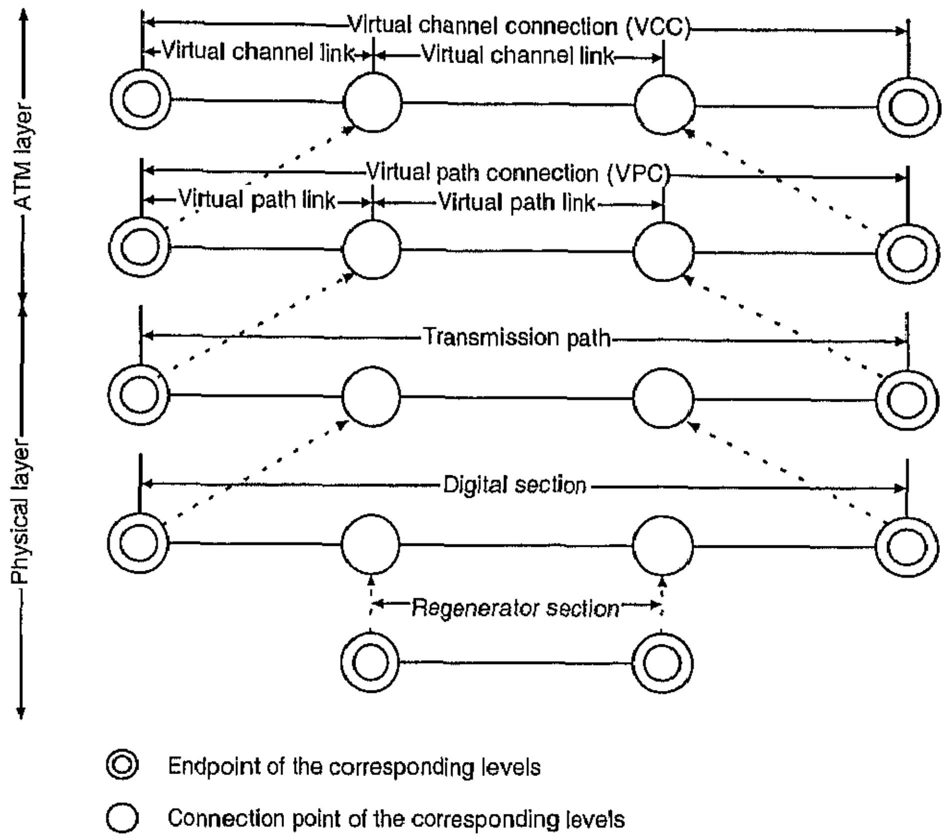


Figure 2.29: ATM transport hierarchy

With respect to ATM layers, an ATM device generally corresponds to either an endpoint or a connecting point for a VP or VC. A VP link or a VC link can exist between an endpoint and a connecting point or between connecting points. A **virtual path connection** (VPC) or a **virtual channel** (VCC) is an ordered list of VP or VC links, respectively. A VPC or a VCC exists only between endpoints.

With regard to VPC and VCC, specific characteristics are specified [STAL95].

- **Quality of service (QoS)**
  A QoS defined by parameters are provided to a user.

- **Switched and semi-permanent virtual channel**
  Both switched connections, which require call-control signaling, and dedicated channels can be provided.

- **Cell sequence integrity**
  The sequence of transmitted cells within a virtual channel is preserved.

- **Traffic parameter negotiation and usage monitoring**
  Traffic parameters can be negotiated between a user and the network for each virtual chan-

nel. The input stream of cells is monitored by the network to ensure that the negotiated parameters are not violated.

In addition to this, one more specific characteristic is appended to the VPC:

- **Virtual channel identifier restriction within a virtual path**
  Several virtual channel identifiers may not be available to the user of the virtual path, but may be reserved for network use.

More details about the VC and the VP will be discussed next.

### Virtual channel level.

A **virtual channel** (VC) is a generic term used to describe unidirectional transport of ATM cells associated by a common unique identifier value. A **virtual channel identifier** (VCI) identifies a particular VC link for a given virtual path connection (VPC). Switching at a VC connecting point is based on this VCI and the VPI. A **virtual channel link** is a means of unidirectional transport of ATM cells between a point where a VCI value is assigned and the point where that value is translated or terminated. A virtual channel connection (VCC) is defined as a concatenation of VC links that extends between two points where the adaptation layer is accessed. It determines a unidirectional flow of ATM cells from one user to one or more other users. VCCs are provided for the purpose of user-user, user-network, or network-network information transfer.

In narrowband ISDN, the D channel is used for control signaling of calls or connections on B and H channels. In B-ISDN, however, there is no fixed-rate structure for control signaling. Instead, more flexible methods of control signaling for virtual channels and virtual paths are provided.

In establishing/releasing the virtual channels at the B-ISDN UNI, one or a combination of four methods can be used [STAL95]:

1. **Semi-permanent virtual channel**
   A semi-permanent virtual channel may be used for control signaling. No signaling procedure is necessary.

2. **Meta-signaling channel**
   If there is no pre-established call-control signaling channel, one must be set up. For this purpose, a control-signaling must be exchanged between the user and the network on some channel. Hence, a permanent channel is needed, which can be used to set up a signaling virtual channel that can be used for call control. Such a permanent channel is called a meta-signaling channel, since the channel is used to set up signaling channels.

3. **User-to-network signaling virtual channel**
   The meta-signaling channel can be used to set up a signaling virtual channel between the user and the network for call-control signaling. This user-to-network signaling virtual channel can then be used to set up end-to-end virtual channel connections to carry user data.

4. **User-to-user signaling virtual channel**
   The meta-signaling channel can also be used to set up a user-to-user signaling virtual channel. If a virtual path already exists between two B-ISDN UNIs, a user-to-user signaling virtual channel can then be used to allow the two end users to set up/release user-to-user virtual channels to carry user data.

In establishing a VCC, traffic parameters to control virtual channels are negotiated between the user and the network. These parameters can be renegotiated during the connection. When a connection is set up, all cells originating from the user are monitored by the network to ensure that the agreed parameters are not violated. This mechanism is called "usage parameter control" (UPC). In addition to the parameter negotiation, the quality of service (QoS) for the connection is also negotiated during the VCC establishment phase. The VCI can be determined by the user or the network, or by the negotiation between user and network, or other standardized value will be used as the VCI. Although the value of the VCI field is generally independent of the service provided over the VC, for simplification, the same VCI value will be used at all B-ISDN UNIs for some fundamental functions such as meta-signaling. Once a VCC is established, cell sequence integrity will be preserved. In other words, the cells must be delivered in the same order in which they are sent.

#### Virtual path level.

A virtual path defines an aggregate bundle of VCs between VP endpoints. A **virtual path identifier** (VPI) in the cell header identifies a bundle of one or more VCs. A **virtual path link** is a group of VC links, identified by a common value of VPI, between a point where a VPI value is assigned and the point where that value is translated or terminated. Switching at a VP connecting point is based on the VPI while the VCI is ignored. A virtual path connection (VPC) is defined as a concatenation of VP links that extends between the point where the VCI values are assigned and the point where those values are translated or removed (i.e., extending the length of a bundle of VC links that share the same VPI). A VPC defines a unidirectional flow of ATM cells from one user to one or more other users. VPCs are provided for the purpose of user-user, user-network, or network-network information transfer.

For virtual paths, the following methods can be used to establish/release a VPC between VPC endpoints [STAL95]:

1. Semi-permanent basis
   A virtual path can be established/released on a semi-permanent subscription basis by prior agreement. In this case, no signaling procedure is necessary.

2. Customer-control basis
   A Virtual path establishment/release may be controlled by the customer. In this case, the customer uses a signaling virtual channel to request the virtual path from the network.

3. Network-control basis
   A Virtual path is established/released by the network using network management procedures. In this case, the network establishes a virtual path for its own convenience. The path may be network-to-network, user-to-network, or user-to-user.

During VPC establishment, the traffic parameters for the VPC are negotiated between the user and the network, which can be re-negotiated afterward if necessary. All input cells from the user to the network are monitored to make sure that the agreed traffic parameters are kept. In addition, in the VPC establishment phase, the QoS of the VPC is selected from the range of QoS classes provided by the network so as to meet the highest QoS of the VCCs carried. Some VCIs within a VPC may be reserved for management such as network OAM and will not be available to the user. Standards do not require a network to preserve cell sequence integrity for a VPC. However, the cell sequence integrity requirement of a VCC still applies.

VP and VC switching.

VCIs and VPIs in general have significance for only one link. In a VCC/VPC, the VCI/VPI value will be translated at VC/VP switching entities. Figures 2.30 and 2.31 illustrate the concepts

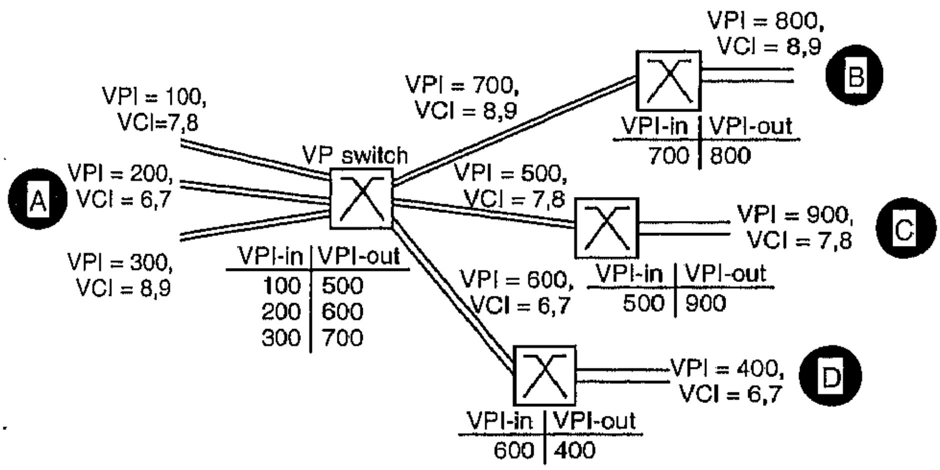


Figure 2.30: Virtual path switching

of VP and VC switching. VP switches terminate VP links and translate incoming VPIs to the corresponding outgoing VPIs according to the destination of the VP connection. VCI values remain unchanged. In Figure 2.30, virtual paths are established between subscriber A and subscriber B, C, and D. Each virtual path has two individual connections or virtual channels. Within VP switches, only VPI values are switched and VCIs are not changed. In the Figure 2.30, same VCI values can be used for different VPIs (e.g., VCI 7 of VPI 100 and VPI 200), since the different VPI values allow the two endpoints to discriminate between the two virtual connections. VC switches in Figure 2.31 terminate VC links, and VP links if necessary. A VC

