Section 3:

ATM Layer Specification

3.1 ATM Layer Services

The ATM layer provides for the transparent transfer of fixed size ATM layer Service Data Units (ATM-SDUs) between communicating upper layer entities (e.g., AAL-entities). This transfer occurs on a pre-established ATM connection according to a traffic contract. A traffic contract is comprised of a QoS class, a vector of traffic parameters, a conformance definition and other items as specified in section 3.6. Each ATM end-point is expected to generate traffic which conforms to these parameters. Enforcement of the traffic contract is optional at the Private UNI. The Public Network is expected to monitor the offered load and enforce the traffic contract.

Two levels of virtual connections can be supported at the UNI:

- A point-to-point or point-to-multipoint Virtual Channel Connection (VCC) which consists of a single connection established between two ATM VCC end-points.
- A point-to-point or point-to-multipoint Virtual Path Connection (VPC) which consist of a bundle of VCCs carried transparently between two ATM VPC end-points.

Note: For VPC at the Public UNI, traffic monitoring and throughput enforcement will be performed across all cells carried on the same VPI independently of the VCI values.

(R) From a single source the relay of cells within a VPC/VCC must preserve cell sequence integrity.

No retransmission of lost or corrupted information is performed by this layer. Flow control over ATM connections is for further study. The ATM layer also provides its users with the capability to indicate the loss priority of the data carried in each cell. The information exchanged between the ATM layer and the upper layer (e.g., the AAL) across the ATM-SAP includes the following primitives:

Primitive	Request	Indicate	Confirm	Respond
ATM-DATA	**X**	**X**		

Figure 3-1 ATM Service Access Point (SAP) Primitives

These primitives make use of the following parameters:

Parameter	Associated Primitives	Meaning	Valid values
ATM-SDU	ATM-DATA.request ATM-DATA.indication	48 byte pattern for transport	Any 48 byte pattern
SDU-type	ATM-DATA.request ATM-DATA.indication	End-to-end cell type indicator	0 or 1
Submitted Loss-priority	ATM-DATA.request	Requested Cell Loss-priority	High or Low priority
Received Loss-priority	ATM-DATA.indication	Received Cell Loss Priority	High or Low priority
Congestion-experienced	ATM-DATA.indication	EFCN indication	True or False

Figure 3-2 ATM-SAP Parameters

The primitives provide the following services:

ATM-DATA.request: Initiates the transfer of an ATM-SDU and its associated SDU-type to its peer entity over an existing connection. The loss priority parameter and the SDU-type parameter are used to assign the proper CLP and PTI fields to the corresponding ATM-PDU generated at the ATM layer.

ATM-DATA.indication: Indicates the arrival of an ATM-SDU over an existing connection, along with a congestion indication and the received ATM-SDU type. In the absence of errors, the ATM-SDU is the same as the ATM-SDU sent by the corresponding remote peer upper layer entity in an ATM-DATA.request.

The following parameters are passed within one or more of the previous primitives:

ATM-SDU: This parameter contains 48 bytes of ATM layer user data to be transferred by the ATM layer between peer communicating upper layer entities.

Submitted Loss Priority: This parameter indicates the relative importance of the requested transport for the information carried in the ATM-SDU. It can take only two values, one for high priority and the other for low priority.

Received Loss Priority: This parameter indicates the relative importance of the transport given to the information carried in the ATM-SDU. It can take only two values, one for high priority and the other for low priority.

Congestion indication: This parameter indicates that the received ATM-SDU has passed through one or more network nodes experiencing congestion.

SDU-type: This parameter is only used by the ATM layer user to differentiate two types of ATM-SDUs associated with an ATM connection.

3.2 Service Expected from the Physical Layer

The ATM layer expects the Physical layer to provide for the transport of ATM cells between communicating ATM-entities. The information exchanged between the ATM layer and the Physical layer across the PHY-SAP includes the following primitives:

Primitive	Request	Indicate	Confirm	Respond
PHY-UNITDATA[1]	X	X		

1: The ATM-entity passes one cell per PHY-UNITDATA.request and accepts one cell per PHY-UNITDATA.indicate.

Figure 3-3 PHY-SAP Services Required by the ATM Layer

3.3 ATM Cell Structure and Encoding at the UNI

(R) Equipment supporting the UNI shall encode and transmit cells according to the structure and field encoding convention defined in T1 LB310 [7]. (see Figure 3-4 and Figure 3-5)

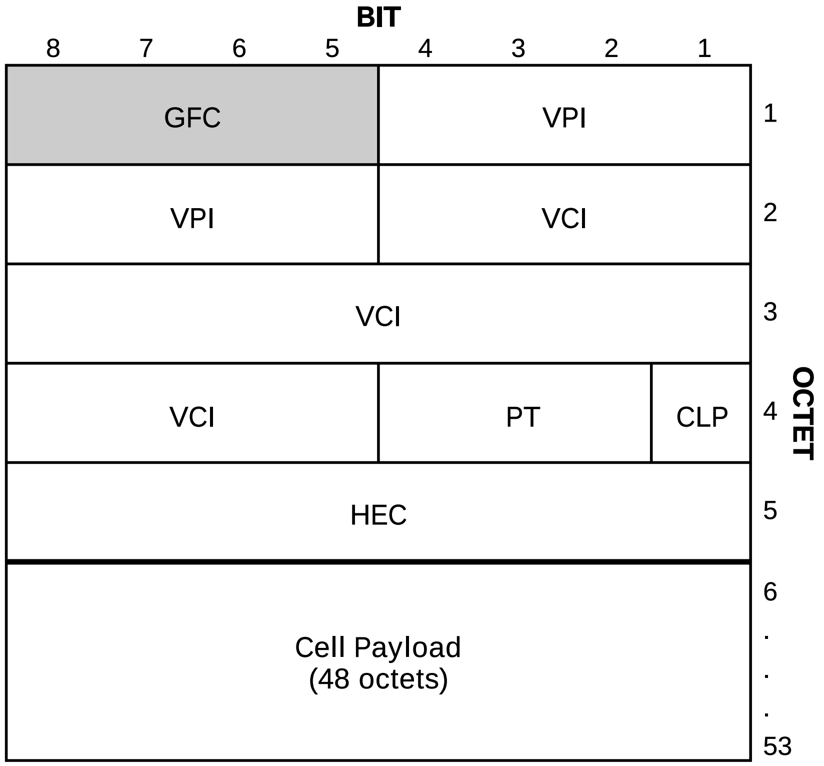

GFC: General Flow Control
VCI: Virtual Channel Identifier
CLP: Cell Loss Priority
VPI: Virtual Path Identifier
PT: Payload Type
HEC: Header Error Check

Figure 3-4 ATM Cell Structure at the UNI

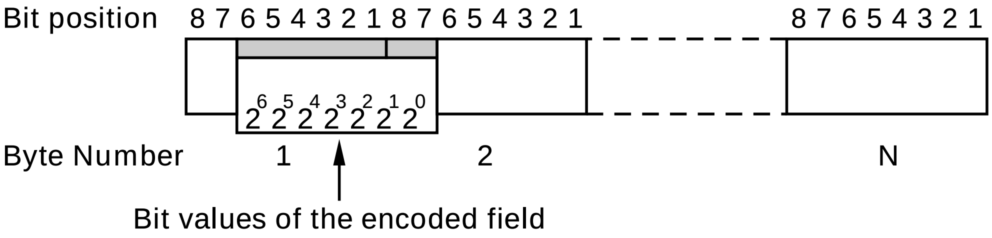

Figure 3-5 ATM Field Encoding Covention

The structure of the ATM cell is shown in Figure 3-4. It contains the following fields:

Generic Flow Control (GFC): This field has local significance only and can be use to provide standardized local functions (e.g. flow control) on the customer site. The value encoded in the GFC is not carried end-to-end and will be overwritten by the ATM switches.

Two modes of operation have been defined for operation of the GFC field. These are“uncontrolled access” and “controlled access”. The “uncontrolled access” mode of operation is used in early ATM environment. This mode has no impact on the traffic which a host generates. Each host transmits the GFC field set to all zeros (0000). In order to avoid unwanted interactions between this mode and the “controlled access” mode where hosts are expected to modify their transmissions according to the activity of the GFC field, it is required that all CPE and public network equipment monitor the GFC field to ensure the attached equipment is operating in “uncontrolled mode”. A count of the number of non-zero GFC fields should be measured for non-overlapping intervals of 30,000 +/- 10,000 cell times. If ten (10) or more non-zero values are received within this interval, an error is indicated to Layer Management.

(R) CPE at the UNI shall encode the GFC value to all zeros (0000).

(R) Public network equipment at the public UNI shall encode the GFC value to all zeros (0000).

(O) CPE shall inform Layer Management if a count of the non-zero GFC fields measured for non-overlapping intervals of 30,000 +/- 10,000 cell times reaches ten (10) or more.

(O) Public network equipment shall inform Layer Management if a count of the non-zero GFC fields measured for non-overlapping intervals of 30,000 +/- 10,000 cell times reaches ten (10) or more.

Virtual Path/Virtual Channel (VPI/VCI) Identifier: The actual number of routing bits in the VPI and VCI subfields use for routing is negotiated between the user and the network, e.g. on a subscription basis. This number is determined on the basis of the lower requirement of the user or the network.

Note: The number of VCI routing bits used in a user-to-user VP is negotiated between the users of the VP.

(R) The bits within the VPI and VCI fields used for routing are allocated using the following rules:

- The allocated bits of the VPI subfield shall be contiguous;
- The allocated bits of the VPI subfield shall be the least significant bits of the VPI subfield, beginning at bit 5 of octet 2;

- The allocated bits of the VCI subfield shall be contiguous;
- The allocated bits of the VCI subfield shall be the least significant bits of the VCI subfield, beginning at bit 5 of octet 4;

(R) Any bits of the VPI subfield that are not allocated are set to 0. For a given VP, any bits of the VCI subfield that are not allocated are set to 0.

Payload Type (PT): This is a 3-bit field used to indicate whether the cell contains user information or Connection Associated Layer Management information (F5 flow). It is also used to indicate a network congestion state or for network resource management. The detailed coding and use of the PT field will be described in section 3.4.4

Cell Loss Priority (CLP): This is a 1-bit field which allows the user or the network to optionally indicate the explicit loss priority of the cell. More details on the use of the CLP bit are given in section 3.4.5.

Header Error Control (HEC): The HEC field is used by the physical layer for detection/ correction of bit errors in the cell header. It may also be used for cell delineation.

3.4 ATM Layer Functions Involved at the UNI (U-plane)

This section describes ATM layer functions that need to be supported at the User-Network Interfaces (see Figure 3-6). It does not cover those ATM functions that are described in standards but have no impact on the UNI specification.

Functions	**Parameters**
Multiplexing among different ATM connections	**VPI/VCI**
Cell rate decoupling (unassigned cells)	**Pre-assigned header field values**
Cell discrimination based on pre-defined header field values	**Pre-assigned header field values**
Payload type discrimination	**PT field**
Loss priority indication and Selective cell discarding	**CLP field, Network congestion state**
Traffic shaping	**Traffic descriptor**

Figure 3-6 Functions supported at the UNI (U-plane)

3.4.1 Multiplexing among different ATM connections

This function multiplexes ATM connections with different QoS requirements. ATM connections may have either a specified or the unspecified QoS class as defined in section 4 of Appendix A.

The QoS class is the same for all cells belonging to the same connection, and remains unchanged for the duration of the connection.

(R) Network equipment supporting the public UNI shall support at least the Specified QoS Class 1 as defined in section 4.1 of Appendix A.

(O) Network Equipment supporting the public UNI may support the unspecified QoS class defined in section 4.2 of Appendix A.

(O) Network equipment supporting the private UNI may support either one or more specified QoS classes and/or the unspecified QoS class.

3.4.2 Cell rate decoupling

The cell rate decoupling[1] function at the sending entity adds unassigned cells to the assigned cell stream (cells with valid payload) to be transmitted, transforming a non-continuous stream of assigned cells into a continuous stream of assigned and unassigned cells. At the receiving entity the opposite operation is performed for both unassigned and invalid cells.The rate at which the unassigned cell are inserted/extracted depends on the bit rate (rate variation) of assigned cell and/or the physical layer transmission rate. The unassigned and invalid cells are recognized by specific header patterns which are shown in Figure 3-7.

Physical layers that have synchronous cell time slots generally require cell rate decoupling (e.g. SONET, DS3 and 8B/10B block-coded interfaces) whereas physical layers that have asynchronous cell times slots do not require this function (e.g. 4B/5B block-coded interface) since no continuous flow of cells needs to be provided. Therefore the requirements in this section only apply to physical layers that require continuous cell streams at the Physical-ATM layers boundary.

(R) Equipment supporting the UNI shall generate unassigned cells in the flow of cells passed to the physical layer in order to adjust to the cell rate required by the payload capacity of the physical layer.

(R) Equipment supporting the UNI shall encode the header fields of unassigned cells in accordance with the pre-assigned header field values defined in T1LB310 [7] and ITU-T Recommendation I.361.

---

[1] The term “cell rate decoupling” has a different meaning within ITU-T and refers to a physical layer process involving physical layer cells (e.g. idle cells) [12]. Further, the term “invalid” when applied to a cell or pattern refers to the illegal appearance of a physical layer cell at the ATM layer. The E3 and E4 interfaces may apply the ITU-T definition of the cell rate decoupling.

(R) The receiving ATM entity shall extract and discard the unassigned and invalid cells from the flow of cells coming from the physical layer.

Note: The cell rate governing the flow between physical and ATM layer will be extracted from Physical layer (e.g. SONET) timing information if required.

3.4.3 Cells discrimination based on pre-defined header field values

The pre-defined header field values defined at the UNI are given in Figure 3-7 (Ref. T1 LB310).

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

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

Figure 3-7 Pre-Defined Header Field Values

Meta-signalling cells are used by the meta-signalling protocol for establishing and releasing signalling virtual channel connections. For virtual channels allocated permanently (PVC), meta-signalling is not used.

(R) Equipment not supporting meta-signalling protocol at the UNI shall discard any cells received with VCI value = 1.

General broadcast signalling cells are used by the ATM network to broadcast signalling information independent of service profiles. For permanent virtual channel (PVC) service, the general broadcast signalling channel is not used since there is no control-plane process involved above the ATM layer.

(R) Equipment not supporting general broadcast signalling at the UNI shall discard any cells received with VCI value = 2.

The Virtual Path Connection (VPC) operation flow (F4 flow) is carried via specially designated OAM cells. F4 flow OAM cells have the same VPI value as the user-data cell transported by the VPC but are identified by two unique pre-assigned virtual channels within this VPC. At the UNI, the virtual channel identified by a VCI value = 3 is used for VP level management functions between ATM nodes on both sides of the UNI (i.e., single VP link segment) while the virtual channel identified by a VCI value = 4 can be used for VP level end-to-end (User <-> User) management functions.

