International Telecommunication Union

ITU-T

TELECOMMUNICATION STANDARDIZATION SECTOR OF ITU

I.371

(03/2004)

SERIES I: INTEGRATED SERVICES DIGITAL NETWORK

Overall network aspects and functions – General network requirements and functions

---

Traffic control and congestion control in B-ISDN

ITU-T Recommendation I.371

ITU-T

ITU-T I-SERIES RECOMMENDATIONS

INTEGRATED SERVICES DIGITAL NETWORK

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation I.371

Traffic control and congestion control in B-ISDN

Summary

This Recommendation addresses functions and parameters for traffic and congestion control in the B-ISDN.

A user-network and an inter-network traffic contract are defined in terms of a traffic descriptor that includes traffic parameters and associated tolerances of an ATM layer transfer capability and of quality of service requirements associated to a QoS class. Relevant traffic parameters and generic conformance definitions for these parameters are specified. ATM transfer capabilities that make use of these traffic parameters to allow for different combinations of QoS objectives and multiplexing schemes and specific conformance definitions applying to these ATM transfer capabilities are provided.

In addition, traffic control and congestion control functions are further specified, among which are traffic parameter control functions at user-network and inter-network interfaces. Some specific traffic control interworking configurations are described.

Finally, procedures for traffic control, congestion control and resource management are defined. This includes specific formats and information supported by resource management cells.

Source

ITU-T Recommendation I.371 was approved on 29 March 2004 by ITU-T Study Group 13 (2001-2004) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

ITU-T Recommendation I.371

Traffic control and congestion control in B-ISDN

1 Scope

This Recommendation describes traffic control and congestion control procedures for the B-ISDN at the ATM layer.

ATM layer traffic control refers to all network actions aiming to meet the network performance objectives and negotiated QoS commitments, and to avoid congested conditions. ATM layer congestion control refers to all network actions to minimize the intensity, spread and duration of congestion.

This Recommendation provides a general description as well as objectives and procedures for traffic control and congestion control. Additionally, it describes the concepts of the traffic contract. It specifies the ATM transfer capabilities (ATCs) including, for each ATC, the applicable source traffic descriptor, associated tolerances and conformance definition.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[1] ITU-T Recommendation I.326 (2003), *Functional architecture of transport networks based on ATM*.

[2] ITU-T Recommendation I.113 (1997), *Vocabulary of terms for broadband aspects of ISDN*.

[3] ITU-T Recommendation I.150 (1999), *B-ISDN asynchronous transfer mode functional characteristics*.

[4] ITU-T Recommendation I.311 (1996), *B-ISDN general network aspects*.

[5] ITU-T Recommendation I.321 (1991), *B-ISDN protocol reference model and its application*.

[6] ITU-T Recommendation I.356 (2000), *B-ISDN ATM layer cell transfer performance*.

[7] ITU-T Recommendation I.357 (2000), *B-ISDN semi-permanent connection availability*.

[8] ITU-T Recommendation I.358 (2003), *Call processing performance for switched virtual channel connections (VCCs) in a B-ISDN*.

[9] ITU-T Recommendation I.361 (1999), *B-ISDN ATM layer specification*.

[10] ITU-T Recommendations I.363.1 (1996), I.363.3 (1996) and I.363.5 (1996), *B-ISDN ATM adaptation layer specification*.

[11] ITU-T Recommendation I.413 (1993), *B-ISDN user-network interface*.

[12] ITU-T Recommendations I.432.1 (1999), I.432.2 (1999), I.432.3 (1999) and I.432.4 (1999), *B-ISDN user-network interface – Physical layer specification*.

[13] ITU-T Recommendation I.610 (1999), *B-ISDN operation and maintenance principles and functions*.

[14] ITU-T Recommendation I.630 (1999), *ATM protection switching*.

[15] ITU-T Recommendation I.731 (2000), *Types and general characteristics of ATM equipment*.

[16] ITU-T Recommendation I.732 (2000), *Functional characteristics of ATM equipment*.

[17] ITU-T Recommendation Q.2650 (1999), *Interworking between Signalling System No. 7 Broadband ISDN User Part (B-ISUP) and Digital subscriber Signalling System No. 2 (DSS2)*.

[18] ITU-T Recommendation Q.2660 (1999), *Interworking between Signalling System No. 7 Broadband ISDN User Part (B-ISUP) and Narrow-band ISDN User Part (N-ISUP)*.

[19] ITU-T Recommendation Q.2761 (1999), *Functional description of the B-ISDN User Part (B-ISUP) of Signalling System No. 7*.

[20] ITU-T Recommendation Q.2762 (1999), *General functions of messages and signals of the B-ISDN User Part (B-ISUP) of Signalling System No. 7*.

[21] ITU-T Recommendation Q.2763 (1999), *Signalling System No. 7 B-ISDN User Part (B-ISUP) – Formats and codes*.

[22] ITU-T Recommendation Q.2764 (1999), *Signalling System No. 7 B-ISDN User Part (B-ISUP) – Basic call procedures*.

[23] ITU-T Recommendation Q.2931 (1995), *Digital subscriber Signalling System No. 2 – User-Network Interface (UNI) layer 3 specification for basic call/connection control*.

[24] ITU-T Recommendations Q.2961.1 (1995) and Q.2961.2 (1997), *Digital subscriber Signalling System No. 2 – Additional traffic parameters*.

[25] ITU-T Recommendation Q.2962 (1998), *Digital subscriber Signalling System No. 2 – Connection characteristics negotiation during call/connection establishment phase*.

[26] ITU-T Recommendation Q.2963.1 (1999), *Digital subscriber Signalling System No. 2 – Connection modification: Peak cell rate modification by the connection owner*.

3 Abbreviations and terminology

3.1 Abbreviations

This Recommendation uses the following abbreviations:

AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ABT	ATM Block Transfer
ACR	Allowed Cell Rate
ADT	ACR Decrease Time
ATC	ATM Transfer Capability
ATM	Asynchronous Transfer Mode
ATM_PDU	ATM Protocol Data Unit
AUU	ATM user to ATM user indication [ITU-T Rec. I.361]
BCR	Block Cell Rate
BECN	Backward Explicit Congestion Notification