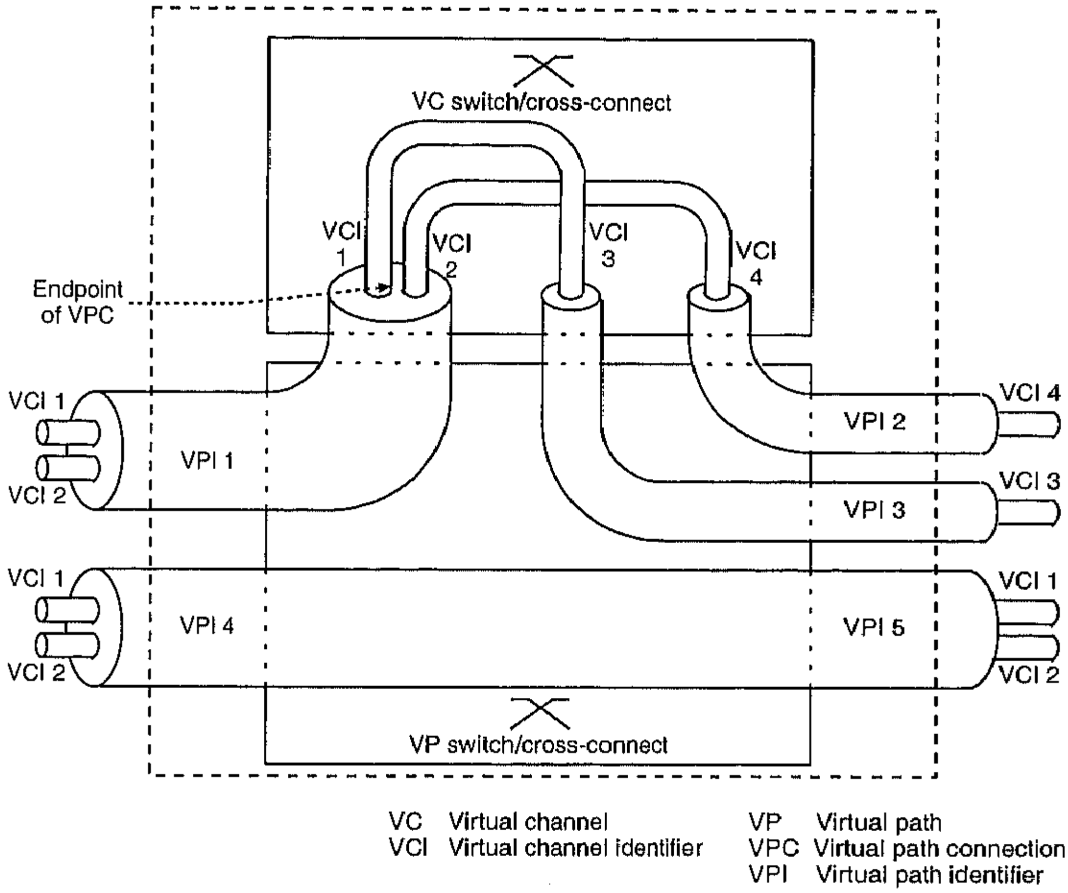


Figure 2.31: Virtual channel/virtual path switching

switch therefore switches both virtual paths and virtual channels, and so both VPI and VCI translation is performed. As VC switching implies VP switching, a VC switch can also handle only VP switching.

### Applications of virtual channel connection.

The endpoints of a VCC/VPC may be end users, network entities or an end user and a network entity. In all cases, all cells associated with an individual VCC/VPC are transported along the same route through the network. Within a VCC/VPC, cell sequence integrity is preserved, i.e., cells are delivered in the same order in which they are sent. The following three examples of a VCC usage can be envisaged:

- User-to-user VCCs are able to carry end-to-end user data and signaling information. A VPC between end users provides them with a transmission *pipe*. The VCC organization of the VPC is up to the two end users, provided the set of VCCs does not exceed the VPC capacity.
- User-to-network VCCs may, for instance, be used to access local connection related functions (user-network signaling). A user-to-network VPC can be used to aggregate traffic from an end user to an network exchange or network server.
- network-to-network VCC applications include network traffic management and routing functions. A network-to-network VPC can be used to organize user traffic according to a predefined routing scheme or to define a common path for the exchange of network management information.

### 2.4.2 ATM cells

The asynchronous transfer mode makes use of small fixed-size packet called a "cell" as the basic transmission element. There are several advantages of using small and fixed-size cells. First, queuing delay may be reduced by means of small cells. Second, fixed-size cells can be switched more efficiently, which is important for the very high data rates of ATM.

A cell consists of a five-octet header and a 48-octet information field. Concerning cell structure, two standard coding schemes exist: the **user-network interface** (UNI) and the **network-network interface** (NNI). The UNI is the interface between the user or **customer premises equipment** (CPE) and the network switch (S or T reference point) while the NNI is the interface between switches or between networks. Each frame format is illustrated in Figure 2.32[13]. Figure 2.32(a) is the format of the ATM cell at the UNI.

The cell header contains a logical address in two parts: an eight-bit virtual path identifier (VPI) and a 16-bit virtual channel identifier (VCI). In addition, a four bits of generic flow control (GFC) field, three bits of payload type indicator (PTI), and one bit of cell loss priority (CLP) indicator are included in the header. The entire header is protected by a one-octet header error control (HEC) field. The format of the ATM cell at the NNI(Figure 2.32(b)) is different from that at the UNI in the use of bits 5-8 of octet one. There is no generic flow control field of four bits in the format of NNI, which are used as a part of the VPI. The fields defined within the cell header have their significance in the ATM layer only.

[13]The octets are sent in increasing order starting with octet one of the header. Therefore, the cell header will be sent first, followed by the information field. Within an octet, the bits are sent in decreasing order starting with bit eight. For all fields in an ATM cell, the most significant bit (MSB) is first sent.

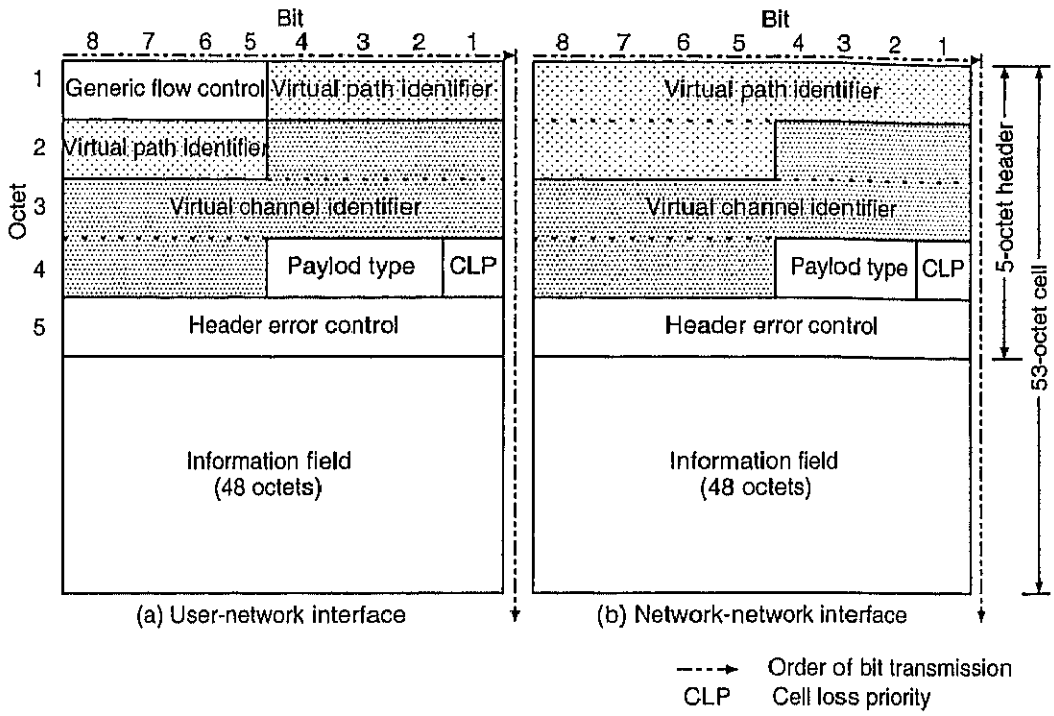


Figure 2.32: ATM cell format

Generic flow control (GFC) field.

The generic flow control (GFC) is a 4-bit field intended to support simple implementation of multiplexing which appear only at the user-network interface. Its default value is 0000 as long as the GFC function is not used. GFC information is carried either in assigned or unassigned cell.

Although the details of its application call for further study, two general functions are listed here: flow control across a T interface and medium access control across an S or T interface. The GFC mechanism helps control the traffic flow for different qualities of service from ATM connections at the B-ISDN UNI. One candidate for the use of this field is a multiple-priority level indicator to control the flow of information in a service-dependent manner. In any case, GFC is used to alleviate short-term overload conditions which may occur in the customer network.

With respect to medium access control, multiple terminals may share a single access link to B-ISDN, as with the basic-rate interface to ISDN. The GFC mechanism supports both point-to-point and point-to-multipoint configurations. In the former configuration, the GFC can be used to reduce the cell flow from each terminal. In the latter configuration, the GFC can only provide aggregate flow control. The GFC might also be used for a common media access control.

The exact GFC procedure is not yet defined. Up to now, standards define the *uncontrolled mode*, where the 4-bit GFC field is always coded as zeros. If any equipment receives too many non-zero GFC values within certain cell times, layer management should be notified.