The detailed layer management procedures making use of the F4 flow OAM cells at the UNI and the specific OAM cells format will be covered in section 3.5.

(R) Equipment supporting VP level management functions at the UNI shall encode the VCI field of the F4 flow OAM cells with the appropriate values as defined in T1 LB310 and shown in Figure 3-7.

(R) Equipment supporting VP level management functions at the UNI shall have the capability to identify F4 flow OAM cells within each VPC.

(R) Equipment not supporting VP level management functions at the UNI shall not transmit cells with VCI values 3 and 4 and shall discard any cells received with VCI value = 3 or 4.

A default header field value has been defined for the carriage of ILMI messages across the UNI. The specific encoding is shown in Figure 3-8.

Use	**Value [1]**			
	Octet 1	**Octet 2**	**Octet 3**	**Octet 4**
Carriage of ILMI message [2]	**00000000**	**00000000**	**00000001**	**0000aaa0**

1: "a" indicates that the bit is available for use by the appropriate ATM layer function

2: The transmitting ATM entity shall set the CLP bit to 0. The receiving ATM entity shall process ILMI cells with CLP=1 as ILMI cells and as any other CLP=1 cell.

Figure 3-8 Default Header Field Value

(R) Equipment supporting the UNI shall support VPI = 0 and VCI = 16 as the default values for the carriage of ILMI messages across the UNI.

3.4.4 Cell discrimination based on Payload Type (PT) Identifier field values

The main purpose of the PT Identifier is to discriminate between user cells (i.e., cell carrying user information) from non-user cells (see Figure 3-9). Code points 0 to 3 are used to indicate user cells. Within these code points, values 2 and 3 are used to indicate that congestion has been experienced in the network (see §3.6). Code points 4 and 5 are used for VCC level management functions. The PT value of 4 is used for identifying OAM cells communicated within the bounds of a VCC segment (i.e., single link segment across the UNI) while the PT value of 5 is used for identifying end-to-end OAM cells

The detailed layer management procedures making use of the F5 flow of OAM cells at the UNI and the specific OAM cells format is covered in section 3.5.

PTI Coding (MSB first)	Interpretation
000	User data cell, congestion not experienced, SDU-type = 0
001	User data cell, congestion not experienced, SDU-type = 1
010	User data cell, congestion experienced, SDU-type = 0
011	User data cell, congestion experienced, SDU-type = 1
100	Segment OAM F5 flow related cell
101	End-to-end OAM F5 flow related cell
110	Reserved for future traffic control and resource management
111	Reserved for future functions

Figure 3-9 Payload Type Indicator Encoding

(R) Equipment supporting VC level management functions at the UNI shall encode the PT field of the F5 flow OAM cells with the appropriate code points as defined in T1 LB310 [7].

(R) Equipment supporting VC level management functions at the UNI shall have the capability to identify F5 flow OAM cells within each VC.

(R) Equipment not supporting VC level management functions,via OAM cells, at the UNI shall ignore PT code points 100, 101.

(R) Where equipment at the UNI is not terminating a VC, it shall ignore and pass through unmodified all valid cells having PT code points which it does not support.

(R) Where equipment at the UNI is not terminating a VP, it shall ignore and pass through unmodified all valid cells having VCI code points which it does not support (e.g. VCI = 7).

3.4.5 Loss priority indication and selective cell discarding

The CLP field may be used for loss priority indication by the ATM end point and for selective cell discarding in network equipment. In a given ATM connection and for each user-data cell in the connection, the ATM equipment that first emits the cell can set the CLP bit equal to zero or one. The CLP bit is used to distinguish between cells of an ATM connection: A CLP bit equal to zero indicates a higher priority cell and a CLP bit equal to one indicates a lower priority cell. Upon entering the network, a cell with CLP value = 1 may be subject to discard depending on network traffic conditions.

The treatment of cells with CLP bit set (low priority cells) by the network traffic management functions is covered in section 3.6.

(O) User equipment supporting the UNI may use the CLP header field to indicate lower priority traffic (cells).

(O) ATM switches may tag CLP=0 cells detected by the UPC to be in violation of the Traffic Contract by changing the CLP bit from 0 to 1 (see 3.6).

3.4.6 Traffic Shaping

Traffic Shaping is expected to be an important function of ATM end-points in order to acheive the desired QoS. Traffic Shaping is covered in section 3.6.3.2.5

3.5 ATM Layer Management Specification (M-plane)

This section identifies the ATM Layer Management functions and procedures at the User Network Interface. Management functions at the UNI require some level of cooperation between customer premises equipment and network equipment. To minimize the coupling required between equipment on both sides of the UNI, the functional requirements have been reduced to a minimal set. The ATM Layer Management functions supported at the UNI are grouped into the following categories (see Figure 3-10):

Functions	Parameters
Fault Management	
• Alarm surveillance (VP)	**OAM cells**
• Connectivity Verification (VP,VC)	**OAM cells**
• Invalid VPI/VCI detection	**VPI/VCI**

Figure 3-10 ATM Layer Management Functions at the UNI

- Fault Management contains Alarm Surveillance and Connectivity Verification functions. OAM cells are used for exchanging related operation information. ATM cells with invalid VPI/VCI values are discarded and Layer Management is informed.

3.5.1 ATM Layer Management Information Flows

Figure 3-11 shows the OAM flows defined for the exchange of operations information between nodes (including customer premises equipment). At the ATM layer, the F4-F5 flows will be carried via OAM cells. The OAM cell flow used for end-to-end management functions may be carried transparently through the private ATM switch and made available to the user. The F4 flow is used for segment[2] or end-to-end (VP termination) management at the VP level using VCI values 3 and 4. The F5 flow is used for segment[3] or end-to-end (VC termination) management at the VC level using PT code points 4 and 5. A detailed explanation on OAM cell flow mechanism is given in T1S1.5/92-029R3, ITU-T Recommendation. I.610.

---

[2] In this case, the segment is defined as the link between the ATM nodes on either side of the UNI.

[3] In this case, the segment is defined as the link between the ATM nodes on either side of the UNI.

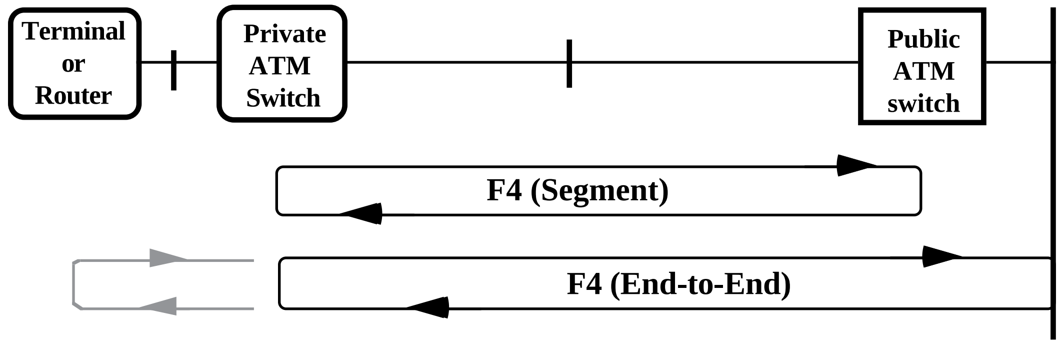

VP Level Service at the UNI

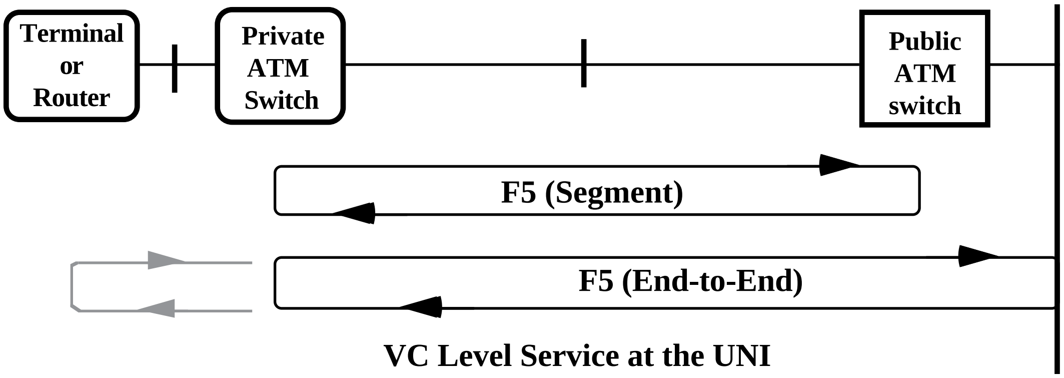

VC Level Service at the UNI

Figure 3-11 ATM Layer OAM flows at the UNI

In case of only Virtual Path (VP) visibility (i.e., VPC service at the Public Network Interface) the OAM operation information exchange will be limited to the F4 flow. Under this scenario, any VC level OAM functions and information exchange (F5 flow) are user-specific and ignored by the network. It is however possible to have VP level service at the Public UNI while maintaining VC visibility at the Private UNI. In this case, the private ATM switch would terminate the F4 flow but could carry transparently the user end-to-end F5 flow. For Virtual Channel (VC) visibility (VCC service), the OAM operation information exchange specified at the Public UNI could be limited to the F5 flow or could invoke both F4 and F5 flows.

(R) Equipment requiring/offering VP level service at the UNI shall support the F4 management flow as defined in T1S1.5/92-029R3 and ITU-T Recommendation I.610 for the functions defined in 3.5.3. Equipment should also be capable of transparently passing/carrying the F5 management flow.

(R) Equipment requiring/offering VC level service at the UNI shall support the F5 management flow as defined in T1S1.5/92-029R3 and ITU-T Rec. I.610 for the required functions defined in 3.5.3.

(R) Equipment requiring/offering VC level service only at a UNI shall either 1) process the F4 flow in accordance with T1S1.5/920029R3 and ITU-T Recommendation I.610, or 2) discard any F4 flow cells (i.e., cells with VCI = 3 and VCI = 4 are considered as cells with invalid VCI value).

The definition of “zero bandwidth” for a particular direction of any connection does not prohibit the transmission of OAM cells (ref. 3.6.3.2.3.7).

(R) For a point-to-multi-point connection, the only allowed use of F4 and F5 OAM flows is segment flows. The leaf node shall not send end-to-end OAM cells.

3.5.2 ATM OAM Cell Format

The virtual path connection (VPC) operational information is carried via the F4 flow OAM cells. These cells have the same VPI value as the user-data cells but are identified by pre-assigned VCI values. Two unique VCI values are used for every VPC as shown in Figure 3-12a. The VCI value = 3 is used to identify the connection between ATM layer management entities (LMEs) on both sides of the UNI (i.e., single link segment) and VCI value = 4 is used to identify connection between end-to-end ATM LMEs.

The virtual channel connection (VCC) operation information is carried via the F5 flow OAM cells. These cells have the same VPI/VCI values as the user-data cells but are identified by pre-assigned code points of the Payload Type (PT) field. Two unique PT values are used for every VCC as shown in Figure 3-12a. The PT value = “100” (4) is used to identify the connection between ATM layer management entities (LMEs) on both sides of the UNI (i.e., single link segment) while the PT value = “101” (5) is used to identify connection between end-to-end ATM LMEs.

End-to-end OAM cells must be passed unmodified by all intermediate nodes. The contents of these cells may be monitored by any node in the path. These cells are only to be removed by the endpoint of the VPC (F4 flow) or VCC (F5 flow). Segment OAM cells shall be removed at the end of a segment where, for the purposes of this specification, segment is defined as a single VP or VC link across the UNI.

The format of the Functions-Specific fields of the Fault Management OAM cell is defined in T1S1.5/92-029R3 and ITU-T Rec. I.610 and shown in Figure 3-12b.

(R) Equipment supporting the F4 and/or F5 flow at the UNI shall encode /interpret the OAM cells according to the format and encoding rules defined in T1S1.5/92-029R1 and ITU-T Rec. I.610 for the functions defined in 3.5.3.

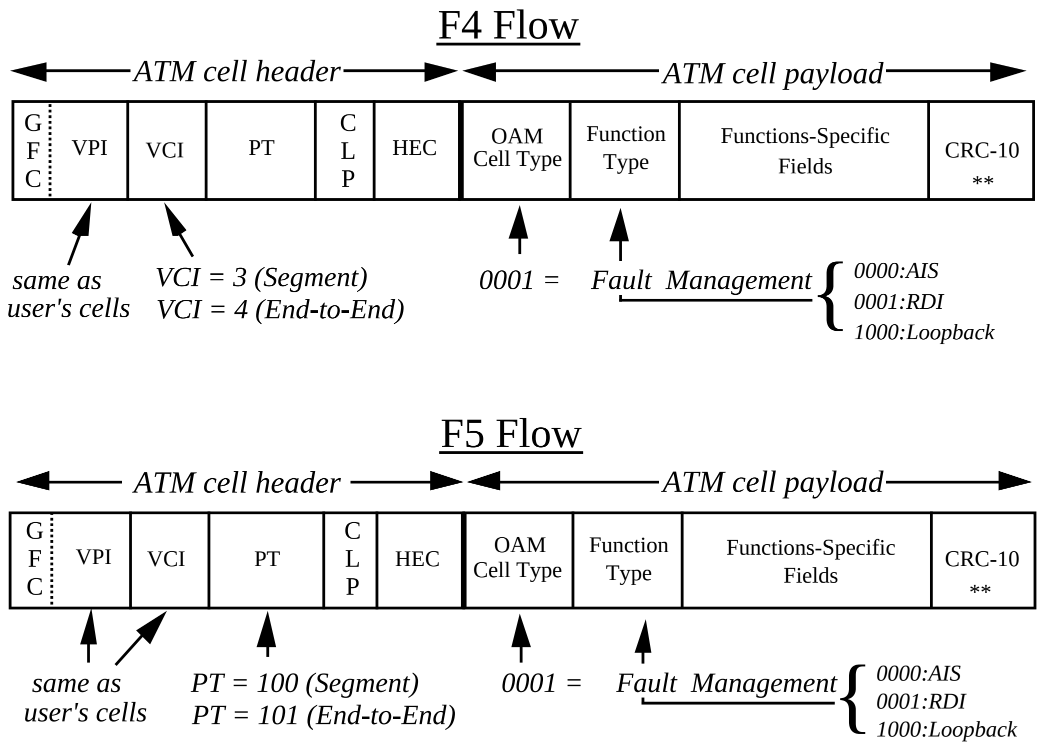

** : $G(x) = x^{10} + x^{9} + x^{5} + x^{4} + 1$

Figure 3-12a Common OAM Cell Format

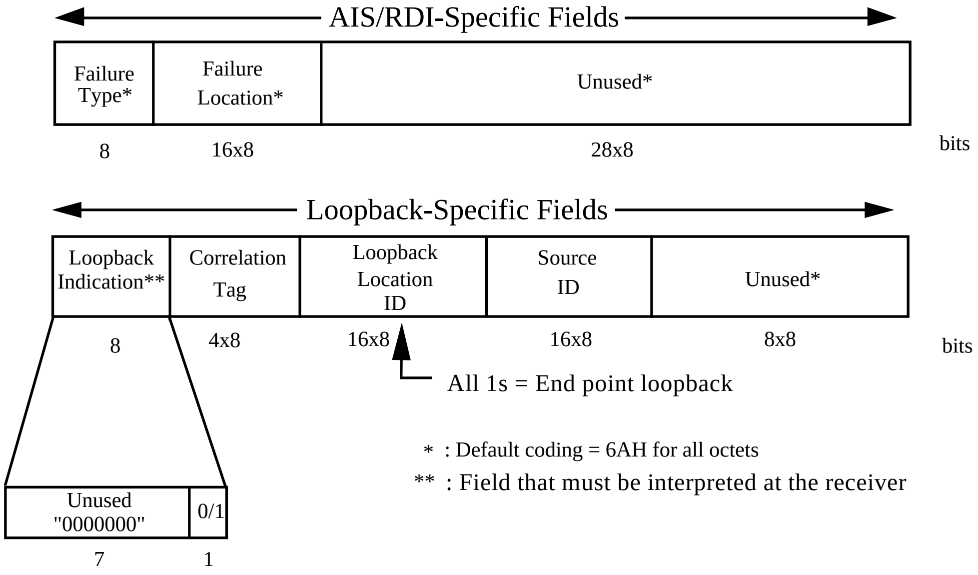

Figure 3-12b OAM Cell Fault Management-Specific Fields

Note: OAM Fault Management functions and cell formats may continue to evolve and software implementations may be advisable.