B-ISDN	Broadband ISDN
B-NT1	Broadband Network Termination 1
B-NT2	Broadband Network Termination 2
B-TE	Broadband Terminal Equipment
CAC	Connection Admission Control
CBR	Constant Bit Rate
CCR	Current Cell Rate
CDV	Cell Delay Variation
CEQ	Customer Equipment
CF-GCRA	Conforming Cell F-GCRA (Appendix XI)
CI	Congestion Indication
CLP	Cell Loss Priority (bit)
CLR	Cell Loss Ratio
CRC	Cyclic Redundancy Check
CRF(VC)	Virtual Channel Connection-Related Function
CRF(VP)	Virtual Path Connection-Related Function
CS	Convergence Sublayer
CT	Connection Termination
CTD	Cell Transfer Delay
DBR	Deterministic Bit Rate
DGCRA	Dynamic GCRA
DIR	Direction
DT	Delayed Transmission
ECR	Explicit Cell Rate
EDC	Error Detection Code
EFCI	Explicit Forward Congestion Indication
F-GCRA	Frame based Generic Cell Rate Algorithm
FIFO	First-In-First-Out
FMBS	Frame Mode Bearer Service
FRM	Fast Resource Management
GCRA	Generic Cell Rate Algorithm
GFC	Generic Flow Control
GFR	Guaranteed Frame Rate
IACR	Initial Allowed Cell Rate
IBT	Intrinsic Burst Tolerance
INI	Inter-Network Interface
IT	Immediate Transmission

ITT	Ideal Transmission Time
IWF	Interworking Function
LCT	Last Conformance Time
LIT	Last Increment Time
LVMT	Last Virtual Modification Time
LVST	Last Virtual Schedule Time
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
MFS	Maximum Frame Size
NE	Network Element
NI	No Increase
NPC	Network Parameter Control
NRM	Network Resource Management
N_{RM}	For each forward RM cell, N_{RM} is the maximum number of in-rate cells (including this particular RM cell) an ABR source may send
OAM	Operation And Maintenance
PACR	Potential Allowed Cell Rate
PC	Priority Control
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PEI	Peak Emission Interval
PHY	Physical Layer
PM	Performance Monitoring
PTI	Payload Type Indicator
QoS	Quality of Service
RDF	Rate Decrease Factor
RIF	Rate Increase Factor
RM	Resource Management
SAP	Service Access Point
SBR	Statistical Bit Rate
SCR	Sustainable Cell Rate
SDU	Service Data Unit
SN	Sequence Number
TAT	Theoretical Arrival Time
TBE	Transient Buffer Exposure
TPT	Transmission Path Termination
UNI	User-Network Interface

UPC	Usage Parameter Control
VBR	Variable Bit Rate
VCC	Virtual Channel Connection
VCCT	Virtual Channel Connection Termination
VCI	Virtual Channel Identifier
VCLT	Virtual Channel Link Termination
VD	Virtual Destination
VPC	Virtual Path Connection
VPCT	Virtual Path Connection Termination
VPI	Virtual Path Identifier
VPLT	Virtual Path Link Termination
VS	Virtual Source
VSA	Virtual Scheduling Algorithm

3.2 Terminology

This Recommendation defines the following terms:

3.2.1 ATM block: See 6.6.

3.2.2 ATM layer Resource Management (RM): See 4.3.

3.2.3 Connection Admission Control (CAC): See 4.3.

3.2.4 conformance: Conformance is the application, at a given standardized interface, of one or more criterion to a cell, an ATM block or a frame.

3.2.5 congestion: In B-ISDN, 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 either the network performance objectives or the negotiated QoS commitments for the already established connections and/or for the new connection requests. (See 4.1.)

3.2.6 connection traffic descriptor: The collection of source traffic descriptor and associated tolerance parameters to capture the traffic characteristics of an ATM connection at a standardized interface. See 5.1.2.

3.2.7 feedback controls: See 4.3.

3.2.8 frame cell sequence: See 7.2.5.

3.2.9 frame discard: See 7.2.5.

3.2.10 inter-network interface: See Figure 1.

3.2.11 Network Resource Management (NRM): See 4.3.

3.2.12 priority control: See 4.3.

3.2.13 standardized interface: A UNI (see ITU-T Rec. I.413) or an INI.

3.2.14 source traffic descriptor: A collection of traffic parameters to capture intrinsic characteristics of a source. See 5.1.2.

3.2.15 traffic contract: A traffic contract specifies the negotiated characteristics of a connection. See 5.3.

3.2.16 traffic descriptor: A set of traffic parameters and associated tolerances to capture the characteristics of an ATM connection. See 5.1.2.

3.2.17 traffic parameter: A traffic parameter describes an inherent characteristic of a traffic source. A traffic parameter may be quantitative or qualitative. See 5.1.1.

3.2.18 U**sage/Network Parameter Control (UPC/NPC)**: See 4.3.

3.2.19 user: An entity that agrees on a traffic contract at the UNI and exchanges ATM cells with the network at the UNI. This is for further study.

3.2.20 user data cell (on a VPC): Any CLP = 0 and CLP = 1 cell generated by the user, exclusive of F4 OAM cells and RM cells with VCI = 6 and PTI = 110.

3.2.21 user data cell (on a VCC): Any CLP = 0 and CLP = 1 cell generated by the user, exclusive of F5 OAM cells and RM cells with PTI = 110.

3.2.22 user generated cell (on a VPC): Any user data, user OAM or user RM cell.

3.2.23 user generated cell (on a VCC): Any user data, user OAM or user RM cell.

3.2.24 user OAM cell (on a VPC): Any F4 end-to-end OAM cell on the VPC generated by the user.

3.2.25 user OAM cell (on a VCC): Any F5 end-to-end OAM cell on the VCC generated by the user.

3.2.26 user RM cell (on a VPC): Any RM cell with VCI = 6 and PTI = 110 on the VPC generated by the user.

3.2.27 user RM cell (on a VCC): Any RM cell with PTI = 110 on the VCC generated by the user.

3.3 External terminology

Term	**Acronym**	**Reference**
Cell delay variation	CDV	ITU-T Rec. I.356
Cell error ratio	CER	ITU-T Rec. I.356
Cell loss ratio	CLR	ITU-T Rec. I.356
Cell transfer delay	CTD	ITU-T Rec. I.356
Cell loss priority	CLP	ITU-T Rec. I.150

4 Introduction

4.1 General

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 and committed QoS. An additional role is to optimize the use of network resources.

In B-ISDN, 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 either the network performance objectives or the negotiated QoS commitments for the already established connections and/or for the new connection requests.

In general, congestion can be caused by:

– unpredictable statistical fluctuations of traffic flows;

– fault conditions within the network.

Congestion is to be distinguished from the state where buffer overflow is causing cell losses, but still meets the negotiated quality of service.

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 Recommendation defines a set of traffic control and congestion control capabilities.

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

In this Recommendation and in consistency with ITU-T Recs I.150 and I.113, ATM connections are unidirectional. Two ATM connections are associated for the two directions of a communication and identified by the same VPI/VCI at a given interface. It should be noted that traffic control procedures that apply to a unidirectional connection (forward direction) may imply cell flows on the associated connection in the other direction (backward direction). Also, traffic control procedures may use cell flows on the forward direction to control the backward direction.