### Virtual path identifier (VPI) field.

The **virtual path identifier** (VPI) field constitutes a routing information for the network. It consists of eight bits at the B-ISDN UNI and 12 bits at the NNI, allowing for more virtual paths to be supported within the network. Pre-assigned VPI values are used for some special purposes. All bits of the VPI field are set to zero in an unassigned cells.

### Virtual channel identifier (VCI) field.

Together with the VPI field, the **virtual channel identifier** (VCI) field constitutes the routing information of a cell. It is used for routing a cell to and from the end user. Thus, it functions much like a *service access point* (SAP). A field of 16 bits is used for the VCI at the B-ISDN UNI as well as the NNI. It also has some pre-assigned values, which will be described later.

### Payload type indicator (PTI) field.

Three header bits are used for the **payload type indicator** (PTI). They indicate the type of information in the information field. Table 2.7 shows the interpretation of the PTI bits. The payload of user information cells contains user information as well as service adaptation

Table 2.7: Payload type indicator (PTI) field coding

| PTI coding | Interpretation |
|---|---|
| 000 | User data cell, AAL_indication = 0, congestion not experienced |
| 001 | User data cell, AAL_indication = 1, congestion not experienced |
| 010 | User data cell, AAL_indication = 0, congestion experienced |
| 011 | User data cell, AAL_indication = 1, congestion experienced |
| 100 | OAM F5 segment associated cell |
| 101 | OAM F5 end-to-end associated cell |
| 110 | Resource management cell |
| 111 | Reserved for future function |

functions.

A value of 0 in the first bit indicates user information (i.e., information from the next higher layer). In this case, the second bit is the **congestion indication** (CI) bit which may be modified by any congested network element to inform the upstream of its state. The third bit, known as the "ATM-layer-user-to-ATM-layer-user indication" (AAL_indication) bit is currently used by AAL type 5 to identify the last cell. The use of this bit in the ATM adaptation layer type 5 (AAL 5) will be discussed in the next section. A value of 1 in the first bit indicates that this cell carries network management or maintenance information. This indication allows the insertion of network-management cells onto a user's VCC without impacting the user's data. Thus, it can provide in-band control information.

### Cell loss priority (CLP) field.

The cell loss priority (CLP) field consists of one bit which is used explicitly to indicate the cell loss priority. It is used to provide guidance to the network in the event of congestion. A value of zero means that the cell is of the highest priority, which is less likely to be discarded. If the value of the CLP field is one, the cell has low priority, which is subject to be selectively discarded depending on the network conditions. Nevertheless, the agreed quality of service (QoS) parameters will not be violated. The value of a CLP bit may be set by the user or by the network as a result of a policing action. The user might set the CLP field to one to send extra information which may be delivered to the destination if the network is not congested. The network may set this field to one for any data cell which violates the traffic agreement. In this case, the switch that sets the CLP flag realizes that the cell exceeds the agreed traffic parameters and that the switch is capable of handling the cell. At a later point in the network, if congestion is encountered, this cell will be marked for discard in preference to cells that fall within agreed traffic limits.

### Header error control (HEC) field.

This field provides error checking of the header. Though it is part of the cell header, it is not used by the ATM layer. It contains the header error control (HEC) sequence which is processed by the *transmission convergence* (TC) sublayer of the physical layer. A description of the HEC algorithm is given in Section 2.3.5.

### Pre-assigned cell header values.

Pre-assigned cell header values for the UNI by ITU-T are shown in Table 2.8. As described earlier, an *idle cell* is only visible at the physical layer and not passed to the ATM layer. It is used for stuffing unused bandwidth at the physical layer. *Unassigned cells* are visible at both the physical layer and the ATM layers, but are treated as ATM layer cells at the physical layer. They occupy available but unused bandwidth positions in the cell stream at the ATM layer. These idle and unassigned cells allow full asynchronous operation by both sender and receiver at the corresponding layers. The distinction between them is the value of bit one of octet four of the header (i.e., CLP bit position) as shown in Table 2.8. In addition, while the idle cells cannot use the GFC field, the unassigned cells may be used for GFC purposes, since GFC is not a physical layer function. The *meta-signaling cell* is used to negotiate on signaling VCI and signaling resources. A *general broadcast signaling cell* carries information which has to be broadcast to all terminals at the UNI. The *point-to-point signaling cell* is used for signaling a point-to-point configuration on a UNI or NNI at the ATM layer (i.e., the network only sees one signaling entity at the other side). The *segment OAM F4 flows* and *end-to-end OAM F4 flows* are coded by VCIs 0x0003 and 0x0004, respectively. Within the virtual path they carry

Table 2.8: Pre-assigned cell header values at the UNI

| Type | GFC | VPI | | VCI | | | | PTI | CLP |
|---|---|---|---|---|---|---|---|---|---|
| Unassigned cell | gggg | 0000 | 0000 | 0000 | 0000 | 0000 | 0000 | xxx | 0 |
| Idle cell* | 0000 | 0000 | 0000 | 0000 | 0000 | 0000 | 0000 | 000 | 1 |
| Reserved for physical layer* | pppp | 0000 | 0000 | 0000 | 0000 | 0000 | 0000 | ppp | 1 |
| Meta-signaling | gggg | yyyy | yyyy | 0000 | 0000 | 0000 | 0001 | 0a0 | c |
| General broadcast signaling | gggg | yyyy | yyyy | 0000 | 0000 | 0000 | 0010 | 0aa | c |
| Point-to-point signaling | gggg | yyyy | yyyy | 0000 | 0000 | 0000 | 0101 | 0aa | c |
| Segment OAM F4 flow cell | gggg | zzzz | zzzz | 0000 | 0000 | 0000 | 0011 | 0a0 | a |
| End-to-end OAM F4 flow cell | gggg | zzzz | zzzz | 0000 | 0000 | 0000 | 0100 | 0a0 | a |
| Segment OAM F5 flow cell | gggg | zzzz | zzzz | vvvv | vvvv | vvvv | vvvv | 100 | a |
| End-to-end OAM F5 flow cell | gggg | zzzz | zzzz | vvvv | vvvv | vvvv | vvvv | 101 | a |
| Resource management cell | gggg | zzzz | zzzz | vvvv | vvvv | vvvv | vvvv | 110 | a |
| User information cell | gggg | zzzz | zzzz | wwww | wwww | wwww | wwww | 0CA | c |
| | Octet 1 | | Octet 2 | | Octet 3 | | Octet 4 | | |

* defined as an invalid pattern by the ATM Forum.
A: ATM-layer-user-to-ATM-layer-user indication bit.
a: Bit is used by appropriate ATM layer function.
C: Congestion indication bit.
c: Originator sets the CLP, but the network may change the value.
g: Bit is used by the GFC protocol.
p: Bit is used by the physical layer.
v: Any VCI value other than 0.
w: Any VCI value above 15.
x: Bit is a *don't-care* bit.
y: Any VPI value. For VPI=0, the VCI value is reserved for user signaling with local exchange.
z: Any VPI value.
VCI values form decimal 8 to 15 are reserved for future use.

maintenance information. The *segment* and *end-to-end OAM F5 flows* are coded by PTIs 0x4 and 0x5 within the virtual channel in which they carry maintenance information. Value 0x6 of the PTI is reserved for resource management on the virtual channel.

The pre-assignment of cell header values by the ATM Forum is slightly different from this recommendation, by making all PTI and CLP bits of meta-signaling and general broadcast headers available for use at the ATM layer. Furthermore, the Forum defines some additional pre-assigned header value at the ATM layer. The interim local management interface (ILMI) VCI is one of them. The aim is to provide any ATM user device with status and configuration information concerning the virtual path and channel connections available at its UNI. In addition, more global operations and network management information may also facilitate diagnostic procedures at the UNI. Moreover, several VPI/VCI pairs for special functions have been reserved (Table 2.9).

Table 2.9: Pre-assigned values of the cell header at the ATM layer by the ATM Forum

| Type | GFC | VPI | | VCI | | | | PTI | CLP |
|---|---|---|---|---|---|---|---|---|---|
| ILMI | gggg | yyyy | yyyy | 0000 | 0000 | 0001 | 0000 | 0aa | c |
| LECS | gggg | yyyy | yyyy | 0000 | 0000 | 0001 | 0001 | 0aa | c |
| PNNI Hello protocol | gggg | yyyy | yyyy | 0000 | 0000 | 0001 | 0010 | 0aa | c |
| | Octet 1 | | Octet 2 | | Octet 3 | | Octet 4 | | |

ILMI: Interim local management interface
LECS: LAN emulation configuration server
PNNI: Private network-network interface

### 2.4.3 ATM service categories

ATM networks are expected to support various types of services. These services are classified into four types labeled Class A, B, C and D. The ATM Forum has introduced a more comprehensive and general service architecture which makes ATM suitable for the user to select specific combinations of traffic and performance parameters. With respect to the perspective of network and service operators, this architecture enables appropriate tariffing strategies to be deployed as well as offers a range of network services with selectable cost/performance levels. The rationale for this architecture is that, while an appropriate AAL is selected to satisfy the requirements of an application at the edges of the network, the ATM layer behavior should not rely on the AAL protocols nor on the higher layer protocols that are application specific. Thus, by providing selectable options within the ATM layer itself, it is possible to achieve a much greater degree of flexibility, fairness, and network utilization in mixed traffic environments having heterogeneous requirements of network resources [LAMB96]. An ATM service category represents a class of connections that have homogeneous characteristics in terms of traffic pattern, QoS requirements, and the possible use of control mechanisms [LAMB96]. Some of these categories require tightly constrained cell transfer delays (CTDs) and cell delay variations (CDVs), and reasonable amounts of cell loss ratios (CLRs) as well, while others are not so sensitive to these delays or cell losses. These service categories are as follows [LAMB96][GADE97][IBE97][KARI00]:

- Constant bit rate (CBR): This category of service is intended for real-time traffic that requires a fixed bandwidth during the connection lifetime of an application. In this case, all cells are subjected to the same delay, if any, and the delay variations across different cells are virtually zero. Hence, tightly constrained CTD and CDV are required. In addition to these, a user specifies the PCR, CDVT (cell delay variation tolerance), and CLR for CBR connections, yet most public ATM providers allow users to establish the PCR with all other parameters at a default setting. The required bandwidth is determined by the

**peak cell rate** (PCR) specified by the user, [14] and as long as the source conforms to the negotiated PCR, the QoS is guaranteed. Congestion control is achieved generally by means of the **call admission control** (CAC) at connection times. In other words, the network accepts a call request only if it can provide the requested bandwidth. Examples of CBR applications are delay-sensitive, real-time applications such as interactive voice, video conferencing, fixed bit rate coded video and **circuit emulation services** (CES).