3.5.3 ATM Fault Management functions at the UNI

The Fault Management functions at the UNI are grouped in two categories: Alarm Surveillance and Connectivity Verification.

3.5.3.1 Alarm Surveillance

Alarm surveillance at the public UNI involves detection, generation and propagation of VPC/VCC failures (failure indications). In analogy with SONET physical layer, the failure indication signals are of two types: Alarm Indication Signal (AIS) and Remote Defect indication (RDI). These signals are carried via OAM cells as defined in section 3.5.2. The VP/VC Alarm Indication Signal (VP-AIS/VC-AIS) is generated by a VPC/VCC node at a connecting point to alert the downstream VPC/VCC nodes that a failure has been detected upstream. The VP-AIS/VC-AIS can be caused by the detection of a VPC/VCC failure or by the notification of a physical layer failure. Upon receiving a VP-AIS/VC-AIS, the VPC/VCC end-point at the public UNI will return a VP-RDI/VC-RDI to alert the upstream nodes that a failure has been detected downstream. VP-AIS and VP-RDI are always carried on VCI = 4. VC-AIS and VC-RDI are always carried over cells with PT = 101.

(R) ATM End-Point at the public UNI shall detect all incoming VP-AIS and generate a VP-RDI in the upstream direction (toward the public network) to alert the ATM nodes about the failure.

(R) Public network equipment supporting the UNI, acting as VP intermediate node, shall generate VP-AIS upon detection of a VPC failure or upon receiving a physical layer failure notification.

(O) End-Point of VCCs traversing the public network may detect an incoming VC-AIS and generate a VC-RDI in the upstream direction (toward the public network).

(O) Public network equipment supporting the UNI may generate VC-AIS upon detection of a VCC failure or upon receiving a physical layer failure notification.

Equipment inserting Alarm Surveillance cells will do so at a rate low enough to insure that Alarm Surveillance cells amount to less than one percent of the capacity of any link in the connection.

The duration of the condition and the rate associated with the generation and removal of alarm signals (VP-AIS, VP-RDI) is to be defined.

3.5.3.2 Connectivity Verification

Connectivity verification is supported by the use of the OAM loopback capability for both VP and VC connections. More complete details on this loopback function can be found in the modified text of I.610 [37]. The VCC or VPC being checked can remain in service while this loopback function is being performed. The OAM Cell Loopback function supported at the UNI uses the following three fields:

- Loopback indication - This eight-bit field identifier for the end point receiving the OAM cell, whether the incoming cell is to be looped back. A value of (00000001) indicates that the cell should be looped back. All other values indicate that the cell is to be discarded. Before the cell is looped back, the end point should decrement the value of the loopback indication field.

- Correlation Tag - At any given time multiple OAM Fault Management cells may be inserted in the same virtual connection. As a result, the OAM cell loopback mechanism requires a means of correlating transmitted OAM cells with received OAM cells. The node inserting the OAM cell may put any value in this 32-bit field and the endpoint looping back the cell shall not modify it.

- Loopback Location ID (optional) - This 96-bit field identifies the point(s) along a virtual connection where the loopback is to occur. The default value of all ones is used by the transmitter to indicate the end point. The receiver is not required to decode this field.

- Source ID (optional) - This 96-bit field can be used to identify the originator of the Loopback cell so the originator can identify the looped back cell when it returns. This may be encoded any way that enables the originating point to be certain that it has received the cell it transmitted. The default value is all ones.

(R) Endpoints receiving OAM cells with a loopback indication value other than (00000001) and Function Type = 1000 shall discard the cell.

(R) Endpoints receiving OAM cells with a loopback indication value of (00000001) and Function Type = 1000 shall decrement the loopback indication value and then loopback the cell within one second.

For connections that do not terminate in the public network, public network equipment will only insert end to end loopback cells when attempting to verify or isolate a fault. Equipment inserting loopback cells will do so at a rate low enough to insure that loopback cells amount to less than one percent of the capacity of any link in the connection. No requirement should be made that loopback cells support delay measurement.

3.5.3.2.1 End to End Loopback

End to end loopback cells are only looped back by the end point of a VPC or VCC. These cells may be inserted by any node in the connection (intermediate or end point) and may be monitored by any node. However, only end points may remove these cells. End to end loopback cells are indicated by a Payload Type value of (101) for VCCs and a VCI value of (4) for VPCs. An example of how the end to end loopback cell is used is shown in Figure 3-13.

3.5.3.2.2 UNI Loopback

UNI loopback is performed using segment loopback cells which are looped back by the end point of a VPC or VCC segment. The segment is defined as the link between the ATM nodes on either side of the UNI. Segment end points must either remove these cells or loop them back depending on the value in the loopback indication field. That is, these cells must not travel beyond the segment in which they are generated. UNI loopback cells are indicated by a Payload Type value of (100) for VCCs and a VCI value of (3) for VPCs. An example of how the segment loopback cell is used is shown in Figure 3-13.

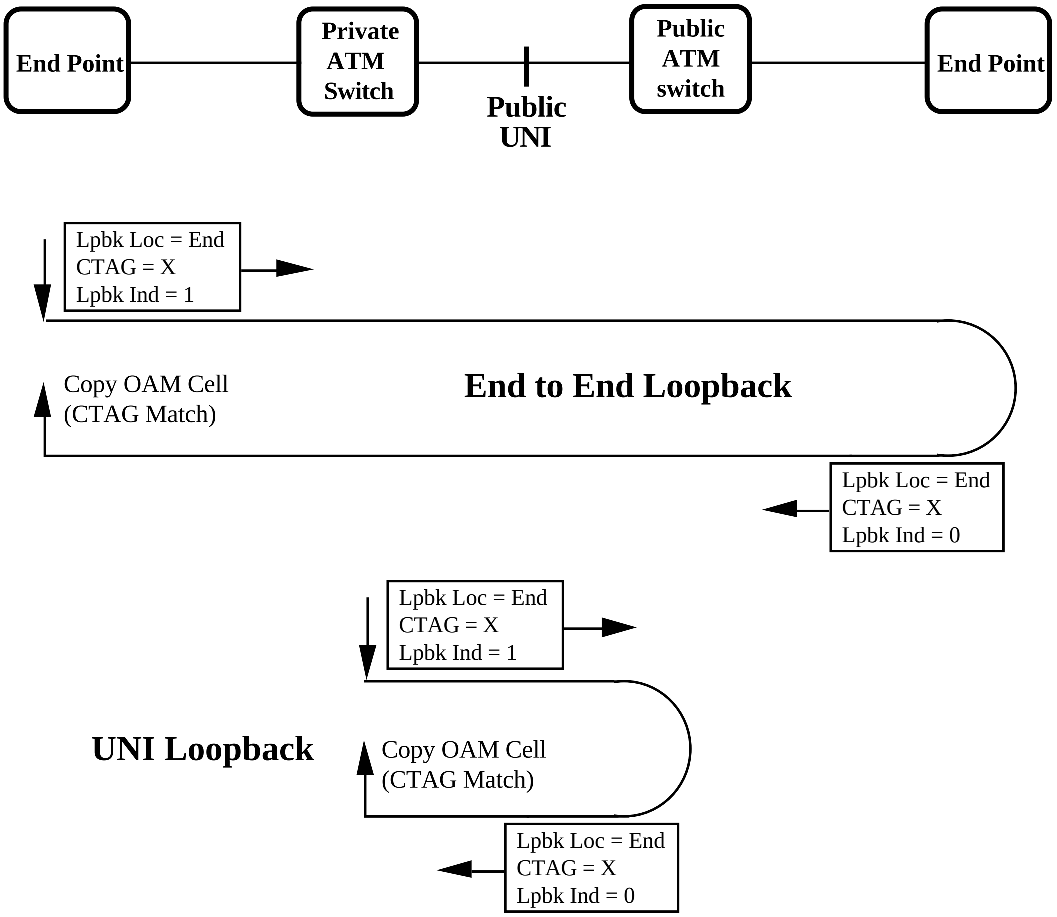

Figure 3-13 Loopback Function

3.6 Traffic Control and Congestion Control

3.6.1 Introduction

The B-ISDN, which is based on the ATM technique, is designed to transport a wide variety of traffic classes satisfying a range of transfer capacity needs and Network Performance objectives.

This section describes aspects of the Traffic Control and Congestion Control procedures relevant to the UNI specification.

In B-ISDN, ATM Layer congestion is defined as a state of Network Elements (e.g. switches, concentrators, cross-connects, and transmission links) in which the network is not able to meet the negotiated Network Performance objectives for the already established connections.

Congestion at the connection/call level is not considered in detail in this version of the specification.

In general, ATM Layer congestion can be caused by:

- unpredictable statistical fluctuation of traffic flows,
- fault conditions within the network.

ATM Layer congestion is to be distinguished from the state where buffer overflow is causing cell losses, but still meets the negotiated Quality of Service.

ATM Layer Traffic Control refers to the set of actions taken by the network to avoid congested conditions.

ATM Layer Congestion Control refers to the set of actions taken by the network to minimize the intensity, spread and duration of congestion. These actions are triggered by congestion in one or more Network Elements.

3.6.1.1 Objectives

The primary role of Traffic Control and Congestion Control parameters and procedures is to protect the network and the user in order to achieve Network Performance objectives. An additional role is to optimize the use of network resources.

The uncertainties of broadband traffic patterns and the complexity of Traffic Control and Congestion Control suggest a step-wise approach for defining traffic parameters and network Traffic Control and Congestion Control mechanisms. This document defines a restricted initial set of Traffic Control and Congestion Control capabilities aiming at simple mechanisms and realistic network efficiency.

It may subsequently be appropriate to consider additional sets of such capabilities, for which additional traffic control mechanisms will be used to achieve increased network efficiency.

The objectives of ATM Layer Traffic Control and Congestion Control for B-ISDN are as follows:

- ATM Layer Traffic Control and Congestion Control should support a set of ATM Layer Quality of Service (QoS) classes sufficient for all foreseeable B-ISDN services; the specification of these QoS classes should be consistent with Appendix A “ATM Bearer Service Quality of Service Objectives”.

- ATM Layer Traffic Control and Congestion Control should not rely on AAL protocols which are B-ISDN service specific, nor on higher layer protocols which are application specific. Protocol layers above the ATM Layer may make use of information which may be provided by the ATM Layer to improve the utility those protocols can derive from the network.

- The design of an optimum set of ATM Layer Traffic Controls and Congestion Controls should minimize network and end-system complexity while maximizing network utilization.

3.6.1.2 Generic Functions

To meet these objectives, the following functions form a framework for managing and controlling traffic and congestion in ATM networks and may be used in appropriate combinations.

- Network Resource Management (NRM): Provisioning may be used to allocate network resources in order to separate traffic flows according to service characteristics.

- Connection Admission Control (CAC) is defined as the set of actions taken by the network during the call set-up phase (or during call re-negotiations phase) in order to determine whether a virtual channel/virtual path connection request can be accepted or should be rejected (or whether a request for re-allocation can be accommodated). Routing is part of CAC actions.

- Feedback controls are defined as the set of actions taken by the network and by the users to regulate the traffic submitted on ATM connections according to the state of Network Elements.

- Usage Parameter Control (UPC) is defined as the set of actions taken by the network to monitor and control traffic, in terms of traffic offered and validity of the ATM connection, at the user access. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior, which can affect the QoS of other already established connections, by detecting violations of negotiated parameters and taking appropriate actions.

- Priority Control: the user may generate different priority traffic flows by using the Cell Loss Priority bit. A Network Element may selectively discard cells with low priority if necessary to protect as far as possible the Network Performance for cells with high priority.

- Traffic Shaping: traffic shaping mechanisms may be used to achieve a desired modification of the traffic characteristics.

- Other control functions are for further study.

All of these functions can make use of information that passes across the UNI. As a general requirement, it is desirable that a high level of consistency be achieved between the above traffic control capabilities.

3.6.1.3 QoS, Network Performance and Cell Loss Priority

The ATM Layer Quality of Service is defined by a set of parameters such as delay and delay variation, cell loss ratio, etc. Other QoS parameters are for further study.

For a complete list of QoS parameters defined in this document, refer to Appendix A.

A user requests one ATM Layer QoS class for each direction of an ATM layer connection. For each direction of an ATM Layer connection, a user requests a specific ATM Layer QoS from the QoS classes that a network provides. These requested QoS classes are a part of the Traffic Contract. The network commits to meet the requested Quality of Service as long as the user complies with the Traffic Contract (see §3.6.2.2).

A requested QoS class may be the “Unspecified QoS class” or may be one of the “Specified QoS classes,” see Appendix A, Section 4. A specified QoS class may contain at most two cell loss ratio objectives. If a specified QoS class does contain two cell loss ratio objectives, then one objective is for the CLP=0 cells and the other objective is for the CLP=1 cells of the ATM connection.