In this Recommendation, QoS requirements refer to QoS classes requested by the user. QoS commitments are referred to where the network actually commits to meet upper bounds for some QoS parameters, assuming the user generated cell flow conforms to a traffic contract. QoS indications pertain when there is no such upper bound specified for a given QoS parameter in the QoS class negotiated in the traffic contract, e.g., in cases where traffic engineering rules are used to operate the network and do not allow for commitments to the user.

Segment OAM flows are not part of the traffic contract negotiated by the user. How they are handled is currently not specified in this Recommendation.

4.2 General objectives

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; these QoS classes are specified in ITU-T Rec. I.356.
– 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 that 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.

4.3 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. This framework is based on the fundamental concept of a traffic contract (see 5.3) that is negotiated between the user and the network and between networks when setting up a connection.

– 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 establishment phase (or during call renegotiation phase) in order to establish whether a virtual channel/virtual path connection request can be accepted or rejected (or whether a request for re-allocation can be accommodated). Choosing a path through the network is part of the connection admission control of the network.

- ATM layer Resource Management (RM) functions make use of resource management cells, e.g., to modify resources that are allocated to ATM connections.
- 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/Network Parameter Control (UPC/NPC) 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 and the network access, respectively. Their main purpose is to protect network resources from malicious as well as unintentional misbehaviour which can affect the QoS of other already established connections by detecting violations of negotiated parameter values or procedures and taking appropriate actions.
- Priority Control: Priority Control functions differentiate how cells are handled relative to each other by the network in terms of time priority (scheduling control) or discard priority.

As a general requirement, it is desirable that a high level of consistency be achieved between the above traffic control capabilities.

A specific subset of these generic functions together with relevant traffic parameters and parameter values, as well as appropriate control functions and procedures are combined to create an ATM Transfer Capability (ATC) (see clause 6). This Recommendation includes a set of such capabilities which is intended to meet the requirements of different sets of broadband applications.

4.4 A reference configuration for traffic control and congestion control

The following reference configuration is used for traffic control and congestion control (Figure 1).

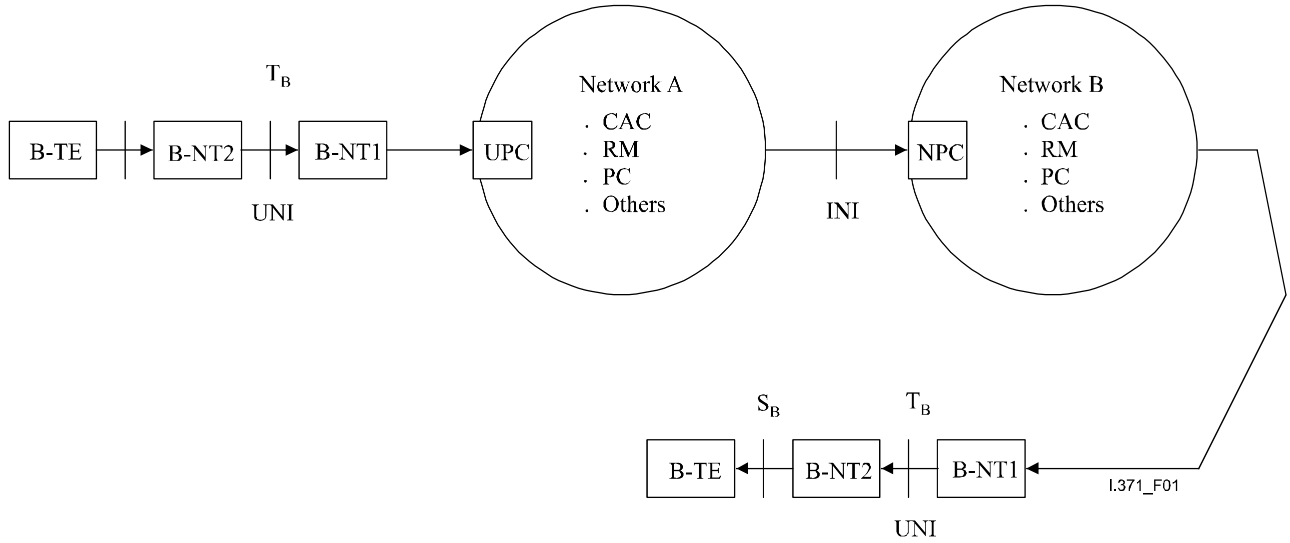

CAC Connection Admission Control
INI Inter-Network Interface
NPC Network Parameter Control
PC Priority Control
RM Resource Management
UNI User-Network Interface
UPC Usage Parameter Control

NOTE 1 – NPC may apply as well at some intra-network interfaces.
NOTE 2 – The arrows indicate the direction of the cell flow.
NOTE 3 – Feedback controls by means of RM cells go in the backward direction.
NOTE 4 – B-NT1 does not have any ATM layer function (see ITU-T Rec. I.413). It is included in the figure for the sake of completeness and consistency with ITU-T Rec. I.413.

NOTE 5 – In this Recommendation, UNI refers to the interface at the T_B reference point.

Figure 1/I.371 – Reference configuration for traffic control and congestion control

4.5 Events, actions, time-scales and response times

Figure 2 illustrates the time-scales over which various traffic control and congestion control functions operate. The response time defines how quickly the controls react. For example, cell discarding can react on the order of the insertion time of a cell. Similarly, feedback controls can react on the time-scale of round trip propagation times. Since traffic control and resource management functions are needed at different time-scales, no single function is likely to be sufficient.