- **Real-time variable bit rate** (rt-VBR): This service category is intended for real-time bursty traffic with variable bandwidths, which is characterized by low CTD, low CDV, and low CLR. It can usually tolerate statistical multiplexing with traffic from other sources. The user specifies PCR, CDVT, SCR, MBS, CTD, and CLR to the network at call setup time. Still, the bandwidth allocation is based on the SCR and possibly PCR and MBS and other parameters generally have default values. Applications using this service category include packetized voice or video.

- **Non-real-time variable bit rate** (nrt-VBR): Like rt-VBR, this service category also requires variable bandwidths. However, because of its non-real-time nature, there are no tight constraints on CTD or CDV requirements, but still it is expected to keep the cell loss ratio low. It permits statistical multiplexing with traffic from other sources. The parameter specifications are the same as in rt-VBR and the network determines its bandwidth requirement on the basis of the PCR, SCR, and MBS. Services such as traditional file transfer, email, and frame relay interworking are in this category.

- **Available bit rate** (ABR): This service category is intended for sources that can reduce or increase their information rate if the network requires them to do so. Thus, this allows them to exploit the available bandwidth produced by changes in the ATM layer transfer characteristics after connection setup. With this service category, an end-point specifies the upper and lower limits of its bandwidth requirement to the network by using PCR and **minimum cell rate** (MCR), which is a guaranteed transmission rate by the network. However, most VCs use zero as a default MCR value, because if an MCR of an ABR application is higher, the connection might be rejected when sufficient bandwidth is not available. A traffic rate control is performed by a closed-loop feedback control using *resource management (RM)* cells, which have been standardized by the ATM Forum. With this scheme, the bandwidth assigned to a user may vary depending on the congestion experienced by the network. Therefore, this service works for non-real-time traffic with no commitments for delay.

- **Unspecified bit rate** (UBR): This service category is intended for non-critical applications, which do not have tightly constrained requirements of delay and delay variation, nor

[14]This is indicated by the user at call request time via a traffic descriptor of the SETUP message.

a specified quality of service. In the service category, no performance parameters are set except the PCR which generally equals the line rate. The application types of UBR are similar to those of nrt-VBR, except that sources do not specify the usual traffic descriptors. Thus, the same applications that are suitable for nrt-VBR services can also be handled by the UBR service. UBR service supports a high degree of statistical multiplexing among sources. Since the user does not specify its cell rate, the network may reduce its bandwidth in the event of congestion. The congestion control may be provided by a switch at the ATM layer, but it could also be up to the upper layers of the protocol.

Table 2.10 is a summary of these service categories, along with their features.

Table 2.10: Summary of service categories and their features

| Service category | Network priority | Traffic descriptors | QoS parameters Delay/Variance | CLR | Burst tolerance | Control |
|---|---|---|---|---|---|---|
| CBR | 1 | PCR | small | small | none | CAC |
| rt-VBR | 2 | PCR, SCR, MBS | small | medium | some | CAC |
| nrt-VBR | 3 | PCR, SCR, MBS | n/a | medium | some | n/a |
| ABR | 4 | PCR, MCR | n/a | medium | n/a | RM cell |
| UBR | 5 | (PCR) | n/a | n/a | n/a | * |

Priority 1 is the highest priority. n/a stands for non-applicable. CAC stands for call admission control, which is performed at connection setup time. RM stands for resource management. Congestion control of UBR is performed in ATM layer or higher layers. Congestion controls of ABR and UBR will be described later.

## 2.5 ATM adaptation layer

### 2.5.1 AAL protocol model

ITU-T has defined the ATM adaptation layer (AAL). The use of ATM creates the need for an adaptation layer to support information transfer protocols not based on ATM. The B-ISDN protocol model adopts the AAL for this purpose. For this reason, the AAL is located between the ATM layer and the higher layers. Its basic function is the enhanced adaptation of services provided by the ATM layer to the requirements of the higher layer. On the contrary, PDUs from the higher layer are mapped into the information field of an ATM cell in the AAL. AAL entities exchange information with their peer AAL entities to support AAL functions. The structure and logical interfaces of the AAL are illustrated in Figure 2.33. PDUs from the ATM layer (the ATM cell payload) via ATM-SAP (ATM service access point) are passed to the AAL and they are provided to the higher layers through an AAL-SAP (AAL service access point), across which primitives regarding the AAL protocol data units (AAL-PDUs) are passed. The

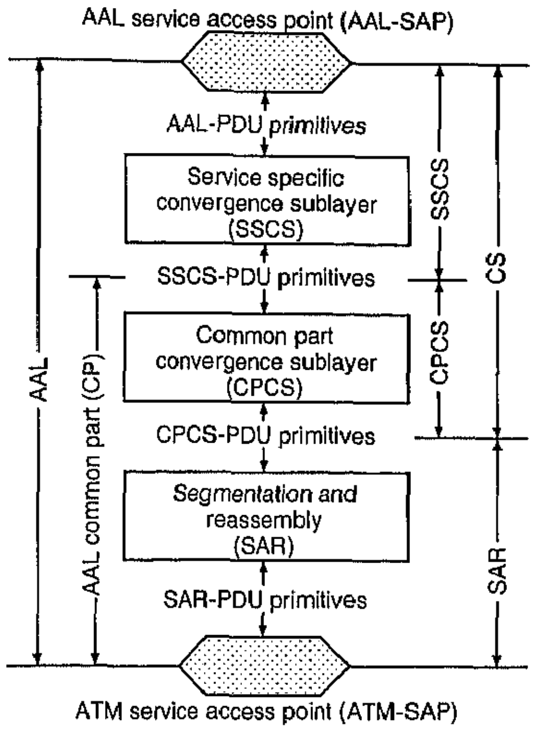


Figure 2.33: Generic AAL protocol sublayer model

AAL is divided into two sublayers: the **convergence sublayer** (CS) and the **segmentation and reassembly** (SAR) sublayer. The CS sublayer is further subdivided into **service specific** (SSCS) and **common part** (CPCS) components. While the CPCS must always be implemented along with the SAR sublayer, the SSCS sublayer may be **null**, which means that it need not be always implemented.

The function of the CS at the source is to divide very long user packets into fixed-length packets called **CS-service data units** (CS-SDUs). It may add header and/or trailer information to the CS-SDU to generate **CS-protocol data units** (CS-PDUs). Finally, it passes the CS-PDUs to the SAR as SAR-SDUs[15]. The function of the SAR sublayer is to package information from CS into cells for transmission and to unpack the information at the other side. As describe above, each ATM cell consists of a 48-octet information field. Thus, SAR must pack any SAR headers and trailers plus CS information into 48-octet blocks.

The ITU-T Recommendation has listed some general examples of services provided by the ATM adaptation layer [STAL95]. They are: i) the handling of transmission errors; ii) segmentation and reassembly to enable larger blocks of data to be carried in the information field of ATM cells; iii) the handling of lost and misinserted cell conditions; And iv) flow

[15] A PDU is a data packet containing user information and control information that is exchanged between two communicating peers in a network. An SDU is a PDU received from the layer directly above the current layer to which the current layer may add control information to form its own PDU.

control and timing control [STAL95]. In order to minimize the number of different AAL protocols, ITU-T has also defined the basic principles and classification of AAL functions [STAL95, HAND94, PRYC93, MCDY94]. This classification is based on whether the timing relationship is required between the source and the destination, whether the bit rate is constant or variable, and whether the transfer is connection-oriented or connectionless. Table 2.11 shows the four original AAL service classes. In Class A, a timing relation between the source and the

Table 2.11: Service classification for AAL

| | Service class | | | |
|---|---|---|---|---|
| | Class A | Class B | Class C | Class D |
| Timing relation between source and destination | Required | | Not required | |
| Bit rate | Constant | Variable | | |
| Connection mode | Connection oriented | | | Connectionless |
| AAL protocol | AAL 1 | AAL 2 | AAL 3/4, 5 | AAL 3/4 |
| Examples | DS1, E1<br>64 kbit/s emulation | Packet video,<br>packet audio | Frame relay,<br>X.25 | IP,<br>SMDS |

destination is required. The bit rate is constant and the service is connection-oriented. A typical example is a voice of 64 Kbit/s as in narrowband ISDN. The offering of this service is sometimes called circuit emulation. Another example is fixed bit rate video. In Class B, a timing relation exists again between the source and the destination for a connection-oriented service. However, in contrast to Class A sources, Class B sources have a variable bit rate. Typical examples are variable bit rate video and audio, as might be used in a teleconference. In the application, the connection mode is connection-oriented and timing is important, but the bit rate varies depending on the amount of activity on the scene. In Class C, there is no time relation between the source and the destination, and the bit rate is variable. Service is connection-oriented. Examples are connection-oriented data transfer such as X.25 and signaling. Class D differs from Class C in being connectionless. An example of such a service is connectionless data transport (e.g., internet protocol (IP) or switched multimegabit data services (SMDS)).

### 2.5.2 AAL definition

Initially, ITU-T defined one AAL protocol type for each service class. Thus AAL 1 through AAL 4 were directly mapped to Class A through Class D depending on their own features. AAL 1 has been defined by the ITU-T and further clarified in the ANSI for continuous bit rate (CBR) applications. On the other hand, for VBR, AAL 3 was initially developed for connection-oriented services and AAL 4 for connectionless services. As the AAL 3 and AAL 4 were common in structure and they were merged into a single type called "AAL 3/4". However, because of the complexity and implementation difficulties in the AAL 3/4, AAL 5 or simple efficient adaptation layer (SEAL) was proposed by the computer industry. AAL 5 has become the predominant AAL type in a great deal of data communications equipment. It is standardized for the transport of signaling messages as well. On the contrary, standardization procedures of initial AAL 2 was stopped. Instead, a new AAL (initially AAL 6) has been conceived for real-time voice transmission recently. It is now described in ITU-T Recommendation as new AAL 2. Thus, currently standardized AALs are:

- AAL 1: Constant bit rate (CBR) traffic
- AAL 2: Variable bit rate (VBR) traffic for Voice over ATM (VOA)
- AAL 3/4: VBR traffic
- AAL 5: Lightweight VBR traffic

Each AAL type is detailed next section.