Network Performance objectives at the ATM Layer Service Access Point are intended to capture the network’s ability to meet the requested ATM Layer Quality of Service. It is the role of the upper layers, including the ATM Adaptation Layer, to translate this ATM Layer QoS to any specific application requested QoS.

3.6.1.4 Relation with other Standard documents

Section 3.6 of this UNI Specification is largely based on the ITU-T Recommendation I.371(Formerly ITU-T Recommendation I.371)[41]. Progress has been made in the ATM Forum on some of the issues that were left for further study in I.371. In addition, some modifications have been made to the text taken from I.371. Main additions, omissions and modifications to I.371 are as follows:

- Recommendation I.371 states that “The use of UPC function is recommended.” However, in this specification, the UPC function is required at the Public UNI (§3.6.3.2.3.3).

- In this specification, in addition to the Peak Cell Rate definition of I.371, two new traffic parameters (namely Sustainable Cell Rate (SCR) and Burst Tolerance) have been defined which are to be used jointly. These are optional traffic parameters which are defined by a Generic Cell Rate Algorithm (GCRA) which is the generalized version of the Peak Cell Rate monitoring algorithm described in Annex 1/I.371.

- I.371 states that when a traffic contract uses the CLP capability, the traffic parameters should be specified in terms of CLP=0 and CLP=0+1 traffic flows of the ATM connection (see section 2.3.2/I.371). This is intended to apply not only for Peak Cell Rate currently defined in I.371, but for any future traffic parameter as well.

  The ATM Forum specification conforms to I.371 for the Peak Cell Rate. However, this specification allows that the SCR and Burst Tolerance Traffic Parameters of an ATM Layer connection can also be specified in terms of CLP=0 and CLP=1 traffic flows. This is mainly due to the need for additional flexibility for mapping the traffic parameter definitions of some existing services, such as the Frame Relay service (see Appendix B, Figure B-2).

- In section 1.5/I.371, it is stated that “A user may request at most two different QoS classes for a single ATM connection, which differ with respect to Cell Loss Ratio objectives. The Cell Loss Priority bit of the ATM cell header allows for two Cell Loss Ratio objectives for a given ATM connection.” Furthermore, in section 2.3.2/I.371, it is stated that “As indicated in section 1.5, the network provides an ATM Layer QoS for each of the components (CLP=0 and CLP=0+1) of an ATM connection. The traffic contract specifies the particular QoS choice (from those offered by the network operator) for each of the ATM connection components. There may be a limited offering of QoS specifications for the CLP=1 component.”

  The text in section 1.5/I.371 is not consistent with sections 3.4.2.1 and 3.4.2.2 of I.150, moreover the text in section 2.3.2/I.371 is open to different interpretations. In an effort to provide a consistent interpretation of the above basic concept, in this specification, the QoS classes have been divided into two categories: Specified QoS class, and Unspecified QoS Class (see Appendix A, Section 4). For each direction of an ATM Layer connection, a user requests one QoS class at connection setup or subscription time. A specified QoS class may contain two cell loss ratio objectives, one for the CLP=0 cells and the other for the CLP=1 cells.

- Fast Resource Management functions are not included in this specification.

- This specification concentrates on Traffic Management and Quality of Service issues involved at the UNI, and therefore does not include these issues at the NNI.

3.6.2 User-Network Traffic Contract

3.6.2.1 Traffic Parameters and Descriptors

Traffic parameters describe traffic characteristics of an ATM connection. For a given ATM connection, Traffic Parameters are grouped into a Source Traffic Descriptor, which in turn is a component of a Connection Traffic Descriptor. These terms are defined below.

Definition: Traffic Parameters

A traffic parameter is a specification of a particular traffic aspect. It may be quantitative or qualitative. Traffic parameters may for example describe Peak Cell Rate, Sustainable Cell Rate, Burst Tolerance, and/or source type (e.g., telephone, video phone).

Definition: ATM Traffic Descriptor

The ATM Traffic Descriptor is the generic list of traffic parameters that can be used to capture the traffic characteristics of an ATM connection.

Definition: Source Traffic Descriptor

A Source Traffic Descriptor is a subset of traffic parameters belonging to the ATM Traffic Descriptor. It is used during the connection set-up to capture the intrinsic traffic characteristics of the connection requested by a particular source. The set of Traffic Parameters in a Source Traffic Descriptor can vary from connection to connection.

Definition: Connection Traffic Descriptor

The Connection Traffic Descriptor specifies the traffic characteristics of the ATM Connection at the Public or Private UNI. The Connection Traffic Descriptor is the set of traffic parameters in the Source Traffic Descriptor, the Cell Delay Variation (CDV) Tolerance and the Conformance Definition that is used to unambiguously specify the conforming cells of the ATM connection. Connection Admission Control procedures will use the Connection Traffic Descriptor to allocate resources and to derive parameter values for the operation of the UPC. The Connection Traffic Descriptor contains the necessary information for conformance testing of cells of the ATM connection at the UNI.

Any traffic parameter and the CDV Tolerance in a Connection Traffic Descriptor should fulfill the following requirements:

- be understandable by the user or terminal equipment; conformance testing should be possible as defined in §3.6.2.2;
- be useful in resource allocation schemes meeting Network Performance requirements described in §3.6.3.2.2;
- be enforceable by the UPC as specified in §3.6.3.2.3.

These criteria should be respected since users may have to provide these traffic parameters and CDV Tolerance at connection set-up. In addition, these traffic parameters and CDV Tolerance should be useful to the CAC procedure so that Network Performance objectives can be maintained once the connection has been accepted. Finally, they should be enforceable by the UPC in case of non-compliant usage in order to maintain Network Performance.

3.6.2.2 Traffic Contract Specification

A Traffic Contract specifies the negotiated characteristics of an ATM Layer Connection at a Private or Public UNI.

(R) The Traffic Contract at the Public UNI shall consist of a Connection Traffic Descriptor and a requested QoS class for each direction of the ATM Layer connection and shall include the definition of a compliant connection.

(O) The Private UNI may optionally support the same traffic contract as the Public UNI or a different traffic contract from the Public UNI.

The Connection Traffic Descriptor consists of all parameters and the Conformance Definition used to specify unambiguously the conforming cells of the ATM connection, i.e.:

- the Source Traffic Descriptor (e.g. Peak Cell Rate, Sustainable Cell Rate and Burst Tolerance),

- the CDV Tolerance,

- the Conformance Definition based on one or more applications of the Generic Cell Rate Algorithm (GCRA). See §3.6.2.4.1 for details on the GCRA.

The Conformance Definition based on the GCRA is used to specify unambiguously the conforming cells of an ATM Connection at the UNI. See §3.6.2.4 for further details on the GCRA and conformance definition. Examples on conformance definitions using the GCRA are provided in Appendix B.

The UNI Specification places no restrictions on the possible combinations that a user may request for (1) QoS class and (2) parameters in the Connection Traffic Descriptor.

The Conformance Definition should not be interpreted as the UPC algorithm. Although traffic conformance at the UNI is defined by the Conformance Definition based on the GCRA, the network provider may use any UPC as long as the operation of the UPC does not violate the QoS objectives of compliant connections.

The values of the traffic contract parameters can be specified either explicitly or implicitly as summarized in Figure 3-14. A parameter value is explicitly specified when it is specified by the user via signalling for SVCs or when it is specified via Network Management System (NMS) for PVCs. A parameter value specified at subscription time is also considered explicitly specified. A parameter value is implicitly specified when its value is assigned by the network operator using default rules, which, in turn, can depend on the information explicitly specified by the user. A default rule is the rule used by a network to assign a value to a traffic contract parameter that is not explicitly specified. In this version no default rules are specified, hence default rules are network-specific.

	explicitly specified parameters		*implicitly specified parameters*
	parameter values set at circuit-setup time	**parameter values specified at subscription time**	**parameter values set using default rules**
	requested by user/NMS	**assigned by network operator**	
SVC	signalling	by subscription	network-operator default rules
PVC	NMS	by subscription	network-operator default rules

Figure 3-14 Procedures Used to Set Values of Traffic Contract Parameters

(R) For switched or permanent ATM connections, traffic contract parameters shall be either explicitly specified or implicitly specified in accordance with the requirements listed in Section 3.6.2.4.

The CAC and UPC procedures are operator specific and should take into account the knowledge of the specified traffic contract to operate efficiently.

(R) In order to accommodate additional, experimental traffic parameters at either the Private or Public UNI, signalling messages shall have the capability to encode proprietary manufacturer or network operator information elements corresponding to the experimental traffic parameters.

(O) Experimental traffic parameters may be supported by the network equipment and the end user device across the UNI (either the Private UNI or the Public UNI) via mutual agreement.

3.6.2.3 Cell Conformance and Connection Compliance

Conformance applies to the cells as they pass the UNI and are in principle tested according to some combination of GCRA algorithms. The first cell of the connection initializes the algorithm and from then on each cell is either conforming or not conforming. As in all likelihood even with the best of intentions a cell or two may be non-conforming, it is inappropriate for the network operator to only commit to the QoS objectives for connections all of whose cells are conforming. Thus, the term “compliant”, which is not precisely defined, is used for connections in which some of the cells may be non-conforming.

(R) The precise definition of a compliant connection is left to the network operator. For any definition of a compliant connection, a connection for which all cells are conforming shall be identified as compliant.

Based on actions of the UPC function the network may decide whether a connection is compliant or not. The commitment by the network operator is to support the QoS for all connections that are compliant. The precise phrasing of the commitment is stated below.

(R) For compliant connections, at the Public UNI, the agreed QoS class shall be supported for at least the number of cells equal to the conforming cells according to the Conformance Definition.

For non-compliant connections, the network need not respect the agreed QoS class.

The Conformance Definition that defines conformity at the public UNI of the cells of the ATM connection uses a GCRA configuration in multiple instances to apply to particular combinations of the CLP=0 and CLP=0+1 cell streams with regard to the Peak Cell Rate and to particular combinations of CLP=0, CLP=1 and CLP=0+1 cell streams with regard to the Sustainable Cell Rate and Burst Tolerance. For example, the Conformance Definition may use the GCRA twice, once for Peak Cell Rate of the aggregate (CLP=0+1) cell stream and once for the Sustainable Cell Rate of the CLP=0 cell stream. Appendix B provides more details and further examples of the Conformance Definition that defines conformity of the cells of the ATM connection. The network operator may offer a limited set of alternative Conformance Definitions (all based on the GCRA) from which the user may choose for a given ATM connection. The minimum set of conformance definitions for early interoperability is defined in §3.6.2.5.

3.6.2.4 Traffic Contract Parameter Specification

(R) Peak Cell Rate for CLP=0+1 is a mandatory traffic parameter in any Source Traffic Descriptor.

(R) For switched ATM Layer connections, the Peak Cell Rate for CLP=0+1 and the QoS class must be explicitly specified for each direction in the connection-establishment SETUP message.

(R) The Cell Delay Variation Tolerance is a mandatory parameter in any Connection Traffic Descriptor.

(R) The Cell Delay Variation Tolerance shall be either explicitly specified at subscription time or implicitly specified.

Explicit specification of the CDV Tolerance within the signalling message is for further study.

(O) Sustainable Cell Rate and Burst Tolerance is an optional traffic parameter set in the Source Traffic Descriptor.

If either Sustainable Cell Rate or Burst Tolerance is specified then the other must be specified within the relevant Traffic Contract.

The Best-Effort-Capability is the label for a parameter in the ATM-User-Cell-Rate information element (see Section 5.4.5.6). The Best-Effort-Capability is used with the Unspecified QoS class (see Section A.4.2) with the only traffic parameter specified being the Peak Cell Rate specified for CLP=0+1.

If the Best-Effort-Capability is selected in the connection-establishment message then the connection admission-control procedures will not reject the call simply because the signalled Peak Cell Rate is greater than the rate of a link in the path (or is greater than the available bandwidth). The user need not conform to the signalled PCR and the network may enforce a PCR different than the signalled PCR. However, the network operator may reject the call request for other reasons, such as the number of such connections already established on a Network Element is at a chosen threshold, or the network operator simply does not support such connections that request the Best-Effort-Capability. The Best-Effort-Capability will be used to support those user terminals that are capable of and adapting to the time-variable available resources.

In future versions of this specification, traffic parameter negotiation, user notification of assigned traffic parameters, flow control mechanisms and further uses of the Best-Effort-Capability label may be defined.

Introduction of additional parameters to enhance the network resource management procedures or to capture traffic characteristics of a new type of connection is left open for further study.

3.6.2.4.1 Generic Cell Rate Algorithm (GCRA)

The Generic Cell Rate Algorithm (GCRA) is a Virtual Scheduling Algorithm or a continuous-state Leaky Bucket Algorithm as defined by the flowchart in Figure 3-15. The GCRA is used to define, in an operational manner, relationship between PCR and the “Cell Delay Variation tolerance” and relationship between SCR and the “Burst Tolerance”, see respectively §3.6.2.4.2.4 and §3.6.2.4.3.3. In addition, for the cell flow of an ATM connection, the GCRA is used to specify the conformance at the public or private UNI to declared values of the above two tolerances, as well as declared values of the Traffic Parameters “Peak Cell Rate” and “Sustainable Cell Rate and Burst Tolerance”, see respectively §3.6.2.4.2.2 and §3.6.2.4.3.2.

For each cell arrival, the GCRA determines whether the cell is conforming with the Traffic Contract of the connection, and thus the GCRA is used to provide the formal definition of traffic conformance to the Traffic Contract. Although traffic conformance is defined in terms of the GCRA, the network provider is not obligated to use this algorithm (or this algorithm with the same parameter values) for the Usage Parameter Control (UPC). Rather, the network provider may use any UPC as long as the operation of the UPC does not violate the QoS objectives of a compliant connection.

The GCRA depends only on two parameters: the increment I and the Limit L. These parameters have been denoted by T and τ respectively in Annex 1 of I.371, but have been given more generic labels herein since the GCRA will be used in multiple instances. The notation “GCRA(I, L)” means the Generic Cell Rate Algorithm with the value of the increment parameter set equal to I and the value of the limit parameter set equal to L.

The GCRA is formally defined in Figure 3-15. Figure 3-15 is a generic version of Figure 1 in Annex 1 of I.371. The two algorithms in Figure 3-15 are equivalent in the sense that for any sequence of cell arrival times, $\{t_a(k), k >= 1\}$ the two algorithms determine the same cells to be conforming and thus the same cells to be non-conforming. The two algorithms are easily compared if one notices that at each arrival epoch, $t_a(k)$, and after the algorithms have been executed, TAT = X + LCT, see Figure 3-15. An explanation of each algorithm follows.