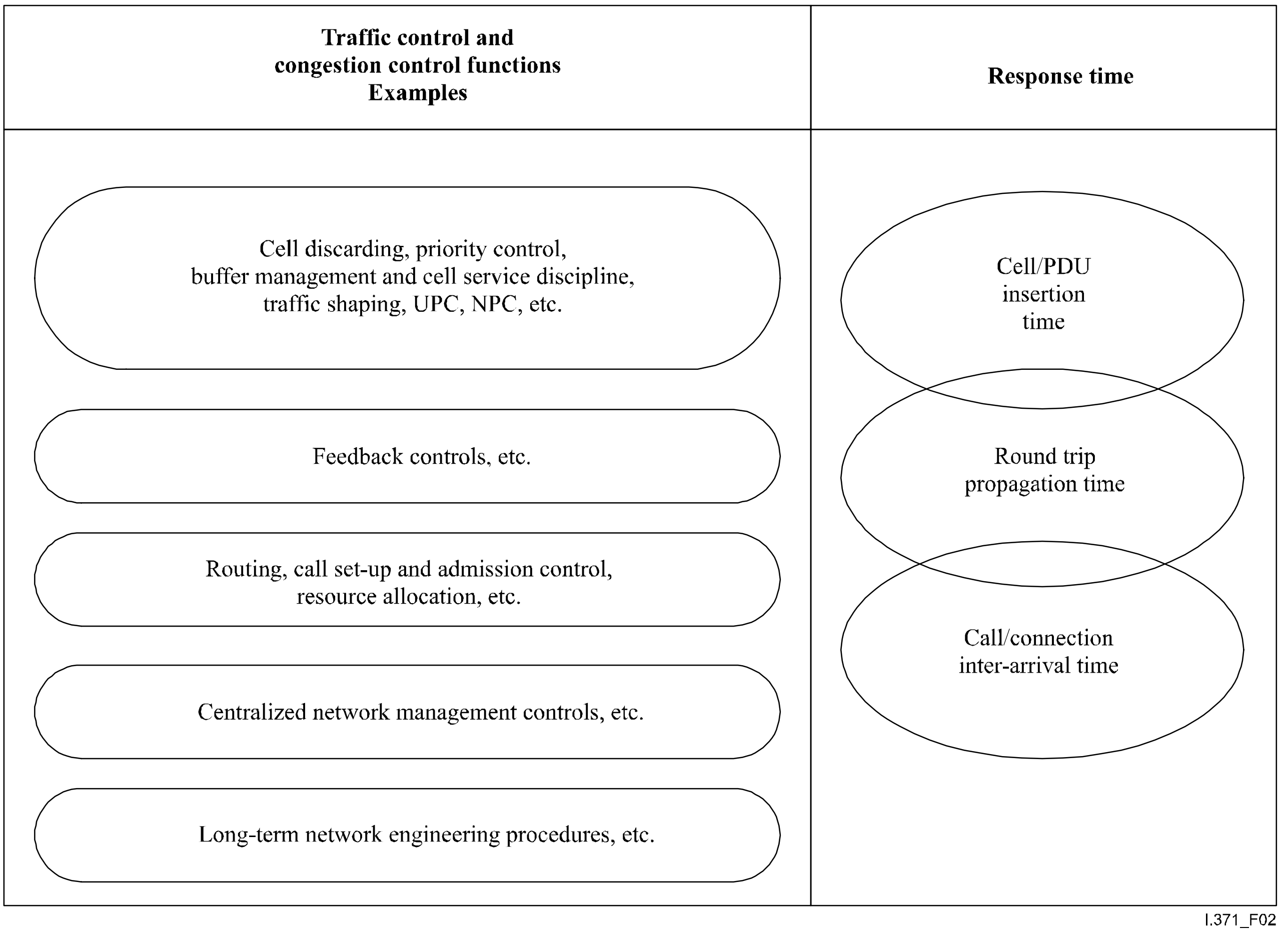

Figure 2/I.371 – Control response times

4.6 Quality of service and network performance

The B-ISDN shall be able to meet different QoS requirements at the ATM layer. These QoS requirements are specified in terms of objective values of some of the network performance parameters specified in ITU-T Rec. I.356. These network performance parameters include Cell Loss Ratio (CLR), Cell Transfer Delay (CTD) and Cell Delay Variation (CDV). In this Recommendation, QoS commitments referring to delays include cell transfer delay and 2-point cell delay variations (see ITU-T Rec. I.356).

Even though the QoS requirements of the B-ISDN users may vary over a continuous spectrum of values, a network can only handle a restricted set of QoS classes corresponding to specific objective values of the relevant network performance parameters. Use of Cell Loss Priority (CLP) is addressed in 5.3.3.

The specification of the different QoS classes in terms of objective values for the relevant network performance parameters is out of the scope of this Recommendation and will be provided in ITU-T Rec. I.356. When the conformance relative to a particular parameter is specified as being "unspecified", the ITU-T establishes no objective for this parameter and any default I.356 objective can be ignored. When the objective for a parameter is set to "unspecified", performance with respect to this parameter may, at times, be arbitrarily poor. Network operators may unilaterally elect to assure some minimum quality level for the unbounded parameters, but the ITU-T will not recommend any such minimum.

Negotiation of a specific ATM layer QoS class takes place at connection establishment. The resulting QoS class is part of the traffic contract (see 5.3). It is a commitment for the network to meet the requested quality of service as long as the user complies with the traffic contract. If the user violates the traffic contract, the network need not respect the agreed QoS (see 5.3).

The QoS class in the traffic contract only captures the ATM layer QoS. It is the role of the upper layers, including the AAL, to translate this ATM layer QoS to any specific application requested QoS.

5 Traffic parameters and descriptors

Many aspects of traffic can be described in terms of qualitative and quantitative parameters. This clause defines a number of specific parameters to describe and specify traffic in accordance with a traffic contract. Traffic parameters describe traffic characteristics of an ATM connection. Traffic parameters are grouped into source traffic descriptors to capture intrinsic characteristics of a source. The source traffic descriptor and associated tolerance parameters are grouped into connection traffic descriptors to capture the traffic characteristics of ATM connections at a standardized interface. Additional parameters have been defined for the operation of ABR. Other ATCs beside ABR do not rely on any parameters outside the connection traffic descriptor.

5.1 Definitions

This Recommendation defines the following terms.

5.1.1 Traffic parameters

A traffic parameter is a specification of a particular traffic aspect. It may be qualitative or quantitative.

Traffic parameters may for example describe peak cell rate, average cell rate, cell delay variation tolerances, burstiness, peak duration.

Some of the above-mentioned parameters are inter-dependent (e.g., the burstiness with the average and peak cell rate).

5.1.2 Traffic descriptors

The ATM traffic descriptor is the generic list of traffic parameters which can be used to capture the traffic characteristics of an ATM connection.

A source traffic descriptor is the set of traffic parameters belonging to the ATM traffic descriptor used during the connection establishment to capture the intrinsic traffic characteristics of the connection requested by the source.

A connection traffic descriptor is the set of traffic parameters belonging to the ATM traffic descriptor used during the connection establishment to capture the traffic characteristics of the connection at a given standardized interface. The connection traffic descriptor consists of the source traffic descriptor and the associated CDV tolerances applicable at that interface (see 5.3.5).

Connection admission control procedures will use the source traffic descriptor and associated CDV tolerances as included in the connection traffic descriptor to accept or reject connection requests.

A description of the characteristics of the traffic that any given requested connection may offer has to be provided by the user at the connection establishment phase.

5.2 Requirements

Any traffic parameter to be involved in a source traffic descriptor should:

– be understandable by the user or the terminal; conformance should be possible;

– participate in resource allocation schemes meeting network performance requirements;

– be enforceable by the UPC and NPC.

These criteria should be respected since users may have to provide these traffic parameters at connection establishment. In addition, these traffic parameters 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/NPC to maintain network performance in case of non-conformance.

5.3 Traffic contract

5.3.1 Traffic contract definition

CAC and UPC/NPC procedures require the knowledge of certain parameters to operate efficiently: they should take into account the ATM transfer capability (see clause 6), the source traffic descriptor, the requested QoS class and the CDV tolerances (see 5.3.5) in order to decide whether the requested connection can be accepted.

An ATM transfer capability, a source traffic descriptor, associated CDV tolerances and a QoS class are declared by the user at connection establishment by means of signalling or subscription.