### 2.5.3 AAL type 0

There are some applications that no AAL functionality is required and the content of the cell information field is directly and transparently transferred to the higher layer. This is unofficially called "AAL type 0", which has empty SAR and CS protocols. Although a detailed description of such a service is not available in the ITU-T standards, it can be considered very important for the prevalence of ATM.

### 2.5.4 AAL type 1

Normally, constant bit rate (CBR) services (Class A) use AAL type 1 because they require information to be transferred between the source and the destination at a constant bit rate after a virtual connection has been set up. The layer services provided by the AAL type 1 to the AAL user are:

- Transfer of service data units (SDUs) with constant bit rate.
- Transfer of timing information between source and destination.
- Transfer of information about the data structure if necessary.
- Indication of lost or erroneous information not recovered within the AAL type 1.

A number of error indications, such as corrupted user information, loss of timing, buffer overflow, and buffer underflow may be passed from the user plane to the management plane.

The functions performed by the AAL 1 are as follows:

- Segmentation and reassembly of user information
- Handling of cell delay variation
- Handling of cell payload assembly delay
- Handling of lost and misinserted cells
- Source clock frequency recovery at the receiver
- Recovery of the source data structure at the receiver
- Monitoring and handling of AAL-PCI (protocol control information) for bit errors
- Monitoring of the user information field for bit errors and possible corrective action

In the case of circuit emulation, the monitoring of the end-to-end QoS is necessary. Therefore, a cyclic redundancy check (CRC) may be calculated at the CS for the information carried in one or more cells. The result is transferred to the receiver within the information field of a cell or in a special OAM cell.

Circuit emulation is an important feature of B-ISDN as it allows existing circuit-based signals (e.g., 1.5 Mbit/s or 2 Mbit/s), thereby satisfying the requirements such as delay, jitter, or bit error rate.

### Segmentation and reassembly sublayer.

The SAR functions are as follows.

- Mapping between CS PDU and SAR PDU
- Indicating existence of the CS function
- Sequence numbering
- Error protection

The SAR sublayer accepts a 47-octet block of data from the CS sublayer and then adds a one octet PCI or SAR-PDU header to each block to form an SAR-PDU. At the destination, the SAR sublayer gets a 48-octet block from the ATM layer, and then, after separating the SAR-PDU header, the remaining 47-octet block is passed to the CS.

Figure 2.34 illustrates the SAR-PDU format of AAL type 1. The PCI is subdivided into a four bit sequence number (SN) field and a four bit sequence number protection (SNP) field. The SN is composed of a convergence sublayer indication (CSI) bit and a three bit sequence count field. The SNP field contains a three bit CRC and an even parity bit which has to be calculated over the resulting seven bit code word.

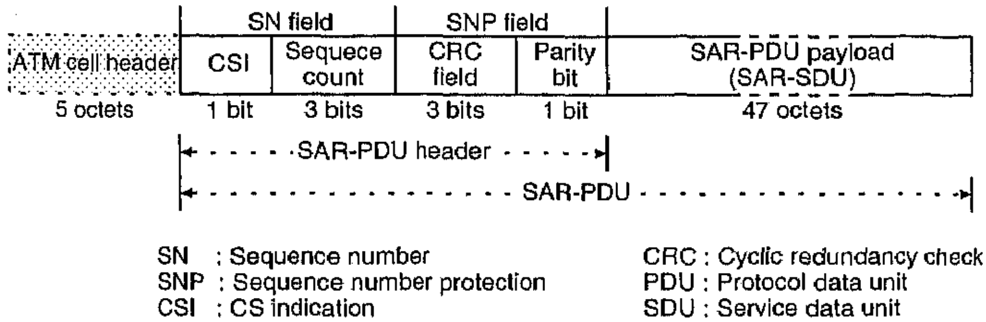


Figure 2.34: SAR structure for AAL type 1

In AAL 1, specifically, there are two transfer modes: the *unstructured data transfer mode* and the *structured data transfer mode*. The former mode uses point-to-point VCs to emulate full T1/E1 connection across an ATM network. In this mode, the user data is seen as one continuous bit stream in which no internal structure is needed for the data. This is equivalent to monopolizing a channel with one application, in which some time slots remain unused. The latter mode contains information which is internally structured by means of internal byte-aligned blocks or internal framing bit patterns. The structured data transfer mode is similar to multiplexing more than one applications in a time-division basis over T1/E1 channels. In this case, framing information is needed at the beginning and end of each structured block. In some applications where the structured data is transferred, one octet of CS-PDU is used as a pointer, which denotes the offset measured in octets, to the start of the structured block consisting of the remaining 46 octets of this PDU and the 47 octets of the next PDU. Thus, it takes values 0-92 and identifies the boundary of the next structure block. The cell which has the pointer is called a "p-format cell", and the cell which does not have the pointer is called a "non-p format cell". Since p-format cell appears only every other cell, the offset value must allow for a structure block beginning on odd-numbered sequence counts. Because there are 46 octets in a p-format cell and 47 in a non-p format cell, the pointer value covers the range of offset values that will occur. The CSI bit is used to indicate the existence of an eight-bit pointer field. CSI=1 indicates the pointer is present, and CSI=0 indicates the pointer is not present. The example of a structured data transfer SAR-PDU and an unstructured data transfer SAR-PDU are shown in Figure 2.35.

Associated with each 47 octet SAR-PDU, the SAR sublayer receives a sequence count value in the SN field (Figure 2.34) from the CS, which is delivered to a peer CS at the receiving end to be used to detect the loss or misinsertion of cells. For systems with high cell loss ratios this method is not very robust since the three bit sequence count field is relatively short.

The data integrity of the SN field (the CSI and the sequence counter) is protected against bit errors by a four bit sequence number protection field. The CRC is calculated by the polynomial $G(x) = x^3 + x + 1$ and is inserted into the CRC field. The resulting seven bit code word is protected by an even parity check (parity bit field).

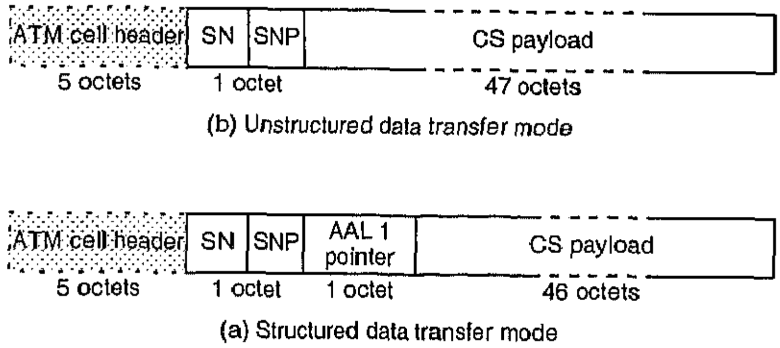


Figure 2.35: Examples of different AAL 1 data transfer mode

Convergence sublayer.

The functions of the CS depend on each service to be supported. They include:

- Handling of **cell delay variation** (CDV): A buffer is used to support this function. Buffer underflow or overflow may lead to the insertion of dummy bits or the dropping of excessive bits, respectively.

- Handling of lost and misinserted cells.

- Source clock recovery at the receiver: For some services, clock recovery at the receiver is needed. Several methods exist for this purpose. The **Synchronous residual time stamp** (SRTS) method is recommended by ITU-T. It uses a **residual time stamp** (RTS) to measure and convey information to the receiver about the frequency difference between a common reference clock derived from the network clock at both the sender and the receiver and the service clock of the sender. The four bit RTS is transferred by the CSI bit in successive SAR-PDU headers with an odd number of the sequence count field (1, 3, 5, 7). A common reference clock is available if both the sender and the receiver communicate via a synchronous network (e.g., an SONET/SDH based network). This method is capable of meeting jitter requirements specified by the ITU-T for 2.048 and 1.544 Mbit/s hierarchies. If a common reference clock is unavailable (e.g., on a PDH based network), an adaptive clock recovery method based on monitoring the buffer filling level at the receiver, may be used.

- Transfer of structure information: A pointer is used for the delineation of structure boundaries as described in 2.5.4. For p-format cells, the CSI value in SAR-PDU headers with an even sequence number (0, 2, 4, 6) is set to 1.

- **Forward error correction** (FEC) for high-quality video and audio: This may be combined with bit interleaving to provide more secure protection against errors. One example is the Reed-Solomon code for unidirectional video services. This method increases the overhead by 3.1% and the introduced delay is 128 cell cycles [HAND94].

- Report end-to-end performance status.

### 2.5.5 AAL type 2

The standardization procedures of initial AAL 2 recommended in I.363 has been withdrawn once. Instead, a new AAL (initially AAL 6) has been conceived recently. It is now described as a new AAL 2.

When a real-time voice application is serviced over an ATM with standard ATM cells, these cells should be partially (or mostly) empty since the actual data rate is very low. In addition, if a voice message suppressed to, for instance, 8 Kbit/s is converted to a cell of the 48 octets payload in an ATM standard cell, it takes quite a long delay of 48 ms in a single cell assembly process ($48octet \times 8[bit/octet]/8[Kbit/s] = 48\ ms$). In order to support such a service, no existing AAL choices are efficient for real-time data that has a requirement for low latency and a low bandwidth. The AAL type 2 satisfies the requirements for the bandwidth-efficient transmission of low-rate, short, and variable length packets in delay sensitive applications.

AAL type 2 provides a 3-octet header and a short variable length payload, and data from multiple users can be multiplexed to a single ATM connection (Figure 2.36).