The virtual scheduling algorithm updates a Theoretical Arrival Time (TAT), which is the “nominal” arrival time of the cell assuming equally spaced cells when the source is active. If the actual arrival time of a cell is not “too” early relative to the TAT, in particular if the actual arrival time is after TAT - L, then the cell is conforming, otherwise the cell is non-conforming.

Tracing the steps of the virtual scheduling algorithm in Figure 3-15, at the arrival time of the first cell $t_a(1)$, the theoretical arrival time TAT is initialized to the current time, $t_a(1)$. For subsequent cells, if the arrival time of the k^{th} cell, $t_a(k)$, is actually after the current value of the TAT then the cell is conforming and TAT is updated to the current time $t_a(k)$, plus the increment I. If the arrival time of the k^{th} cell is greater than or equal to TAT - L but less than TAT (i.e., as expressed in Figure 3-15, if TAT is less than or equal to $t_a(k)$ + L), then again the cell is conforming, and the TAT is increased by the increment I. Lastly, if the arrival time of the k^{th} cell is less than TAT -L (i.e., if TAT is greater than $t_a(k)$ + L), then the cell is non-conforming and the TAT is unchanged.

The continuous-state leaky bucket algorithm can be viewed as a finite-capacity bucket whose real-valued content drains out at a continuous rate of 1 unit of content per time-unit and whose content is increased by the increment I for each conforming cell. Equivalently, it can be viewed as the work load in a finite-capacity queue or as a real-valued counter. If at a cell arrival the content of the bucket is less than or equal to the limit value, L, then the cell is conforming, otherwise the cell is non-conforming. The capacity of the bucket (the upper bound on the counter) is L + I.

Tracing the steps of the continuous-state leaky bucket algorithm in Figure 3-15, at the arrival time of the first cell $t_a(1)$, the content of bucket, X, is set to zero and the last conformance time (LCT) is set to $t_a(1)_.$ At the arrival time of the k^{th} cell, $t_a(k)$, first the content of the bucket is provisionally updated to the value X', which equals the content of the bucket, X, after the arrival of the last conforming cell minus the amount the bucket has drained since that arrival, where the content of the bucket is constrained to be non-negative. Second, if X' is less than or equal to the limit value L, then the cell is conforming, and the bucket content X is set to X' plus the increment I for the current cell, and the last conformance time LCT, is set to the current time $t_a(k)$. If, on the other hand, X' is greater than the limit value L, then the cell is non-conforming and the values of X and LCT are not changed.

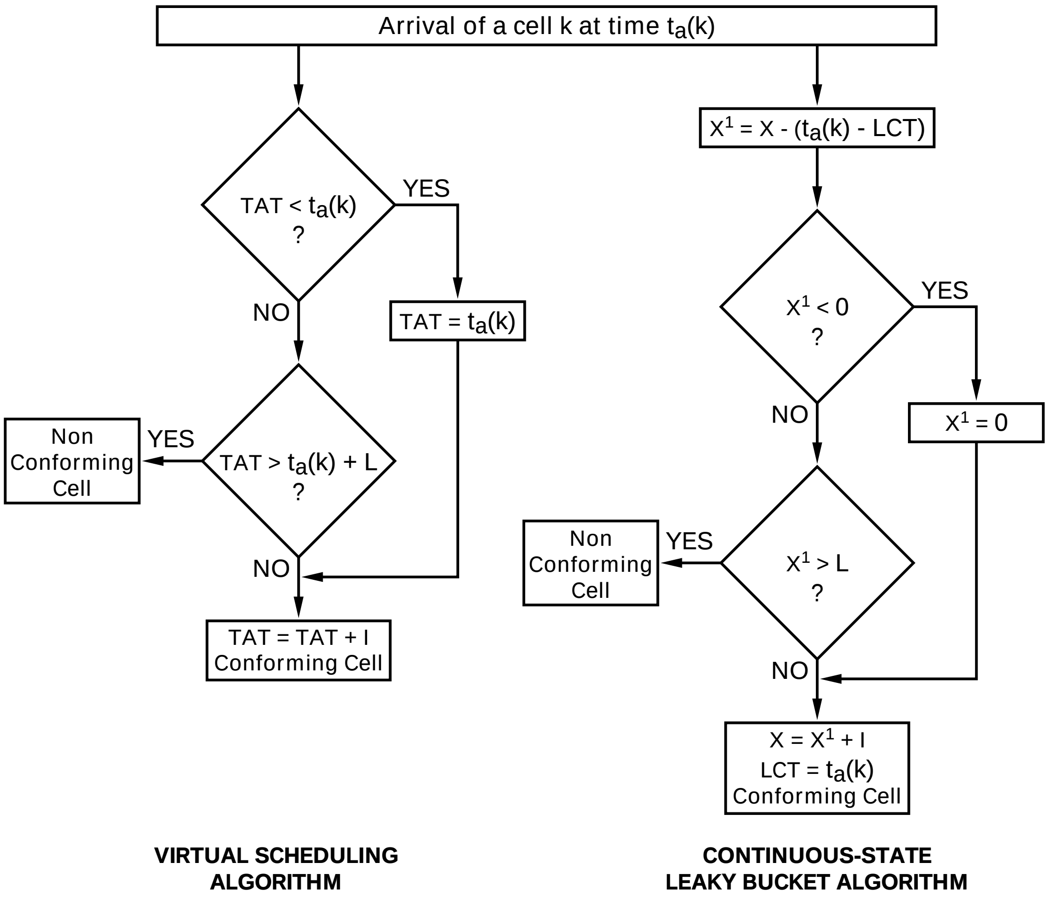

TAT	Theoretical Arrival Time
$t_a(k)$	Time of arrival of a cell

X	Value of the Leaky Bucket counter
X^1	Auxilary variable
LCT	Last Compliance Time

I	Increment
L	Limit

At the time of arrival ta of the first cell of the connection, TAT = $t_a(1)$

At the time of arrival ta of the first cell of the connection, X = 0 and LCT = $t_a(K)$

Figure 3-15 Equivalent versions of the Generic Cell Rate Algorithm

3.6.2.4.2 Peak Cell Rate

The Peak Cell Rate (PCR) traffic parameter specifies an upper bound on the traffic that can be submitted on an ATM connection. Enforcement of this bound by the UPC allows the network operator to allocate sufficient resources to ensure that the Network Performance objectives (e.g., for Cell Loss Ratio) can be achieved.

3.6.2.4.2.1 Peak Cell Rate Reference Model

The Equivalent Terminal configuration is given in Figure 3-16. Traffic sources, the multiplexer and the shaper define the Equivalent Terminal. This is only a model and does not preclude any particular implementation of the CPE or of the Terminal Equipment.

All traffic sources (AALs, FRM, etc.) offering cells to a connection are put together in the Equivalent Terminal. Each source generates Requests to send ATM cells at its own rate. All Requests are multiplexed in a Multiplexer (MUX in Figure 3-16) on a single link before entering the Virtual Shaper.

The Virtual Shaper is intended to reflect some smoothness in the cell flow offered to the ATM connection: at the PHY_SAP, the minimal inter arrival time between two consecutive Requests is greater than or equal to T which is called the Peak Emission Interval of the Connection. The output of the Virtual Shaper at the PHY_SAP of the Equivalent Terminal conforms to GCRA(T, 0). This conformity cannot be required at the Private or Public UNIs since CDV is allowed in the CPE as well as in the Terminal Equipment (TE). The output of the Virtual Shaper is affected by functions in the Equivalent Terminal that cause CDV characterized by τ^*. The value of τ^* is chosen such that the output cell flow conforms to GCRA(T, τ^*)

The output of the Equivalent Terminal is affected by functions in other CPE which may modify the CDV at the Public UNI characterized by τ. The value of τ is chosen such that the output cell flow is conforms to GCRA(T, τ)

The value of the Peak Emission Interval T is left to the discretion of the user to allow for intelligent multiplexing within the Customer and Terminal Equipment. For instance, AALs producing sporadic traffic may be synchronized to share the same transmission capacity. In other cases, T may be set to account for the combined activity of all traffic sources, e.g. the PCR of a VPC may be the sum of the PCRs of the VCCs contained in the VPC.

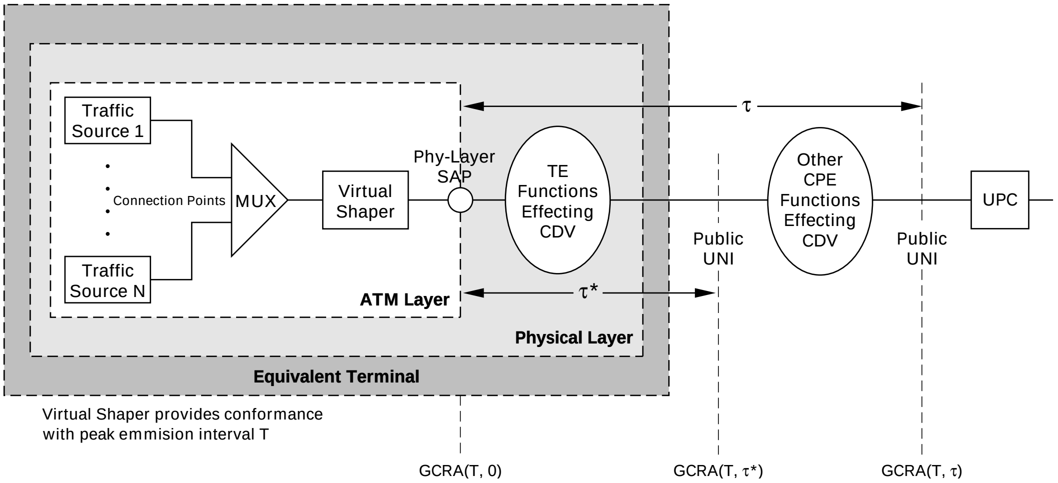

Figure 3-16 PCR Reference Model

The PCR traffic parameter is defined at the PHY_SAP within an Equivalent Terminal. The CDV Tolerance specified at the private UNI (τ*) that directly connects an end system to a private network accounts for the cell clumping introduced by the end system. The CDV Tolerance specified at the public UNI (τ) that connects an end system to a public network through a private ATM network via a private UNI accounts for the cell clumping introduced by the end system and the private ATM network.

3.6.2.4.2.2 Peak Cell Rate Parameter Definition

The following definition applies to ATM connections supporting both CBR and VBR services.

The PCR definition for a VPC/VCC is as follows:

Location: At the Physical Layer SAP in an Equivalent Terminal representing the VPC/VCC (this is only a reference configuration; see Figure 3-16).

Basic Event: Request to send an ATM_PDU in the Equivalent Terminal.

Definition: The PCR (R_p) of the ATM connection is the inverse of the minimum inter-arrival time T between two basic events above. T is called the Peak Emission Interval of the ATM connection.

In the signalling message, the PCR is coded as cells per second. The granularity supported by the signalling message is 1 cell/s. The defined coding for the Peak Cell Rate in the signalling message does not imply that any UPC mechanism has to support the same linear granularity for the PCR across the complete defined cell rate range.

3.6.2.4.2.3 Interpretation of the Definition of Peak Cell Rate and Equivalent Terminal

A “natural” or “intuitive” definition for Peak Cell Rate is the reciprocal of the minimum spacing of cells of an ATM connection on a transmission link. This intuitive definition is a rough approximation of the definition given in the previous section; however, this intuitive definition has technical flaws. These flaws are resolved by the above Equivalent-Terminal definition, at the possible expense of some obtuseness.

A simple, technical flaw of the “intuitive” definition is that for a slotted transmission medium, it constrains the possible values of Peak Cell Rate to be the reciprocal of an integral number of cell slot times. For example and using round numbers, if 150Mb/s is provided to the ATM layer, then the next possible peak rate below 150Mb/s is 75Mb/s and the next possible one below that is 50Mb/s, and so on. This granularity is too coarse.

Note that the equivalent-terminal is a conceptual model (or reference configuration). The definition does not imply that the real Customer Premises Equipment must do the shaping. The shaping by the Equivalent-Terminal may be viewed as a “thought experiment”. The fact that the source may not have shaped the traffic does not imply that the traffic is non-conforming since the criterion for conformance is at the UNI and is defined in terms of the GCRA.

Lastly, note that a definition of Peak Cell Rate does not tell the user the proper choice for the rate that meets the specific needs of the user. In particular, for VBR traffic sources, the equivalent-terminal model does not uniquely determine the value the user should pick for T. In the equivalent terminal, T must be chosen so that the queue in the buffer (“mux”) behind the shaper is stable. This allows T to be chosen so that its reciprocal is any value greater than the sustainable rate, up to the link rate.

3.6.2.4.2.4 Cell Delay Variation Tolerance

ATM Layer functions (e.g. cell multiplexing) may alter the traffic characteristics of ATM connections by introducing Cell Delay Variation as illustrated in Figure 3-16. When cells from two or more ATM connections are multiplexed, cells of a given ATM connection may be delayed while cells of another ATM connection are being inserted at the output of the multiplexer. Similarly, some cells may be delayed while physical layer overhead or OAM cells are inserted. Consequently with reference to the Peak Emission Interval T (i.e., the inverse of the contracted PCR R_p), some randomness may affect the inter-arrival time between consecutive VPC/VCC cells (i.e., the inverse of the contracted PCR) as monitored at the UNI (public or private). The upper bound on the “clumping” measure is called CDV Tolerance.

The CDV Tolerance allocated to a particular VPC/VCC at the private UNI(denoted by the symbol τ^*) represents at this interface a bound on the VPC/VCC cell clumping phenomenon due to the slotted nature of the ATM, the physical layer overhead and the ATM Layer Functions, i.e., cell multiplexing performed within the source terminal equipment. The CDV Tolerance allocated to a particular VPC/VCC at the public UNI (denoted by the symbol τ) represents at this interface a bound on the VPC/VCC cell clumping phenomenon due to the slotted nature of the ATM, the physical layer overhead, and the ATM layer functions performed within the Customer Premises Network before the public UNI.

The CDV Tolerance is defined in relation to the Peak Cell Rate according to the GCRA. In particular, the CDV Tolerance at the public UNI, τ, is defined in relation to the PCR according to the algorithm GCRA(T, τ), where T is the inverse of R_p (the PCR). Likewise, the CDV Tolerance at the private UNI, τ^*, is defined in relation to the PCR according to the algorithm GCRA(T, τ^*).