The selected ATM transfer capability (including associated procedures and options such as tagging), the source traffic descriptor, the QoS class for any given ATM connection and the CDV tolerances allocated to the Customer Equipment (CEQ) agreed upon at connection establishment define the traffic contract at the T_B reference point. A similar contract applies at the Inter-Network Interface (INI). CDV tolerances belonging to a traffic contract at an INI account for the CDV introduced by the upstream portion of the connection, including the CEQ.

For a given ATM connection, the source traffic descriptor belonging to the traffic contract and all parameter values of this source traffic descriptor are the same at all standardized interfaces along the connection.

In order for QoS commitments to be met, a conformance definition is specified at T_B for any given ATM transfer capability (see clause 6). A conformance definition also pertains at each standardized inter-network interface. A traffic contract may apply to a VP or to a VC connection. As a consequence, conformance definition at an interface applies at the level where the traffic contract is defined (VP or VC). Additionally, a traffic contract for a connection may imply a cell flow on the connection of the reverse direction of a communication. In such a case, a conformance definition also pertains for the reverse connection.

The Connection Admission Control (CAC) and Usage/Network Parameter Control (UPC/NPC) procedures are operator specific. Once the connection has been accepted, the value of the CAC and UPC/NPC parameters are set by the network on the basis of the network operator's policy.

NOTE 1 – All ATM connections handled by network connection related functions (CRF, see ITU-T Recs I.311 and I.732) have to be declared.

NOTE 2 – Individual VCCs inside a user end-to-end VPC are neither declared nor enforced at the UPC (VP) and hence an ATM layer QoS can only be assured to the VPC.

5.3.2 Traffic contract and quality of service

ITU-T Rec. I.356 specifies QoS objectives for the end-to-end connection and apportionment rules establishing QoS objectives for each standardized connection portion. ATM layer QoS is a long-term commitment. The ATM layer QoS is assured to all cells when all user cells (or blocks, see 6.6) are conforming to the relevant conformance tests.

It should be noted that QoS is a 2-point concept applying to a connection or a connection portion, whilst conformance is a 1-point concept applying at an interface (see ITU-T Rec. I.356). As a result, conformance definition at an interface applies to cell flows that are submitted at this interface; the upstream portion of the connection is globally responsible for conformance at that interface, i.e., the user at the UNI, the user and upstream networks at an INI.

The ATM layer QoS need not be delivered to any connection that a network provider has determined to be non-compliant. Network providers can unilaterally decide when a connection containing non-conforming cells (or blocks) is non-compliant. The exact definition of non-compliance is a network provider responsibility. Even when a connection is non-compliant, a network operator may choose to offer some QoS commitments, see for example 6.4.6.

5.3.3 Traffic contract and cell loss priority

Depending on the ATM transfer capability, a user may request for an ATM connection a QoS class that involves two levels of priority, as indicated by the CLP bit value. The intrinsic traffic characteristics of both cell flow components have to be characterized in the source traffic descriptor. This is by means of a set of traffic parameters associated with the CLP = 0 component and a set of traffic parameters associated with the aggregate CLP = 0 + 1 cell flow component.

The network may provide a cell loss ratio objective for each of the components (CLP = 0 and CLP = 0 + 1) of an ATM connection. The traffic contract specifies the particular CLR objectives from those offered by the network operator for each of the ATM connection components. This Recommendation currently limits the use of that capability to two cases:

– there is a CLR objective for the CLP = 0 + 1 cell flow, irrespective of the CLP bit value;

– there is a CLR objective for the CLP = 0 cell flow, whilst CLR for the CLP = 0 + 1 cell flow is unspecified.

5.3.4 Traffic contract and tagging

Whether tagging (see 7.2.3.6) is applicable to a connection depends on the ATC negotiated for the connection (see clause 6). When tagging is applicable to a connection, it can be performed at all standardized interfaces along the connection. Networks may tag cells only if the ATM transfer capability specifies that tagging applies. In this Recommendation, the only ATCs for which tagging applies are SBR3 and GFR2.

5.3.5 Impact of cell delay variation on UPC/NPC and resource allocation

ATM layer functions (e.g., cell multiplexing) may alter the traffic characteristics of ATM connections by introducing cell delay variation. 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. Therefore, some randomness affects the time interval between reception of ATM cell data-requests at the end-point of an ATM connection to the time that an ATM cell data-indication is received at the UPC/NPC. Besides, AAL multiplexing may originate CDV (e.g., when a 2-layer coded video signal would consist in two flows transferred by ATM cells that would differ by the CLP bit).

Origins of cell delay variation are illustrated on Figure 3.

The conformance definition to the source traffic descriptor at a given interface (e.g., a UNI or an inter-network interface), as well as the execution of the UPC/NPC functions, require that the CDV attributed to the upstream portion of the connection and affecting each of the relevant parameters be specified.

UPC/NPC should accommodate the effect of the maximum CDV allowed on ATM connections which arises from the accumulated CDV attributed to upstream subnetworks (including CEQ).

In general, each component of a connection (e.g., the CLP = 0 user data component, the CLP = 0 + 1 user data component, the user OAM component and the user RM component) may require the specification of a different value of CDV tolerance for each of its traffic parameters (e.g., peak cell rate, sustainable cell rate parameter set). Therefore, the number of CDV tolerance values relevant to a connection depends on the source traffic descriptor of the connection and, finally, on the ATM transfer capability required by the connection. It is expected that only a subset of the possible tolerances would be needed. The modality of negotiation of the relevant CDV tolerance values between the user and the network and between two networks (e.g., on a subscription basis or on a per connection basis) depends on the traffic parameter the CDV tolerance refers to. More detailed information is provided respectively in 5.4.1.3 for the peak cell rate and 5.4.2.3 for the sustainable cell rate.

Traffic shaping partially compensates for the effects of CDV and results in a reduced CDV tolerance to be applied at subsequent interfaces on the ATM connection. Examples of traffic shaping mechanisms are re-spacing cells of individual ATM connections according to their peak cell rate or suitable queue service schemes.

The definition of a source traffic descriptor and the standardization of maximum allowable CDV tolerances may not be sufficient for a network to allocate resources properly. When allocating resources, the network should take into account the worst-case traffic passing through UPC/NPC in order to avoid impairments to other ATM connections. This worst-case traffic depends on the specific implementation of the UPC/NPC. The trade-offs between UPC/NPC complexity, worst-case traffic and optimization of network resources are made at the discretion of network operators. The quantity of available network resources and the network performance to be provided for meeting QoS requirements can influence these trade-offs.