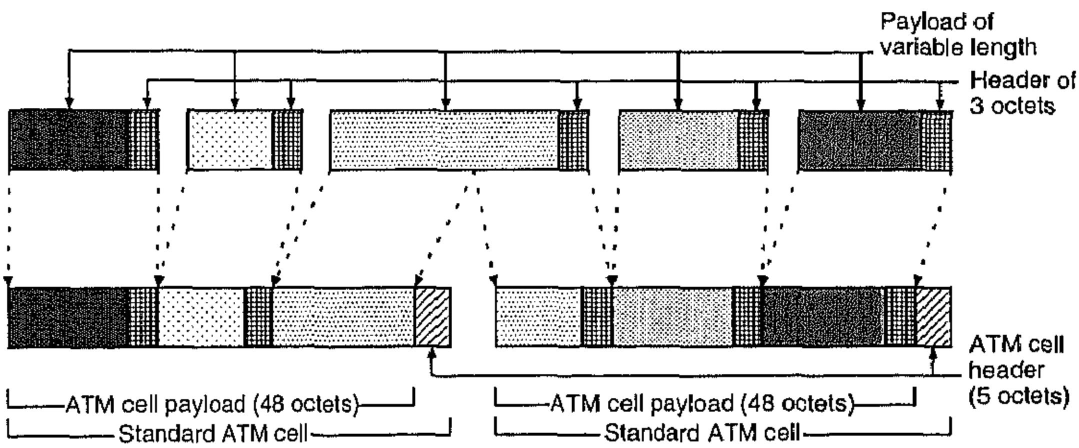


Figure 2.36: AAL type 2

### 2.5.6 AAL type 3/4

AAL 3 and AAL 4 are combined into a single common part AAL 3/4 to support VBR traffic. The ITU-T recommends the use of AAL 3/4 for the transfer of data which is sensitive to loss but not to delay. The AAL may be used for connection-oriented as well as for connectionless data communication. Support for connectionless service can be provided at the service specific convergence sublayer (SSCS) level. However, the AAL itself does not perform all functions required by a connectionless service, since functions like routing and network addressing are performed at the network layer.

The two modes of AAL 3/4 are defined: *message mode* and *streaming mode*.

- Message-mode service: A single AAL-SDU is transported in one CS-PDU or (optionally) more than one CS-PDUs, which may build one or more SAR-PDUs. Figure 2.37 shows

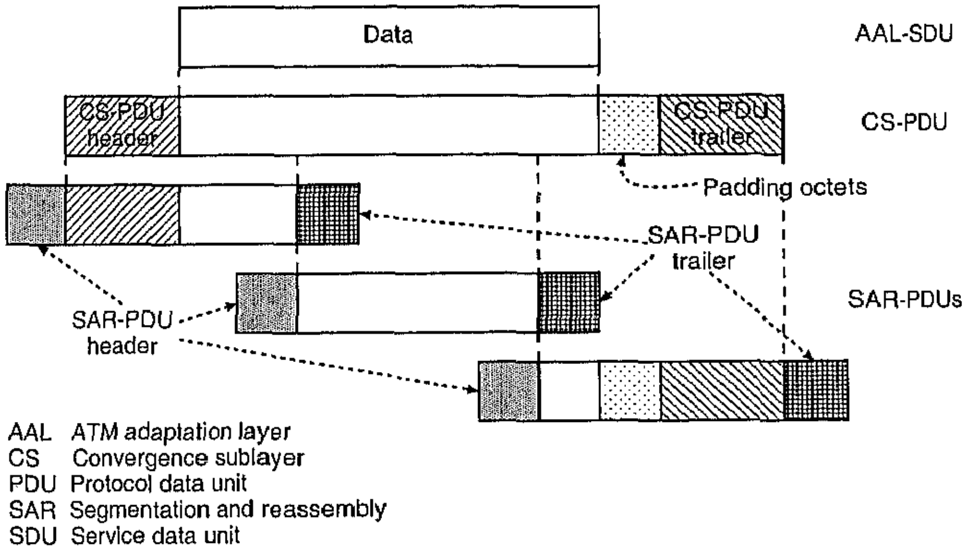


Figure 2.37: Message mode service

  the operation of this mode. This service can be used for framed data transfer of fixed or variable length AAL-SDUs.

- Streaming-mode service: One or more fixed-size AAL-SDUs are transported in one CS-PDU. The AAL-SDU may be as small as one octet and is always delivered as one unit, because only this unit will be recognized by the application (one SAR-SDU contains at most one AAL-SDU). The operation of this mode is illustrated in Figure 2.38.

In both modes, the SAR sublayer provides error detection and both these modes may offer the following peer-to-peer operational procedures:

- Assured operation: Retransmission of missing or erroneous AAL-SDUs are provided as a mandatory feature. In addition, flow control is supported between the endpoints. Flow control may be restricted to point-to-point connections at the ATM layer.

- Non-assured operation: No retransmission of missing or erroneous SAR-PDUs is provided. Optionally erroneous PDUs are delivered to the user. Flow control can be provided for point-to-point connections, but no flow control is provided for point-to-multipoint ATM layer connections.

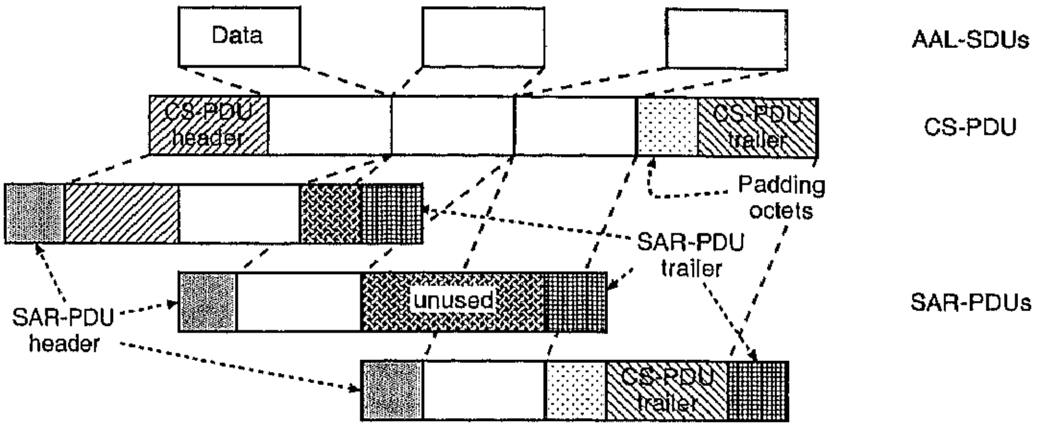


AAL ATM adaptation layer
CS Convergence sublayer
PDU Protocol data unit
SAR Segmentation and reassembly
SDU Service data unit

Figure 2.38: Streaming mode service

**Segmentation and reassembly sublayer.**

The SAR structure is shown in Figure 2.39. Four octets (two for the SAR-PDU header and 2

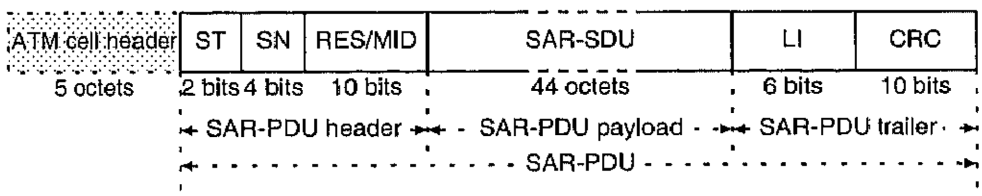


ST : Segment type
SN : Sequence number
RES : Reserved (type 3)
MID : Multiplexing identifier (type 4)
LI : Length indicator
CRC : Cyclic redundancy check

Figure 2.39: SAR structure for AAL 3/4

for the trailer) of protocol control information are appended. The SAR provides the following functions:

- Segmentation and reassembly of variable length CS-PDUs: The SAR-PDU contains two fields for this purpose:
  - **Segment type** (ST) field (two bits): The ST identifies the segment type of the SAR-PDU: 10 for the beginning of the message (BOM); 00 for continuation of message (COM), 01 for end of the message (EOM); and 11 for a single segment message (SSM).
  - **Length indicator** (LI) field (six bits): The LI field contains the number of valid octets in the SAR-PDU payload field with a maximum of 44 octets. Thus, the last

segment and single segment SAR-PDUs may contain a value of less than 44. An LI value of 63 is associated with an ST indicating an EOM that leads to an abortion of partially transmitted CS-SDUs at the receiver.

- Error detection and sequence integrity: This function includes detecting bit errors in the SAR-PDU as well as detecting lost or misinserted SAR-PDUs. To detect bit errors in the SAR-PDU, a cyclic redundancy check (CRC) field of 10 bits is defined, which is based on the generating polynomial $G(x) = x^{10} + x^9 + x^5 + x^4 + x + 1$. In addition, four bits of the sequence number (SN) field are used for detection of lost or misinserted cells. This sequence number is incremented by one relative to the SN of the previous SAR-PDU belonging to the same AAL connection (numbering modulo 16).

- Multiplexing/demultiplexing of multiple CS-PDUs: Multiplexing is supported by a 10 bit multiplexing identifier (MID) in the SAR-PDU. The use of the MID allows to multiplex and interleave $2^{10}$ AAL-user-to-AAL-user connections on a single user-to-user ATM layer connection for connection-oriented data communication. SAR-PDUs with an identical MID value belong to a particular CS-PDU. For connectionless data communication, the MID field is used to interleave SAR-PDUs from (up to $2^{10}$) different CS-PDUs on the same semi-permanent ATM layer virtual connection. If multiple AAL connections use the same ATM layer connection, these AAL connections must have identical QoS characteristics. Multiplexing/demultiplexing is done on an end-to-end basis. An ATM layer connection which is used by different AAL connections is administered as a single entity.

### Convergence sublayer

The common part convergence sublayer (CPCS) transfers user data frames with any length between 1 and 65535 octets. CPCS connections are established by the management or the control plane. One or more CPCS connections may be established between two peer CPCS entities. As the CPCS provides non-assured service, the integrity of the CPCS-SDU sequence must be guaranteed on each CPCS connection. The sublayer performs the following functions:

- Error detection and handling: Corrupted CPCS-SDUs are either discarded or optionally delivered to the service specific convergence sublayer (SSCS). Detected errors means errors detected at the CPCS layer as well as those detected at the SAR layer.