Figures 3-17 - 3-20 show a few examples that illustrate the potential cell clumping allowed at the Public UNI for a given value of τ and for a given value of T (the inverse of the contracted PCR) of an ATM connection according to the GCRA (T, τ). In all these examples, it is assumed that $T = 4.5\delta$, where δ is the time required to send 53 octets at the ATM layer data rate of 150 Mb/s (i.e., a peak bit rate of 33.3 Mb/s, including the cell header, is assumed).The notation in Figures 3-17 - 3-20 is defined in §3.6.2.4.1, where it is noted that “X + LCT” equals the “TAT” at each cell arrival time and after the GCRA has been executed.

From Figure 3-17, it can be observed that the minimum value of τ to be accommodated at the UNI is 0.5δ. From Figures 3-18 to 3-20, we observe that as τ increases, the minimum inter arrival time between conforming cells decreases. When τ is greater than or equal to $T - \delta$, the maximum number N of conforming back-to-back cells, i.e., at the full link rate, equals:

$$N = \left\lfloor 1 + \frac{\tau}{T - \delta} \right\rfloor \quad \text{For } T > \delta \text{ where } \lfloor x \rfloor \text{ stands for the integer part of x.}$$

This result of back-to-back cell clumping is illustrated in Figures 3-19 and 3-20.

The value of the CDV Tolerance may have an impact on the allocation of network resources for a particular VPC/VCC. It is therefore recommended that the CDV Tolerance (at both private and public UNI) be upper bounded, as a function of the PCR.

(R) A user shall explicitly or implicitly select a value for the CDV Tolerance at the public UNI for an ATM connection from a set of values supported by the network. Whether the set of values is to be standardized or to be determined by the network operator is for further study.

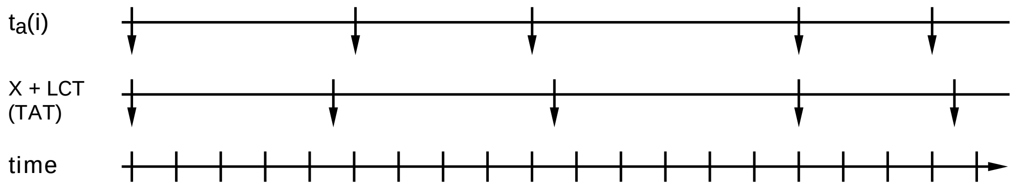

Figure 3-17 Ideal Cell Arrival at the Public UNI ($\tau = 0.5\delta$)

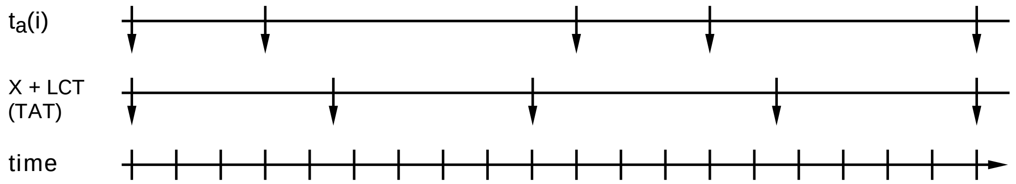

Figure 3-18 Possible Cell Arrival at the Public UNI ($\tau = 1.5\delta$)

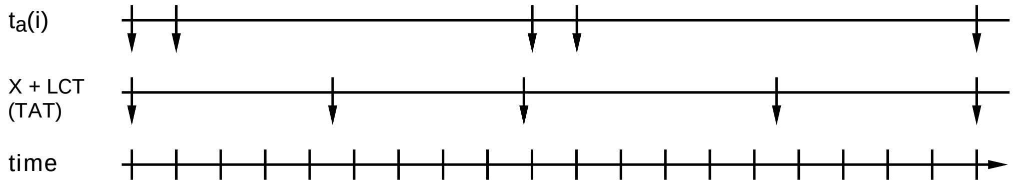

Figure 3-19 Possible Cell Arrival at the Public UNI ($\tau = 3.5\delta$)

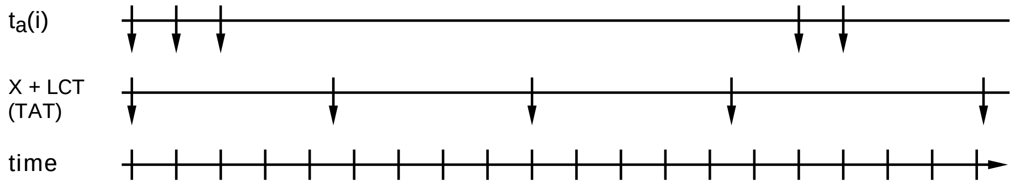

Figure 3-20 Possible Cell Arrival at the Public UNI ($\tau = 7\delta$)

3.6.2.4.3 Sustainable Cell Rate and Burst Tolerance

The Sustainable Cell Rate is an upper bound on the conforming average rate of an ATM connection. Enforcement of this bound by the UPC could allow the network operator to allocate sufficient resources, but less than those based on the Peak Cell Rate, and still ensure that the performance objectives (e.g., for Cell Loss Ratio) can be achieved.

3.6.2.4.3.1 Sustainable Cell Rate and Burst Tolerance Reference Model

The Sustainable Cell Rate (SCR) Reference Model is defined with reference to Figure 3-21.

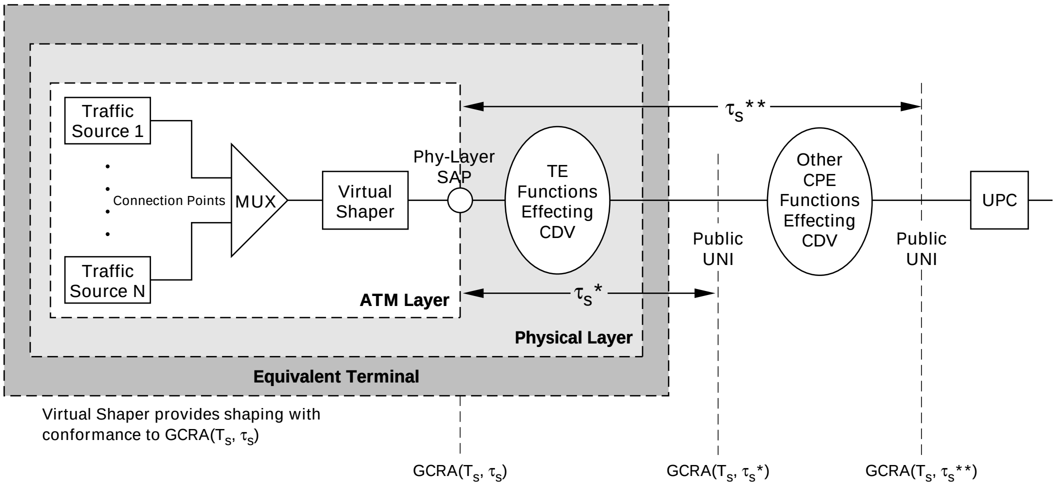

Figure 3-21 SCR and Burst Tolerance Reference Model

The SCR and Burst Tolerance traffic parameters are defined at the PHY_SAP within an Equivalent Terminal.

3.6.2.4.3.2 Sustainable Cell Rate and Burst Tolerance Parameter Definitions

The following definition applies to ATM connections supporting VBR services. The SCR and Burst Tolerance parameters for a VPC/VCC are defined according to the GCRA (§3.6.2.4.1) as follows:

Location:	At the Physical Layer SAP in an Equivalent Terminal representing the VPC/VCC (this is only a reference configuration; see Figure 3-21.)
Basic Event:	Request to send an ATM_PDU in the Equivalent Terminal.

Definition: The SCR, denoted as R_s, and the Burst Tolerance denoted as τ_s, of the ATM connection are defined by the GCRA(T_s, τ_s) based on the arrivals of the Basic Event above. R_s is the inverse of T_s, the increment parameter of the GCRA and τ_s is the limit parameter of the GCRA.

In the signalling message, the SCR is coded as cells per second. The granularity supported by the signalling message is 1 cell/s. The defined coding for the Sustainable Cell Rate in the signalling message does not imply that any UPC mechanism has to support the same linear granularity for the SCR across the complete defined cell rate range.

3.6.2.4.3.3 Interpretation of the Definition of Sustainable Cell Rate and Burst Tolerance in Conjunction with Peak Cell Rate

The SCR is an upper bound on the possible conforming “average rate” of an ATM connection, where “average rate” is the number of cells transmitted divided by the “duration of the connection”; where in this case, the “duration of the connection” is the time from the emission of the first cell until the state of the GCRA for the SCR returns to zero after the emission of the last cell of the connection. Relative to the Peak Cell Rate parameter, T_s is greater than T.

The SCR and Burst Tolerance traffic parameters are optional traffic parameters a user may choose to declare jointly, if the user can upper bound the realized average cell rate of the ATM connection to a value below the PCR. Note that as specified in §3.6.2.4 the PCR must be specified for every connection. To be useful to the network provider, the value of the SCR must be less than the PCR. For CBR connections, the user would not declare a SCR and would only declare a PCR.

The SCR and Burst Tolerance traffic parameters enable the end-user/terminal to describe the future cell flow of an ATM connection in greater detail than just the PCR. If an end-user/terminal is able to specify the future cell flow in greater detail than just the PCR, then the network provider may be able to more efficiently utilize the network resources. This directly benefits the network provider whether public or private, and in the case of public ATM networks, benefits the end-user with possible reduced charges for the connection.

If the source wants to submit traffic that conforms to the SCR ($R_s = 1/T_s$) and the Burst Tolerance (τ_s) and the Peak Cell Rate (1/T) at the PHY-SAP of the equivalent terminal, then it offers traffic that is conforming to the GCRA(T_s, τ_s) and the peak emission interval T (i.e., GCRA(T, 0)).

The Burst Tolerance together with the SCR and the GCRA determine the maximum burst size (MBS) that may be transmitted at the peak rate and still be in conformance with the GCRA(T_s, τ_s). The maximum burst size in number of cells is given by

$$\text{MBS} = \left\lfloor 1 + \frac{\tau_S}{T_S - T} \right\rfloor$$

where $\lfloor x \rfloor$ stands for the integer part of x.

In the signalling message, the Burst Tolerance is conveyed through the MBS which is coded in number of cells. The granularity supported by the signalling message is 1 cell. The MBS is used to derive the value of τ_s. The MBS and τ_s apply at the PHY_SAP of the Equivalent Terminal. Note that in order to determine τ_s from the MBS, the Peak Cell Rate also needs to be specified. By convention, the peak rate used in the calculation of τ_s is the Peak Cell Rate of the CLP=0+1 cell flow. This convention holds whether τ_s is associated with the SCR for the CLP=0, or the CLP=1, or the CLP=0+1 cell flow of the connection. Also, given the MBS, T, and T_s, then τ_s is not uniquely determined, but can be any value in the half-closed interval:

$$[\ (\text{MBS} - 1)(T_s - T), \text{MBS}(T_s - T)\ ].$$

Hence, in order for all parties to derive a common value for τ_s, by convention, the minimum possible value is used. Thus, given the MBS, T, and T_s, then τ_s is set equal to:

$$\tau_s = (\text{MBS} - 1)(T_s - T)$$

Note that over any closed time interval of length t, the number of cells, N(t), that can be emitted with spacing no less than T and still be in conformance with GCRA(T_s, τ_s) is bounded by:

$$N(t) \leq \min\left(\left\lfloor 1 + \frac{t + \tau_S}{T_S} \right\rfloor, \left\lfloor 1 + \frac{t}{T} \right\rfloor \right)$$

Observe that if t is greater than or equal to the MBS × T, then the first term of the above equation applies; otherwise, the second term applies.

Note that the maximum conforming burst size, defined above, does not imply that bursts of this size with arbitrary spacing between the bursts would be conforming with the GCRA(T_s, τ_s). Rather, in order for a burst this large to be conforming, the cell stream needs to be idle long enough for the state of the GCRA associated with SCR to become zero (i.e., long enough for the continuous-state leaky bucket to become empty) prior to the burst.

If a user chooses to specify a value for the SCR and Burst Tolerance traffic parameters and wishes to emit conforming bursts at the peak rate, then the appropriate choice of T_s and τ_s depends on the minimum spacing between bursts as well as the burst size. For a cell flow of an ATM connection, if the minimum spacing between bursts at the equivalent terminal is T_I

and if the maximum burst size (with inter-cell spacing T) is B, then the cell flow is conforming with GCRA(T_s, τ_s), if T_s, τ_s are chosen at least large enough to satisfy the following equation:

$$B = 1 + \left\lfloor \frac{\min(T_I - T_S, \tau_S)}{T_S - T} \right\rfloor$$

where $\lfloor x \rfloor$ stands for the integer part of x.

The traffic pattern conforming with the GCRA(T_s, τ_s) is in general not unique. Two traffic patterns are equivalent in relationship with the GCRA(T_s, τ_s) if they both conform at the PHY-SAP with the GCRA(T_s, τ_s) within the equivalent terminal. Therefore, any cell stream that complies with the GCRA(T, 0) and GCRA(T_s, τ_s) at the PHY_SAP has a Peak Cell Rate of $R_p = 1/T$, a mean cell rate which is bounded by $R_s = 1/T_s$ and a burst length which is bounded by B. Note that the bounds R_s and B are achievable. For example, a periodic cell stream with period B * T_s which transmits B cells at the peak rate with inter burst spacing $T_I = B * (T_s - T) + T$ has Peak Cell Rate R_p, mean cell rate R_s and burst length B, and is compliant with both GCRAs.

3.6.2.4.3.4 Relationship of CDV Tolerance, SCR and Burst Tolerance

ATM Layer functions (e.g. cell multiplexing) may alter the characteristics of a connection's cell flow between the Equivalent Terminal and the public or private UNI. Thus, as with the Peak Cell Rate, some tolerance for Cell Delay Variation may need to be considered in order that cells conforming to the GCRA(T_s, τ_s) at the Equivalent Terminal are also conforming at the public UNI.

It can be shown that if a terminal emits cells such that the emission epochs are conforming with GCRA(T_s, τ_s) and if the cells pass through a customer premises ATM network that introduces a random delay, but which is within the interval [d_{min}, d_{max}], then the cell arrival process at the public UNI is conforming with GCRA(T_s, $\tau_s + d_{max} - d_{min}$). Thus if τ, the CDV Tolerance parameter for the Peak Cell Rate is chosen to be $d_{max} - d_{min}$ (or is chosen to be a small quantile, e.g. 10^{-9}, of the possible delay variation), then τ could be used for the CDV Tolerance for the SCR as well. Thus, for simplicity, in the present UNI Specification, the same value for CDV Tolerance is used for the Peak Cell Rate and for the Sustainable Cell Rate of an ATM connection. Note that although a user may choose to select the CDV Tolerance from the set of values supported by the network to be greater than or equal to $d_{max} - d_{min}$, there is no requirement to do so.