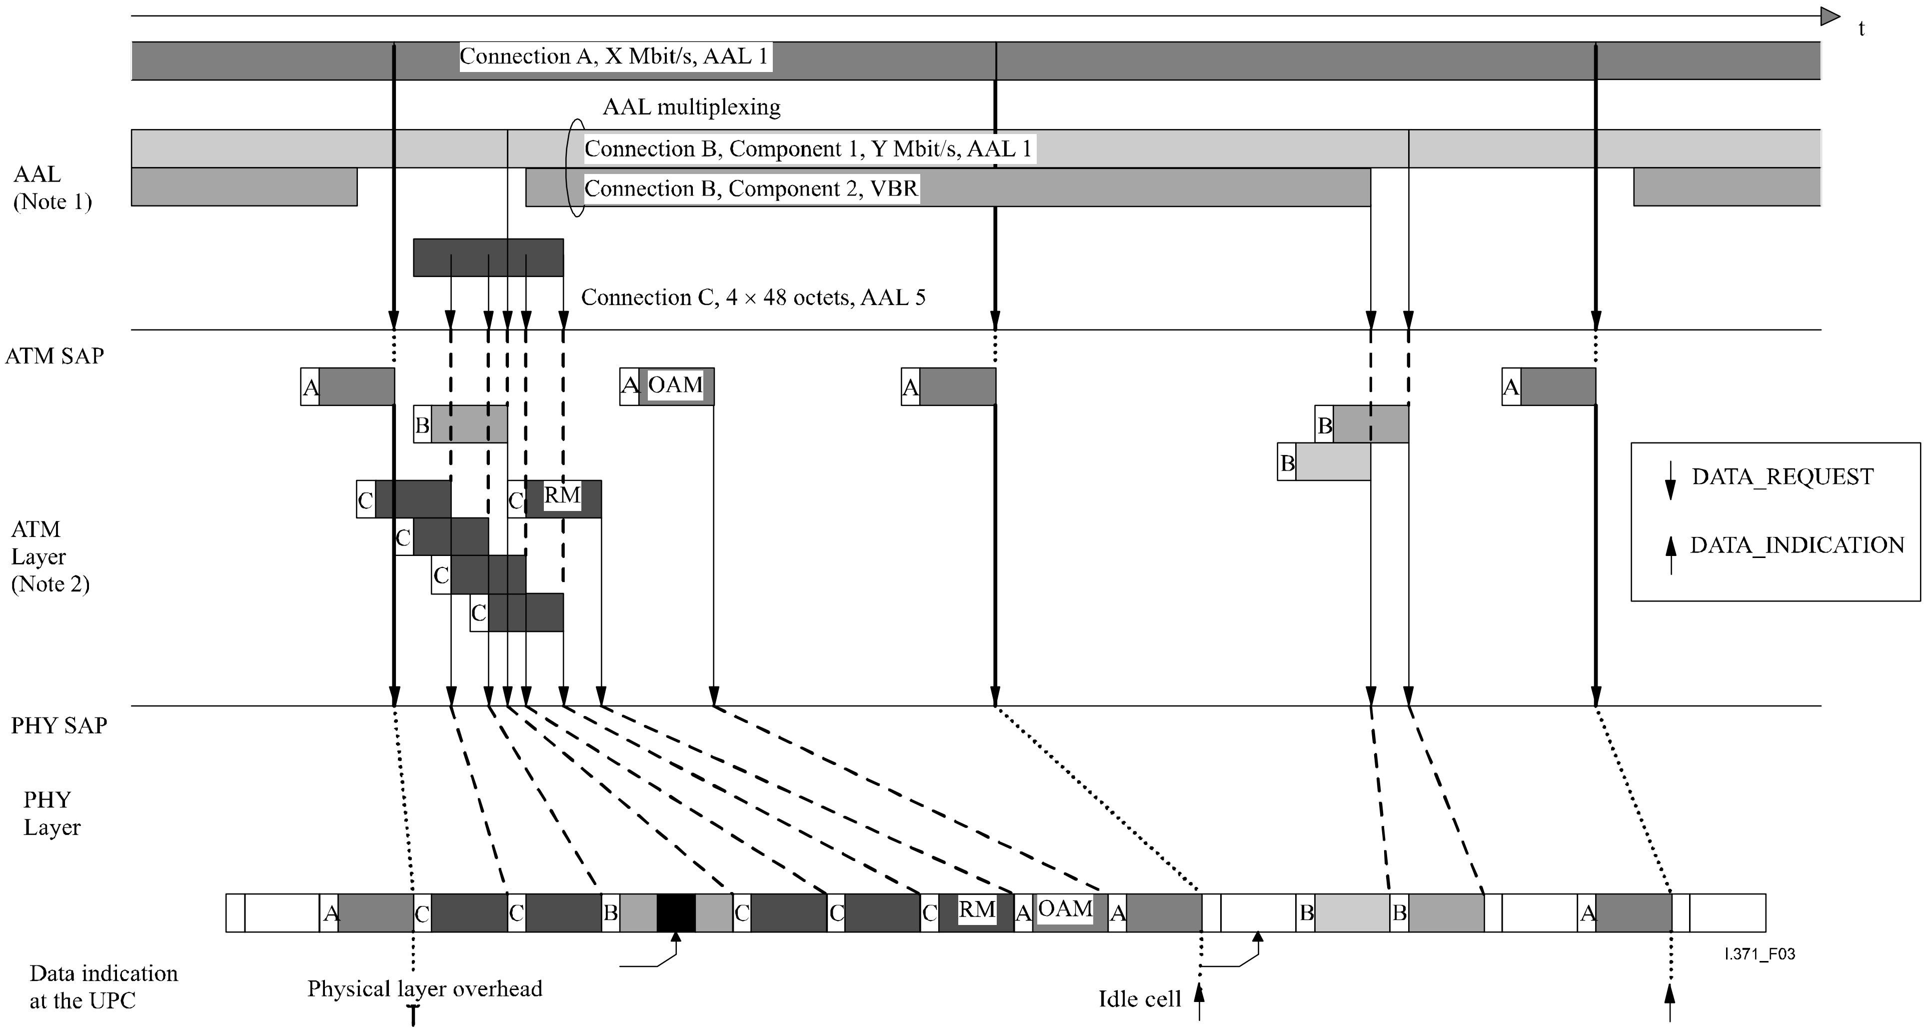

NOTE 1 – ATM SDUs are accumulated at the upper layer service bit rate. Besides, CDV may also originate in AAL multiplexing.
NOTE 2 – GFC delay and delay variation is part of the delay and delay variation introduced by the ATM layer.
NOTE 3 – CDV may also be introduced by the network because of random queueing delays which are experienced by each cell in concentrators, switches and cross-connects.

Figure 3/I.371 – Origins of cell delay variation

Figure 3 is intended only to illustrate functions generating CDV. It does not imply any relationship to the definition of traffic parameters (e.g., peak emission interval).

5.4 Traffic parameter specifications

Conformance definition to a given traffic contract relies on a non-ambiguous specification of traffic parameters. Relevant traffic parameters and conformance to any given traffic parameter may depend on the ATM transfer capability that applies to a connection (see clause 6).