- Buffer allocation size: Each CPCS-PDU carries an indication of maximum buffer requirements to the receiving peer entity to receive the CPCS-PDU.

These functions are implemented by the CPCS-PDU of AAL 3/4 illustrated in Figure 2.40.

- Common part indicator (CPI) field: This field is used to interpret the remaining fields in the CPCS-PDU header and trailer. Currently, it indicates the counting units (octets when $CPI = 0$) for the values specified in the *BASize* and *Length* fields.

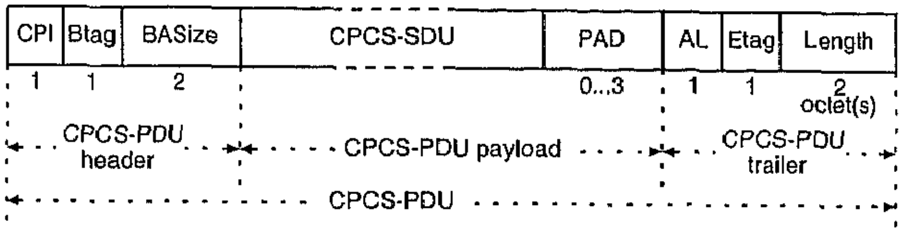


| | | | |
|---|---|---|---|
| CPI | Common part Indicator | PAD | Padding |
| Btag | Beginning tag | AL | Alignment |
| BASize | Buffer allocation size | Etag | Ending tag |
| CPCS | Common part convergence sublayer | Length | Length of CPCS-PDU payload |
| SDU | Service data unit | PDU | Protocol data unit |

Figure 2.40: CPCS-PDU format for AAL type 3/4

- Beginning tag (Btag) and Ending tag (Etag) field: These fields allow the proper association of the CPCS-PDU header and trailer at the receiver. The same numerical value is set into both fields by the sender for a given CPCS-PDU, which is changed (e.g. incremented) for each successive CPCS-PDU.

- Buffer allocation size (BASize) field: This field indicates to the receiving peer entity how much buffer space should be reserved to reassemble the CPCS-PDU before the data arrives. Its unit is indicated by the CPI. In message mode, the BASize is encoded equal to the CPCS-PDU payload length. In streaming mode, it is encoded equal to or greater than the CPCS-PDU payload length.

- Padding (PAD) field: The field ensures that the CPCS-PDU payload is an integer multiple of four octets. It may be 0 to 3 octets long and does not convey any information.

- Alignment (AL) field: This field is only used for 32 bit alignment of the CPCS-PDU trailer. It does not convey any information again, and shall be set to 0.

- Length field: This field is used to encode the length of the CPCS-PDU payload field, whose unit is indicated by the CPI value. It is also used by the receiver to detect the loss or gain of information.

### AAL 3/4 multiplexing example

Figure 2.41 illustrates a terminal that has two inputs with two 98-octet packets arriving simultaneously destined for a single ATM output port using the AAL 3/4 protocol. Two parallel instances of the CPCS sublayer encapsulate the packets with a header and trailer. These CPCS-PDUs are passed to the two parallel segmentation and reassembly (SAR) instances and are segmented into SAR-PDUs by properly arranging MIDs and segment types (BOM, COM and EOM). The processing of two packets occurs in parallel and the ATM cells are interleaved

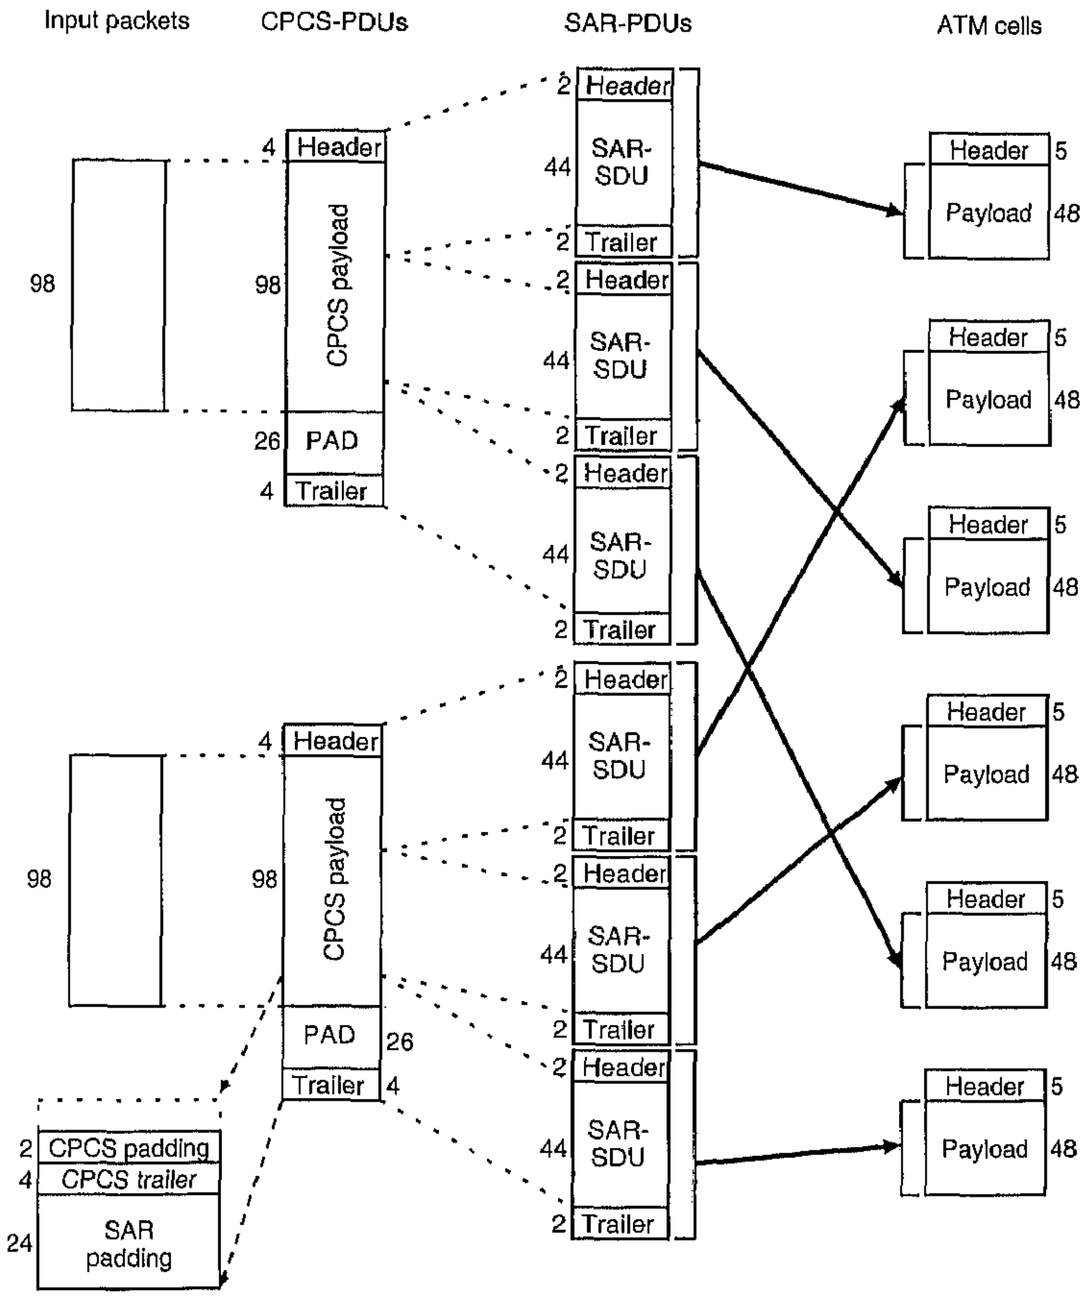


Figure 2.41: Multiplexing example using AAL 3/4

on the output port. This is a unique function to AAL 3/4 over AAL 5, but it also introduces an additional complexity.

### 2.5.7 AAL type 5

The AAL 3/4 has a high overhead of 4 octets per SAR-PDU of 48 octets. In addition, the CRC bits for detecting corrupted segments and the sequence number for detecting lost and misinserted segments may not be enough for protection of very long blocks of data. Therefore, the ATM Forum has proposed a new type of AAL, called "AAL type 5", for variable bit rate sources without a timing relation between a source and destination. The objective is to offer a service with less overhead and better error detection below the CPCS layer. At this layer, its message mode service, streaming mode service, and assured/non-assured operation are identical to the service provided by the CPCS of AAL 3/4. However, no multiplexing is supported by AAL 5. If multiplexing is required at the AAL layer, it will occur in the SSCS layer. In addition to those describe services, AAL type 5 will be used for signaling and frame relay over ATM.

#### Segmentation and reassembly sublayer

The SAR sublayer accepts variable length SAR-SDUs which are integer multiples of 48 octets from the CPCS. No more fields are additionally appended to the received SDUs at the SAR sublayer. It only generates SAR-PDUs containing 48 octets of SAR data. The preservation or delineation of the SAR-SDU happens by an end of SAR-SDU indication, which is carried by an AAL_indication bit in the PTI field in the ATM header. A value of 1 means the end of an SAR-SDU, while a value of 0 means the beginning or continuation of an SAR-SDU. Thus, a segment type field is not needed. Moreover, this makes the reassembly design simpler and makes more efficient use of ATM bandwidth[16]. This AAL type 5, however, makes use of the ATM cell header information on the ATM layer, which can be considered as an infringement of the protocol reference model (PRM) for ATM. Nevertheless, it is adopted because of its simplicity and efficiency, and this infringement is a key to the congestion control schemes described later.

#### Convergence sublayer

The functions implemented by the AAL type 5 CPCS are almost the same as those offered by the AAL 3/4 CPCS. However, the AAL 5 CPCS does not have a BASize indication function. In addition, error protection in the AAL 5 is handled at this CPCS layer, instead of being shared between SAR and CPCS as in AAL 3/4. Figure 2.42 depicts the format of the CPCS-PDU for AAL 5. The AAL 5 CPCS-PDU payload may be any integer length from 1 to 65535 ($= 2^{16} - 1$) octets. The padding (PAD) field provides for a 48 octet alignment of the CPCS-PDU so that it can be directly segmented into ATM cell payloads. The CPCS user-to-user indication

[16]This is because the original AAL 5 was called the **simple efficient adaptation layer** (SEAL).