In analogy with the PCR, the criterion for conformance to the SCR and the Burst Tolerance is specified at the UNI (both public and private). At the public UNI, the criterion for conformance is specified in terms of the Generic Cell Rate Algorithm with the argument T_s and $\tau_s + \tau$, GCRA(T_s, $\tau_s + \tau$).

With regard to conformance to the SCR at the public UNI, note that conformance depends on τ_s and on τ only via their sum. Thus, the constraint of a common CDV Tolerance τ for both Peak Cell Rate and SCR is not unduly restrictive as a user still has freedom in the choice of τ_s (by the choice of MBS), and thus can choose τ_s so that $\tau_s + \tau$ has the desired value. However, note that a negative consequence of choosing τ_s dependent on τ is that it violates the modeling principle that traffic parameters in the Source Traffic Descriptor are chosen based solely on the characteristics of the source and do not consider the equipment and traffic between the source and the UNI. For example, Peak Cell Rate is a traffic parameter, but CDV Tolerance is not a traffic parameter. Thus, although Burst Tolerance is defined herein as a traffic parameter, if the user chooses its value based on factors besides the source traffic, then the modeling principle for source traffic descriptors is violated. This matter will be revisited in the next edition of the UNI Specification.

Also note that to apply the equations for “MBS”, “N(t)” and “B” in the previous section to burst sizes at the public UNI, as opposed to at the equivalent terminal, one simply needs to replace the “τ_s” with “$\tau_s + \tau$.”

The text in this section is also applicable to the private UNI, one simply needs to replace τ with τ^*.

3.6.2.5 Conformance Definitions Supported at the Public UNI

The conformity of cells of an ATM connection at the public UNI is defined according to the GCRA algorithm in relation to the corresponding parameters specified in the Connection Traffic Descriptor. This Conformance Definition is specified in the Traffic Contract. The set of Conformance Definitions that will be supported at the public UNI is the network provider’s choice.

The Conformance Definitions include only traffic parameters for user data traffic of an ATM connection. For specification of OAM traffic see §3.6.3.2.3.7.

(R) For switched connections the signalling message shall be capable of conveying information that identifies at least the following set of Conformance Definitions. For permanent connections the Conformance Definition shall be explicitly identified at subscription time.

3.6.2.5.1 Conformance Definition for PCR

The following is a Conformance Definition for a Source Traffic Descriptor that specifies PCR for the CLP=0 cell stream and PCR for the CLP=0+1 cell stream:

Conformance Definition:

1. One GCRA(T_{0+1}, τ) defining the CDV tolerance in relation to the PCR of the CLP=0+1 cell stream.

2. One GCRA(T_0, τ) defining the CDV tolerance in relation to the PCR of the CLP=0 cell stream.

A CLP=0 cell that is conforming to both GCRAs (1) and (2) above is said to be conforming to the Connection Traffic Descriptor. A CLP=1 cell that is conforming to GCRA (1) above is said to be conforming to the Connection Traffic Descriptor. If the user requests tagging and if tagging is supported by the network, a CLP=0 cell that is not conforming to GRCA (2) above but is conforming to GCRA (1) above is considered to have the CLP bit changed to 1 and said to be conforming to the Connection Traffic Descriptor.

For networks that handle cells of the connection independent of the value of the CLP bit, the above Conformance Definition reduces to GCRA (1). The tagging option is not applicable to this Conformance Definition.

If the same value is specified for both PCR of the CLP=0 cell stream and PCR of the CLP=0+1 cell stream, this Conformance Definition could be used by a connection that only wants to send CLP=0 cells at its PCR (e.g. the example 3 given in Appendix B, the Constant Bit Rate Services). If the PCR of CLP=0 cell stream is set to zero, this Conformance Definition could be used by a connection that only wants to send CLP=1 cells at its PCR. Proper specification of the values of the PCRs allow this Conformance Definition to be used to accommodate all traffic mixes of a connection that only uses PCR traffic parameter.

3.6.2.5.2 Conformance Definition for PCR CLP=0+1 and SCR CLP=0

The following is a Conformance Definition for a Source Traffic Descriptor that specifies PCR for the CLP=0+1 cell stream and SCR for the CLP=0 cell stream:

Conformance Definition:

1. One GCRA(T_{0+1}, τ) defining the CDV tolerance in relation to the PCR of the CLP=0+1 cell stream.

2. One GCRA(T_{s0}, $\tau_{s0}+\tau$) defining the sum of the CDV tolerance and the Burst Tolerance in relation to the SCR of the CLP=0 cell stream.

A CLP=0 cell that is conforming to both GCRAs (1) and (2) above is said to be conforming to the Connection Traffic Descriptor. A CLP=1 cell that is conforming to GCRA (1) above is said to be conforming to the Connection Traffic Descriptor. If the user requests tagging and if tagging is supported by the network, a CLP=0 cell that is not conforming to GRCA (2) above but is conforming to GCRA (1) above is considered to have the CLP bit changed to 1 and said to be conforming to the Connection Traffic Descriptor.

This conformance definition allows a connection to send CLP=1 cells at a PCR equal to the specified PCR of the CLP=0+1 cell stream.

3.6.2.5.3 Conformance Definition for PCR CLP=0+1 and SCR CLP=0+1

The following is a Conformance Definition for a Source Traffic Descriptor that specifies PCR for the CLP=0+1 cell stream and SCR for the CLP=0+1 cell stream:

Conformance Definition:

1. One GCRA(T_{0+1}, t) defining the CDV tolerance in relation to the PCR of the CLP=0+1 cell stream.

2. One GCRA(T_{s0+1}, t_{s0+1}+t) defining the sum of the CDV tolerance and the Burst Tolerance in relation to the SCR of the CLP=0+1 cell stream.

A cell that is conforming to both GCRAs (1) and(2) above is said to be conforming to the Connection Traffic Descriptor. The tagging option is not applicable to this Conformance Definition.

3.6.3 Functions and Procedures for Traffic Control and Congestion Control at the UNI

3.6.3.1 Introduction

Generic Traffic Control and Congestion Control functions are defined as the set of actions taken by the network in all the relevant Network Elements.

Under normal operation, i.e., when no network failures occur, functions referred to as traffic control functions in this Specification are intended to avoid network congestion.

However, congestion may occur, e.g. because of malfunctioning traffic control functions caused by unpredictable statistical fluctuations of traffic flows or by network failures. Therefore, additionally, functions referred to as congestion control functions in this Specification are intended to react to network congestion in order to minimize its intensity, spread and duration.

A range of traffic and congestion control functions will be used in the B-ISDN to maintain the QoS of ATM connections. The following functions are described in this Specification:

Traffic Control Functions:

i Network Resource Management (§3.6.3.2.1)

ii Connection Admission Control (§3.6.3.2.2)

iii Usage Parameter Control (§3.6.3.2.3)

iv Selective Cell Discarding (§3.6.3.2.4)

v Traffic Shaping (§3.6.3.2.5)

vi Explicit Forward Congestion Indication (§3.6.3.2.6)

Congestion Control Functions:

vii Selective Cell Discarding (§3.6.3.3.1)

viii Explicit Forward Congestion Indication (§3.6.3.2.6)

Additional Control Functions:

Possible useful techniques that require further study to determine details are:

ix Connection Admission Control that reacts to and takes account of the measured load on the network

x Variation of usage monitored parameters by the network. For example, reduction of the peak rate available to the user.

xi Other traffic control techniques (e.g. re-routing, connection release, OAM functions) are for further study.

xii Fast Resource Management.

The impact on standardization of the use of these additional techniques (e.g. the impact on the ATM layer management, user-network signalling and control plane) requires further study.

Different levels of Network Performance may be provided on ATM connections by proper routing, Traffic Shaping, Priority Control and Resource Allocation to meet the required ATM Layer QoS for these connections.

3.6.3.2 Traffic Control Functions

3.6.3.2.1 Network Resource Management

The section on Network Resource Management of I.371 is included in the UNI Specification because Virtual Paths may be used at the UNI, even though the underlying structures are created and managed in the Network.

The use of Virtual Paths is described below. Other networking techniques are for further study.

Use of Virtual Paths

Virtual Paths are an important component of Traffic Control and Resource Management in the B-ISDN. With relation to Traffic Control, VPCs can be used to:

- simplify CAC,
- implement a form of priority control by segregating traffic types requiring different QoS,

- efficiently distribute messages for the operation of traffic control schemes (for example to indicate congestion in the network by distributing a single message for all VCCs comprising a VPC),

- aggregate user-to-user services such that the UPC can be applied to the traffic aggregate.

VPCs also play a key role in Network Resource Management. By reserving capacity on VPCs, the processing required to establish individual VCCs is reduced. Individual VCCs can be established by making simple connection admission decisions at nodes where VPCs are terminated. Strategies for the reservation of capacity on VPCs will be determined by the trade-off between increased capacity costs and reduced control costs. These strategies are left to the decision of network operators.

The peer-to-peer Network Performance on a given VCC depends on the performances of the consecutive VPCs used by this VCC and on how it is handled in Virtual Channel Connection Related Functions (CRF(VC)s). See Figure 3-22. A Connection Related Function may be a switch, concentrator or other network equipment.

If handled similarly by CRF(VC)s, different VCCs routed through the same sequence of VPCs experience similar expected Network Performance - e.g. in terms of Cell Loss Ratio, Cell Transfer Delay and Cell Delay Variation - along this route.

When VCCs within a VPC require a range of QoS, the VPC performance objective should be set suitably for the most demanding VCC carried. The impact on resource allocation is for further study.

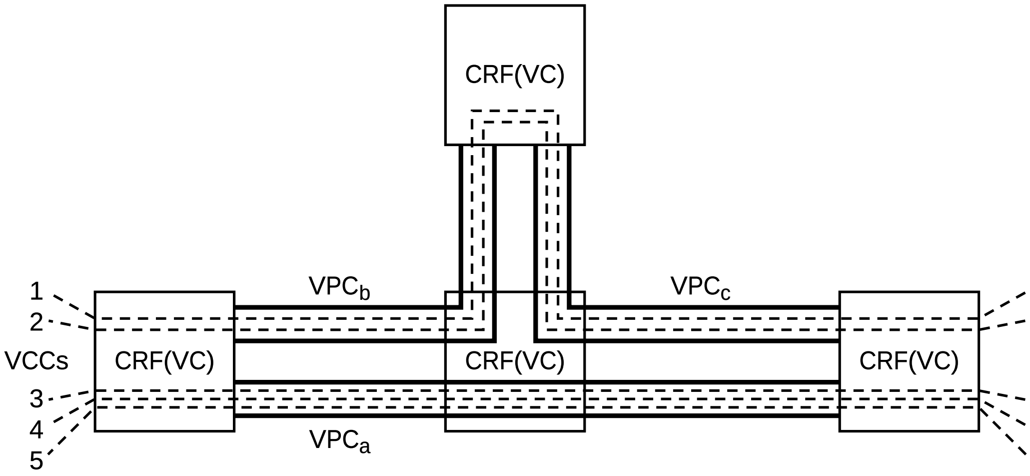

Figure 3-22 Mapping Cell Loss Ratios for VCC and VPC

Notes:

- VCCs 1 and 2 experience a Network Performance which depends on Network Performance on VPCs b and c and on how these VCCs are handled by Cell Relay Function CRF(VC)s. It may differ from Network Performance experienced by VCCs 3, 4 and 5, at least due to different Network Performances provided by VPCs.

- VCCs 3, 4 and 5 experience similar Network Performances in terms of Cell Delay and Cell Delay Variation if handled similarly by CRF(VC)s.

- On a user-to-user VPC, the QoS experienced by individual VCCs depends on CPE traffic handling capabilities.

On the basis of the applications of VPCs contained in ITU-T document I.311 [35] see §2.3.2, namely:

A) User-user application: the VPC extends between a pair of UNIs,

B) User-network application: the VPC extends between a UNI and a network node,

C) Network-Network application: the VPC extends between network nodes.

The above implies:

In case A: because the network has no knowledge of the QoS of the VCCs within the VPC, it is the user's responsibility to determine in accordance with the network capabilities the necessary QoS for the VPC.

In case B and C: the network is aware of the QoS of the VCCs carried within the VPC and has to accommodate them.

Statistical multiplexing of VC links within a VPC, where the aggregate peak of all VC links may exceed the Virtual Path Connection capacity, is only possible when all Virtual Channel links within the Virtual Path Connection can tolerate the QoS that results from this statistical multiplexing. The way this is managed is for further study.

As a consequence, when statistical multiplexing of Virtual Channel links is applied by the network operator, Virtual Path Connections may be used in order to separate traffic thereby preventing statistical multiplexing with other types of traffic. This requirement for separation implies that more than one Virtual Path connection may be necessary between network origination/destination pairs to carry a full range of QoS between them. Further implications of this are for further study.

3.6.3.2.2 Connection Admission Control

Connection Admission Control is defined as the set of actions taken by the network at the call set up phase (or during call re-negotiation phase) in order to establish whether a Virtual Channel Connection or a Virtual Path Connection can be accepted or should be rejected.

(R) The information contained in the Traffic Contract (§3.6.2.2) shall be accessible to the CAC function.

On the basis of Connection Admission Control in an ATM based network, a connection request is accepted only when sufficient resources are available to establish the connection through the whole network at its required Quality of Service (QoS) and to maintain the agreed QoS of existing connections. This applies as well to re-negotiation of connection parameters within a given call.

For each connection request the CAC function shall be able to derive the following information from the Traffic Contract (See §3.6.2.2):

- Values of parameters in the Source Traffic Descriptor (§3.6.2.1);
- the requested QoS class (§3.6.1.3);
- the value of the CDV Tolerance (§3.6.2.4.2.3);
- the requested Conformance Definition (§3.6.2.5).

Connection Admission Control makes use of the derived information and the network operator’s defintion of a compliant connection to determine:

- whether the connection can be accepted or not;
- traffic parameters needed by usage parameter control;
- routing and allocation of network resources.

Different strategies of network resource allocation may be applied for CLP=0 and CLP=1 traffic flows. If tagging is not requested, the network may still tag non-conforming cells, therefore CAC may accept a call even if tagging is not requested. In addition, information such as the measured network load may be used when performing CAC. This may allow a network operator to achieve higher network utilization while still meeting the Network Performance objectives.

Resource allocation schemes are currently left to network operator’s decision.

3.6.3.2.3 Usage Parameter Control

Recommendation I.371 states that “The use of UPC function is recommended.” However, in this specification, the UPC function is required at the Public UNI (§3.6.3.2.3.3).

(R) The UPC function shall be provided at the Public UNI.

3.6.3.2.3.1 UPC Functions

Usage Parameter Control is defined as the set of actions taken by the network to monitor and control traffic in terms of traffic offered and validity of the ATM connection, at the user access. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior which can affect the QoS of other already established connections by detecting violations of negotiated parameters and taking appropriate actions.