When relevant, a tolerance will be specified in combination with each traffic parameter at an interface where a conformance definition applies in order to account for the impact of upstream multiplexing functions on source traffic parameter values.

For a traffic parameter specification to be non-ambiguous, a list of discrete values that any given parameter and CDV tolerance may take, expressed in the appropriate unit, is standardized in this Recommendation. In addition, translation rules are standardized, e.g., when a parameter needs more than one representation in more than one unit at the ATM layer or within management and control planes (e.g., translation from peak cell rate to peak emission interval, translation from maximum burst size to intrinsic burst tolerance).

Parameters referring to time intervals will be specified as subsets of a single generic list of values specified by a floating point coding scheme with a 10-bit mantissa and a 6-bit exponent as follows:

$$2^{e-32}\left(1+\frac{w}{2^{10}}\right)[\text{seconds}]$$

$$0 \le e \le 41$$

$$0 \le w \le 1023$$

This Recommendation defines the Peak Cell Rate (PCR, see 5.4.1) and the sustainable cell rate parameter set (SCR/IBT, see 5.4.2). Additional parameters based upon these definitions also appear in ATM transfer capabilities specifications when relevant (see clause 6). Additional standardized parameters which may be specified in the future should provide for a significant improvement of network utilization.

Peak cell rate is a mandatory traffic parameter to be explicitly or implicitly declared in any source traffic descriptor. In addition to the peak cell rate of an ATM connection, it is mandatory for the user to declare either explicitly or implicitly the corresponding cell delay variation tolerance τ_{PCR} at the UNI within the relevant traffic contract.

Reference configuration and equivalent terminal for traffic parameter specification

The reference configuration of Figure 4 applies to the specification of traffic parameters and associated tolerances at the UNI.

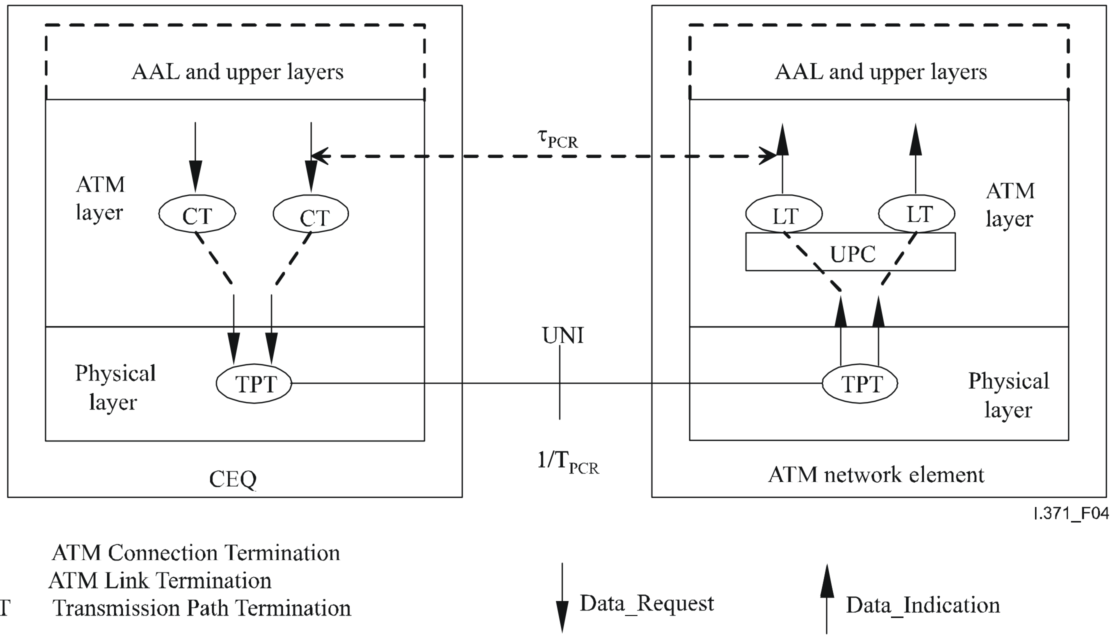

NOTE 1 – ATM connection termination may be a VP Connection Termination (VPCT) or a VC Connection Termination (VCCT).
NOTE 2 – ATM link termination may be a VP Link Termination (VPLT) or a VC Link Termination (VCLT).
NOTE 3 – For further details on TPT, VPLT, VPCT, VCLT and VCCT, see ITU-T Recs I.731 and I.732.

Figure 4/I.371 – Reference configuration for the specification of traffic parameters

When an ATM connection consists in a number of connection components (e.g., VCCs within a VPC) that are generated by different sources at different locations, Figure 5 illustrates the cell rate of an ATM connection and its associated CDV tolerance by the means of an equivalent source and of an equivalent terminal. In this figure, ATM PDU Data_Requests from individual sources are virtually merged and spaced out at the emission interval corresponding to the cell rate of the connection. The ATM PDU Data_Requests resulting from the equivalent source would be ideally conforming to a GCRA(T,0) (see Annex A). Cell delay variation, as produced by the different terminal equipment (AAL or ATM layer multiplexing, physical layer functions, incorporated in an equivalent terminal) and by the customer equipment (CEQ), is captured by a CDV tolerance τ_{UNI} at the UNI, so that the cell flow at the UNI is conforming to a GCRA(T, τ_{UNI}). Similarly, a CDV tolerance τ_{INI} accounts for the CDV introduced by the upstream portion of the connection at a given INI.

It should be noted with respect to Figure 5 that the equivalent source may consist of a single traffic source and a virtual shaper (case where there is an intrinsic tolerance attached to the source, see for example sustainable cell rate parameter set) or of a single traffic source without a shaper (case where the source actually produces ATM PDU Data_Request at interval T).