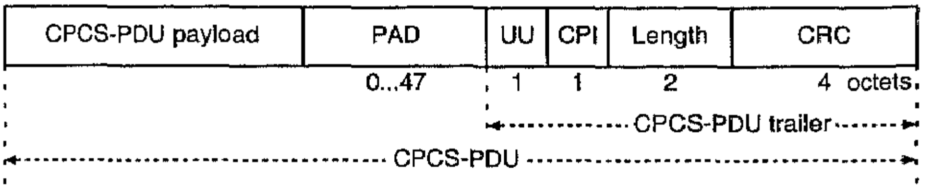


| | | | |
|---|---|---|---|
| CPCS | Common part convergence sublayer | CPI | Common part indicator |
| PDU | Protocol data unit | Length | Length of CPCS-PDU payload |
| PAD | Padding | CRC | Cyclic redundancy check |
| UU | CPCS user-to-user indication | | |

Figure 2.42: *CPCS-PDU format for AAL type 5*

(UU) field is used for the transparent transfer between AAL users. The common part indicator (CPI) functions similar to that of AAL 3/4. It was merely used to align the trailer to a 64-bit boundary. The length field identifies the length of the CPCS-PDU payload so that the PAD can be removed. It is also used by the receiver to detect the loss or gain of information. As 16 bits are allocated to the length field, the maximum payload length is $2^{16} - 1 = 65535$ octets. If the length field is set to zero, the partially transmitted CPCS-SDUs are aborted. The CRC-32 field is a new field. It detects errors in the CPCS-PDU by a calculation performed over the entire contents of the CPCS-PDU by the generating polynomial:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

#### AAL 5 multiplexing example

Figure 2.43 shows the same example previously used for the AAL 3/4(Figure 2.41) to illustrate a major difference in AAL 5. Two 98-octet packets arrive at a terminal simultaneously. They are destined for a single ATM output port using the AAL 5 protocol this time. A trailer having CPI is appended to each packet by the two parallel instances of the CPCS sublayer. The two parallel SAR processes segment the CPCS-PDU into SAR-SDU and construct ATM cells. In this protocol, no buffer allocation size (BASize) field is needed to be received beforehand. Neither are segment types needed. In this example, as the packets are destined for the same VPI/VCI, only one cell can be sent at a time. This implementation is simpler than the AAL 3/4. But it is unable to keep the link as fully occupied as the AAL 3/4 which multiplexes several packets, when the packets arrive much faster than the rate at which SAR and ATM cell transmission occur.

## 2.6 Quality of service

ATM should support different types of service with different requirements of performance levels. This performance level of a connection is represented by the QoS parameters. In addition,

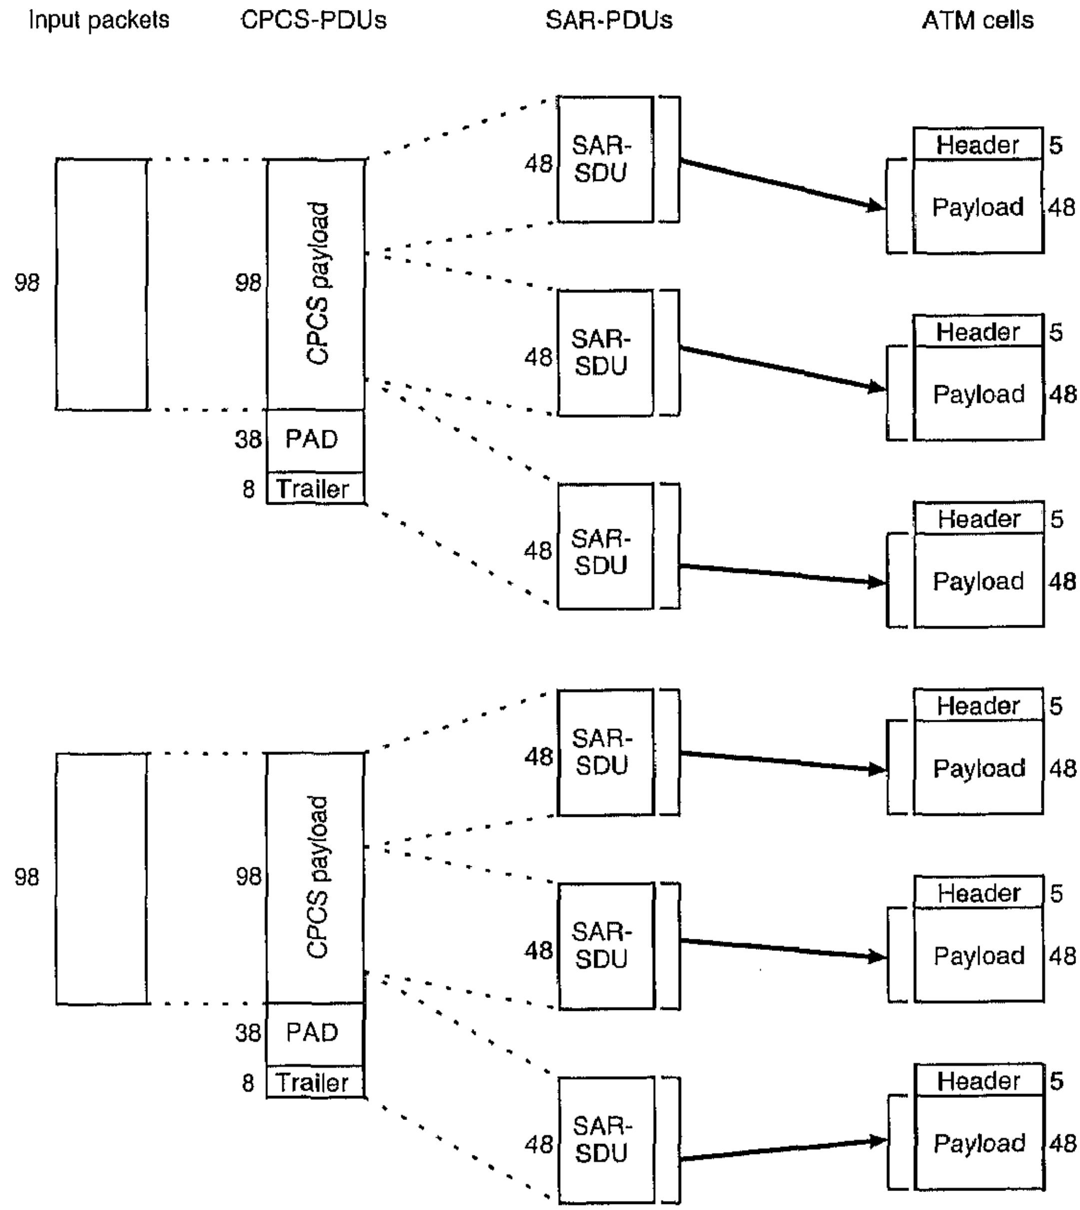


Figure 2.43: Multiplexing example using AAL 5

the QoS requirements is one of the factors to decide the service category of a class of ATM connections.

Simply, **quality of service** (QoS) refers to a set of user-perceivable ATM performance parameters that characterize the traffic over a given virtual connection. Each connection has a unique QoS within the network, which determines how the network treats the cells. As the ATM networks are needed to carry a variety of traffic types, such as voice, video, and data, different quality of services are required by different connections. This QoS negotiation is achieved by exchanging a set of QoS parameters between the user and the network.

Even if a number of QoS parameters have been defined, they are not always used. The QoS parameters mainly used are as follows [GADE97]:

**Cell delay variation (CDV)**: The difference between the actual transfer delay of an arbitrary cell and the expected transfer delay of that cell. Buffering, scheduling and multiplexing of a cell can introduce variable delays in the cell stream in both the CPE and network, making this CDV necessary. Additionally, cell delay variation tolerance (CDVT) represents the maximum allowed value for CDV on a connection.

**Cell transfer delay (CTD)** : The average time delay for a cell to be transferred from its source to its destination over a virtual connection as measured on an end-to-end basis.

**Cell loss ratio (CLR)** : The ratio of lost cells to the total number of transmitted cells. It is the allowed percentage of cells which can be lost in the network, which is measured on an end-to-end basis.

**Burst tolerance** : The maximum length of time that the user can transfer at PCR. If a user sends traffic for the full burst tolerance, the cell rate must be decreased to meet the requirements of the SCR parameter some time. Burst tolerance can be used as roughly the same as **maximum burst size** (MBS) measured in number of cells.

**Peak cell rate (PCR)** : The maximum number of cells per second the connection can transfer into the network. The PCR is often set to the maximum number possible for the given line rate, but it cannot be set higher than the line capacity.

**Sustained cell rate (SCR)** : The average number of cells per second that the connection can transfer into the network. The burst tolerance determines the length of time over which the network measures this average.

**Minimum cell rate (MCR)** : The smallest cell transfer rate that the connection must always support.

Table 2.12: Relation of applications and the QoS parameters

| Services | Characteristics | QoS Parameters |
|---|---|---|
| Voice | Sensitive to delays<br>Tolerant to random,<br>but not bursty bit errors<br>Fixed bandwidth for CBR-encoded audio<br>Variable bandwidth for VBR-coder outputs | Low CDVs and<br>low CLRs |
| Interactive video | Sensitive to delays<br>Tolerant to random,<br>but not bursty bit errors<br>Fixed or variable bandwidth<br>depending on the coding used | |
| Interactive data | Variable bandwidth<br>Tolerant to delays or delay variations<br>No errors | A reasonable delays<br>and/or CDVs,<br>no CLR |
| Variable data transfer | Tolerant to delays and delay variations<br>No errors | |

How to use these parameters differs according to service categories. Each service category uses a subset of the QoS parameters to achieve a required service quality. Table 2.12 shows the characteristic applications and the corresponding QoS parameters [KARI00].