Connection monitoring encompasses all connections crossing the Public UNI. Usage Parameter Control applies to both user VCCs/VPCs and signalling virtual channels. Methods for monitoring meta-signalling channels and OAM flows are for further study.

The monitoring task for usage parameter control is performed for VCCs and VPCs respectively by the following two actions:

1. checking the validity of VPI and VCI (i.e., whether or not VPI/VCI values are associated with active VCCs) and monitoring the traffic entering the network from active VCCs in order to ensure that parameters agreed upon are not violated;

2. checking the validity of VPI (i.e., whether or not VPI values are associated with active VPCs), and monitoring the traffic entering the network from active VPCs in order to ensure that parameters agreed upon are not violated.

3.6.3.2.3.2 UPC Requirements

(R) The operation of the UPC mechanism, utilized by a network operator, shall not violate the QoS objectives of a compliant connections, see §3.6.2.2.

A number of desirable features of the UPC algorithm can be identified as follows:

- capability of detecting any non-compliant traffic situation;

- selectivity over the range of checked parameters (i.e., the algorithm could determine whether the user’s behavior is within an acceptance region);

- rapid response time to parameter violations;

- simplicity of implementation.

There are two sets of requirements relating to the UPC:

- those which relate to the quality of service impairments the UPC might directly cause to the user cell flow;
- those which relate to the resource the operator should allocate to a given connection and the way the network intends to protect those resources against misbehavior from the user side (due to fault conditions or maliciousness).

Two performance parameters are identified that could be considered when assessing the performance of UPC mechanisms. Methods for evaluating UPC performance and the need to standardize these methods are for further study.

- Response time: the time to detect a given non-compliant situation on a VPC/VCC under given reference conditions.
- Transparency: for given reference conditions, the accuracy with which the UPC initiates appropriate control actions on a non-compliant connection and avoids inappropriate control actions on a compliant connection.

Additional UPC performance parameters are for further study.

A specific UPC mechanism may commit errors by taking policing actions on a compliant connection, e.g. the number of discarded cells of a compliant connection is more than the number of non-conforming cells according to the Traffic Contract. The UPC can also fail to take the appropriate policing actions on a non-compliant connection.

Excessive policing actions of the UPC on a compliant connection are part of the overall Network Performance degradation. Safety margins may be provisioned depending upon the UPC algorithm to limit the degradation introduced by the UPC.

Policing actions performed on the excess traffic in case of Traffic Contract conformance violation are not to be included in the Network Performance degradation allocated to the UPC.

Impact of UPC on cell delay should also be considered. Cell Delay and Cell Delay Variation introduced by the UPC is also part of the delay and delay variation allocated to the network.

Performance of UPC

A method to determine whether a traffic flow is conforming to a negotiated PCR at a given interface is currently considered for Network Performance purposes. Non-conformance is measurable by a 1-point measurement process in terms of the ratio γ_M between the number of cells exceeding the traffic contract and the total number of submitted cells.

An ideal UPC implementing the 1 point-measurement process would just take policing actions on a number of cells according to this ratio. Although the process allows for a cell-based decision, it is not possible to predict which particular cells of a connection with non-conforming cells will suffer from the policing action (this is because of measurement phasing).

According to the definition of the conformance of a traffic flow to a PCR, the transparency of a UPC mechanism can be defined by the accuracy with which this mechanism approaches the ideal mechanism, i.e., the difference between the reference policing ratio γ_M and the actual policing ratio γ_p. A positive difference means that the UPC is taking less policing action than allowed. A negative difference means that policing action are unduly taken by the UPC.

The above discussion can also be applied to the SCR and the Burst Tolerance traffic parameters. The exact way of measuring the transparency of a given mechanism for the UPC and its dependence on time requires further study.

3.6.3.2.3.3 UPC location

Usage parameter control is performed on VCCs or VPCs at the point where the first VP or VC links are terminated within the network. Three possibilities are shown in Figure 3-23.

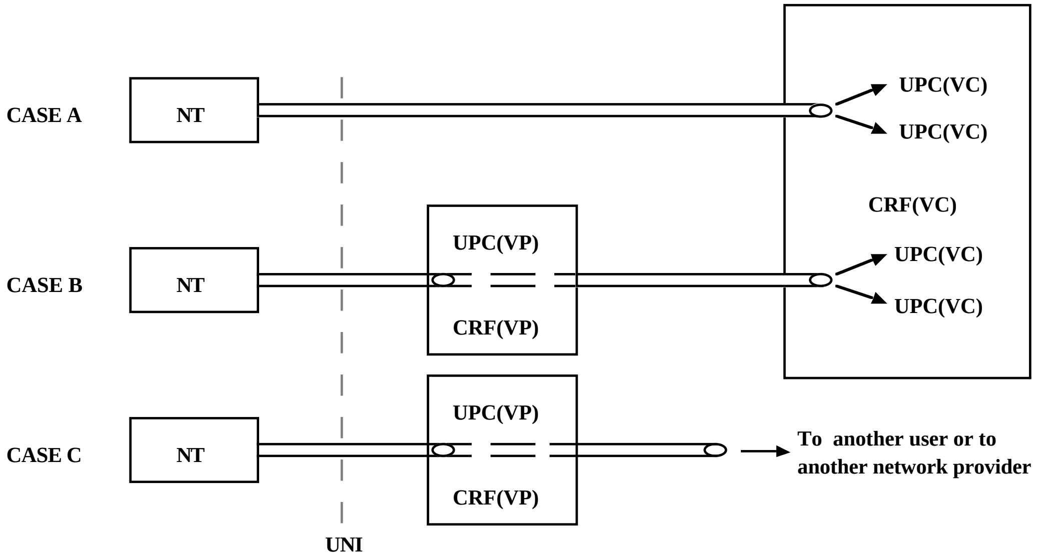

NT : Network Termination

CRF : Connection Related Function

CRF(VC) : Virtual Channel Connection Related Function

CRF(VP) : Virtual Path Connection Related Function

Figure 3-23 Location of the UPC Functions

Notes:

- In case A, the VPI value does not identify a negotiated VPC.
- Provision of UPC at other locations is for further study.

In the following, CRF(VC) stands for Virtual Channel Connection Related Function, and CRF(VP) stands for Virtual Path Connection Related Function. The following CRF(VC)s and CRF(VP)s refer to the first CRFs on the public network side of the Public UNI. A CRF(VC) or a CRF(VP) may respectively be a VC or VP concentrator.

(CR) If connected directly to CRF(VC) (CASE A of figure 3-23), the UPC function shall be performed within the CRF(VC) on VCCs before the switching function is executed (action 1, §3.6.3.2.3.1).

(CR) If connected directly to CRF(VC) via CRF(VP) (CASE B of figure 3-23), the UPC function shall be performed within the CRF(VP) on VPCs only (actions 2, §3.6.3.2.3.1) and within the CRF(VC) on VCCs only (action 1, §3.6.3.2.3.1).

(CR) If connected to user or to another network provider via CRF(VP) (CASE C of figure 3-23), the UPC function shall be performed within the CRF(VP) on VPCs only (action 2, §3.6.3.2.3.1).

In CASE C of figure 3-23, the VCC usage parameter control will be done by the first public network provider (if any) where CRF(VC) is present.

3.6.3.2.3.4 Traffic parameters subject to control at the UPC

Traffic Parameters that may be subject to control are those included in the Source Traffic Descriptor.

(R) The Peak Cell Rate of the CLP=0+1 cell flow shall be controlled for all types of connections at the Public UNI.

(O) Even when specified, control of Sustainable Cell Rate and Burst Tolerance is at the discretion of the network operator.

3.6.3.2.3.5 UPC actions

The UPC is intended to control the traffic offered by an ATM connection to ensure conformance with the negotiated Traffic Contract. The objective is that a user will never be able to exceed the Traffic Contract.

At the cell level, actions of the UPC function may be:

- cell passing;
- cell tagging (network operator's option); cell tagging operates on CLP=0 cells only, by overwriting the CLP bit to 1;
- cell discarding.

If the tagging option is used by the network operator, CLP=0 cells identified by the UPC function to be non-conforming to the CLP=0 cell stream are converted to CLP=1 cells. (For terminology, these cells with the converted CLP bit are called tagged cells). A tagged cell that is identified by the UPC function to be conforming is passed; otherwise, it is discarded. Likewise, a user submitted CLP=1 cell that is identified by the UPC function to be conforming is passed; otherwise, it is discarded.[4]

(R) Cell passing shall be performed on cells that are identified by the UPC as conforming. (For terminology, if the tagged cell is passed, then it is said to be conforming to the UPC function.)

(R) Cell discarding shall be performed on cells that are identified by the UPC as non-conforming.

(O) Following the UPC function, traffic shaping may be used to perform cell re-scheduling (e.g. to reduce cell clumping) on cells identified by the UPC as conforming.

Besides the above actions at the cell level, as an option, one other action performed at the connection level may be initiated by the UPC:

(O) At the option of the network provider, the UPC function may initiate the release of an identified non-compliant SVC connection.

3.6.3.2.3.6 Relationship between UPC, CLP and Network Performance

When an ATM connection utilizes the CLP capability on user request, network resources are allocated to CLP=0 and CLP=1 traffic flows as described in §3.6.3.2.2. By controlling the connection traffic flows, allocating adequate resources and suitably routing, a network operator may provide the requested QoS class for CLP=0 and CLP=1 cell flows.

When no additional network resource has been allocated for CLP=1 traffic flow (either on user request or due to network provisioning), CLP=0 cells identified by the UPC as non-conforming are discarded. In this case, tagging is not applicable.

---

[4] Note that the in principle cell flow that is passed by the UPC with the tagging option is no greater than what might be passed by a UPC without the tagging option when the end terminal sets the CLP bit to 1 on those cells that would otherwise have been tagged under the tagging option.

Section 3.6.3.2.3.2 addresses undue UPC actions on compliant ATM connections. This is part of the Network Performance degradation allocated to the UPC and should remain of a very low probability.

When cells of the aggregate CLP=0+1 flow are non-conforming to the parameters negotiated for the aggregate stream, the UPC function performed on the aggregate flow may discard CLP=0 cells that would not be considered in excess by the UPC function performed on the CLP=0 cell stream.

3.6.3.2.3.7 Relationship Between UPC and OAM

For OAM cell flows across the UNI, the network may require from the user the knowledge of some traffic parameters such as the Peak Cell Rate and knowledge of some clumping tolerance t_{OAM} for the OAM traffic of an ATM connection. Regardless of whether a user explicitly or implicitly specifies the OAM cell stream, the network shall police OAM cell flows together with user data cell streams. However, the initial release of the user-network signalling protocol will not allow a user to explicitly specify traffic parameters of OAM flow (i.e. only the combined user data cell and OAM cell parameters are allowed). Traffic parameters for OAM cell flow across the UNI may be explicitly specified at subscription time, or implicitly by a default rule (see §3.6.2.2).

(R) OAM features for a VCC or a VPC shall be selected by the end user at service subscription from amongst the features supported by the public network.

(R) All cells of a connection that are transported transparently by a network shall be policed together.

> Note: User cells are transported transparently end-to-end. OAM cells (i.e., PT=101 or VCI=4) are transported transparently. See 3.4.4 for other cells that may be transported transparently.

3.6.3.2.4 Selective Cell Discard

Network Elements may selectively discard cells of the CLP=1 flow while still meeting Network Performance objectives on both the CLP=0 and CLP=1 flows.

(R) For a given ATM connection the Cell Loss Ratio objective for CLP=1 cells shall be greater than or equal to the Cell Loss Ratio objective for CLP=0 cells.

3.6.3.2.5 Traffic Shaping (O)

When used in the source ATM end-point, traffic shaping is a mechanism that attains desired characteristics for the stream of cells emitted into a VCC or a VPC. When used in a private ATM Switch, traffic shaping is a mechanism that alters the traffic characteristics of a stream of cells on a VCC or a VPC to achieve a desired modification of those traffic characteristics. Traffic shaping must maintain cell sequence integrity on an ATM connection. Examples of traffic shaping are Peak Cell Rate reduction, burst length limiting and reduction of cell clumping due to CDV by suitably spacing cells in time.

(O) CPE may perform traffic shaping to be in conformance with the Connection Traffic Descriptor and associated parameter values that were negotiated with the public network.

Traffic shaping is an optional function. For example, an ATM end-point may choose to shape to the negotiated Peak Cell Rate for the aggregate cell stream of CLP=0 and CLP=1 cells and may choose not to shape to the negotiated Peak Cell Rate for the CLP=0 cell stream and instead to allow the network's UPC mechanism to tag as CLP=1 the non-conforming CLP=0 cells.

The algorithm used by the traffic shaping function is not specified. However, when used for conformance with negotiated parameters of a traffic descriptor, a natural choice for the algorithm is one that mimics the Conformance Definition (§3.6.2.2).

3.6.3.2.6 Explicit Forward Congestion Indication

The EFCI is a congestion notification mechanism that the ATM layer service user may make use of to improve the utility that can be derived from the ATM layer. Since the use of this mechanism by the CPE is optional, the network operator should not rely on this mechanism to control congestion.

A network element in an impending-congested state or a congested state may set an explicit forward congestion indication in the cell header so that this indication may be examined by the destination CPE. For example, the end user's CPE may use this indication to implement protocols that adaptively lower the cell rate of the connection during congestion or impending congestion. A network element that is not in a congested state or an impending congested state will not modify the value of this indication. An impending-congestion state is the state when a network equipment is operating around its engineered capacity level.

The mechanism by which a network element determines whether it is in an impending-congested or a congested state is an implementation issue and is not subject to standardization. The mechanism by which the congestion indication is used by the higher layer protocols in the CPE is for further study.

The impact of explicit forward congestion indication on the Traffic Control and Congestion Control functions requires further study.

3.6.3.3 Congestion Control Functions

For low priority traffic, some adaptive rate control facilities at the ATM layer or above may be used. Such cell-based reactive techniques are for further study.

The following congestion control functions have been identified. Other congestion control functions are for further study.

3.6.3.3.1 Selective Cell Discard

A congested Network Element may selectively discard cells explicitly identified as belonging to a non-compliant ATM connection and/or those cells with CLP=1. This is to primarily protect, as long as possible, CLP=0 flows.

3.6.3.3.2 Reaction to UPC failures

Due to equipment faults (e.g. in usage parameter control devices and/or other network elements) the controlled traffic characteristics at the UPC/NPC could be different from the values agreed to during the call set-up phase. To cope with these situations, specific procedures of the management plane should be designed (e.g. in order to isolate the faulty link). The impact of these malfunctioning situations on the usage parameter control needs further study.