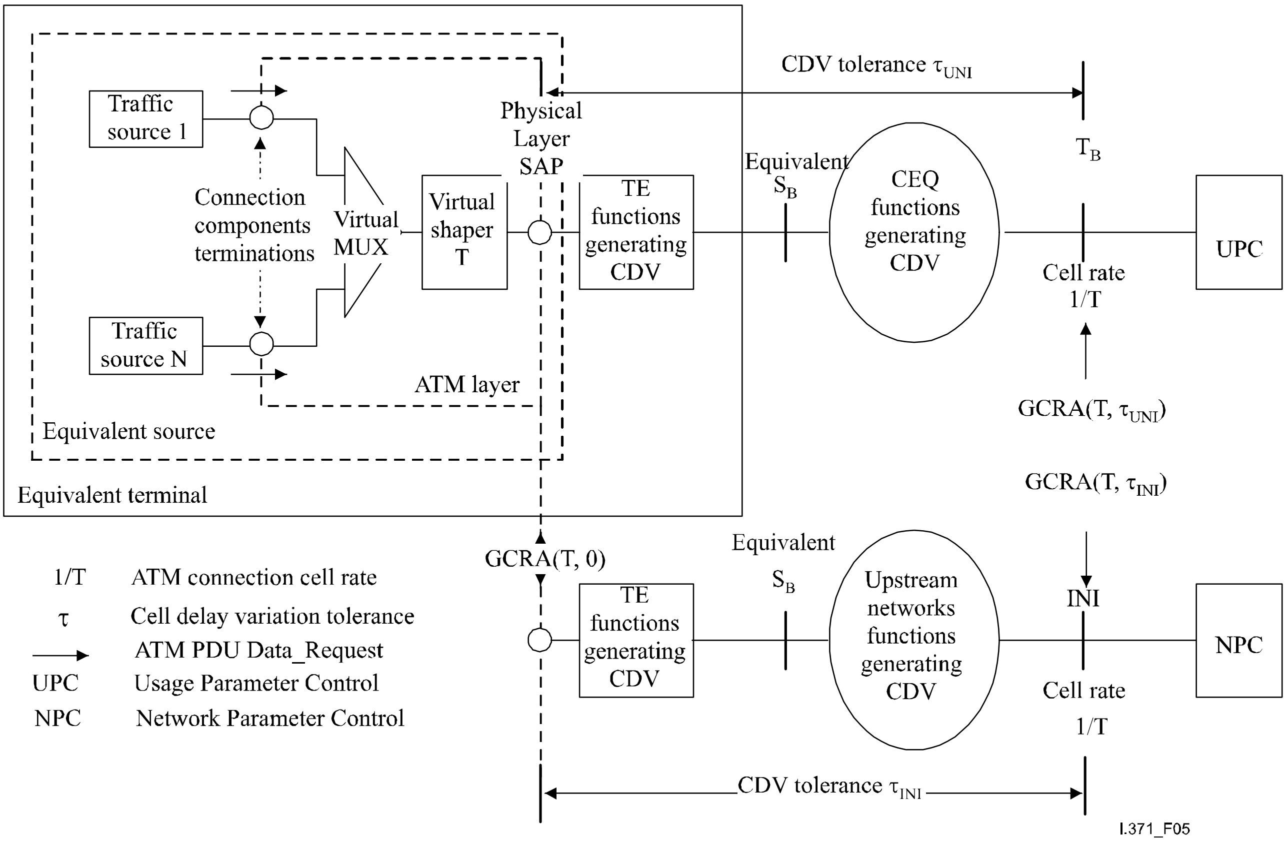

NOTE – This figure is illustrative only and does not imply any implementation.

Figure 5/I.371 – Equivalent source and equivalent terminal for the definition of a cell rate of an ATM connection

5.4.1 Peak cell rate

The following definition applies to ATM connections for any ATM transfer capability defined in clause 6.

The peak cell rate in the source traffic descriptor specifies an upper bound on the traffic that can be submitted on an ATM connection. Enforcement of this bound by the UPC/NPC allows the network operator to allocate sufficient resources to ensure that the performance objectives (e.g., for cell loss ratio) can be achieved.

The peak cell rate value as negotiated and agreed upon at connection establishment or subsequently modified via signalling or network management procedures shall be the same along a given ATM connection. The CDV tolerance τ_{PCR} associated to the peak cell rate may be different at different interfaces along that ATM connection. τ_{PCR} may not be the same for all cell flows of an ATM connection at a given interface.

5.4.1.1 Peak cell rate definition for a VPC/VCC

Location

– at the physical layer SAP for an equivalent terminal representing the VPC/VCC (see Figure 5); or

– equivalently at the Transmission Path Termination (TPT) for the reference configuration representing the VPC/VCC, see Figure 4.

Basic event

Request to send an ATM_PDU.

Definition

The peak cell rate of the ATM connection is the inverse of the minimum inter-arrival time T_{PCR} between two basic events defined above. T_{PCR} is the peak emission interval of the ATM connection.

On a terminal with a single AAL entity and with neither ATM layer OAM nor RM flows, assuming that no shaping is performed at the ATM layer, location and basic event are equivalent to the following ones:

Location

- at the ATM layer SAP for an equivalent terminal representing the VPC/VCC (see Figure 5); or
- equivalently at the Virtual Path/Virtual Channel Connection Termination (VPCT/VCCT) for the reference configuration representing the VPC/VCC, see Figure 4.

Basic event

Request to send an ATM_SDU.

Referring to ATCs currently specified in clause 6 and in order to properly allocate resources to a VPC/VCC, a peak cell rate as defined above has to be specified for each component of the ATM connection, i.e., the user data (CLP = 0 + 1) component, the user OAM component and the RM component. Some components may be aggregated (e.g., user OAM with user data). For each component, a corresponding CDV tolerance τ_{PCR} accounts for cell delay variation (see 5.3.5).

Examples of application of the peak cell rate definition to specific configurations are in Appendix I.

5.4.1.2 Specification of peak cell rate

The following list of joint values of peak cell rates and Peak Emission Intervals (PEIs) defines the ATM layer peak cell rate granularity, which is used for conformance definition. A requirement relating to the accuracy of cell rate control of a UPC/NPC function is defined in 7.2.3.2.1.

Specification of the peak cell rate values

The following formula provides for a list of 16 384 peak cell rate values Λ_{PCR} ranging between 1 cell/s and 4.29077 Gcell/s. A floating point coding scheme with a 9-bit mantissa and a 5-bit exponent is used. The relative difference between any pair of successive values is quasi-constant over the full range and always smaller than .19%.

$$\Lambda_{PCR} = 2^{m_{PCR}} \cdot \left(1 + \frac{k_{PCR}}{512}\right) \text{cells per second}$$

$$0 \le m_{PCR} \le 31$$

$$0 \le k_{PCR} \le 511$$

Specification of peak emission interval values

The following formula gives the corresponding list of 16 384 peak emission interval values T_{PCR} ranging between .9995 seconds and 2.33 10^{-10} seconds. Relative difference between any pair (same values of m_{PCR} and k_{PCR}) of peak cell rate and 1/peak emission interval is smaller than .0977%. This list is a subset of the generic list specified above for time intervals. With respect to the 9-bit mantissa plus 5-bit exponent coding scheme for Λ_{PCR}, one extra bit is needed for the exponent to code the sign and one extra bit for the mantissa in order to increase the precision of the coding because of the non-linearity of ($x \rightarrow 1/x$).

$$T_{PCR} = 2^{-(m_{PCR}+1)} \cdot \left(1 + \frac{1023 - k'_{PCR}}{1024}\right) \text{second}$$

$$k'_{PCR} = \left\lfloor \frac{2047 k_{PCR} - 512}{k_{PCR} + 512} \right\rfloor + 1$$

$$0 \le m_{PCR} \le 31$$

$$0 \le k_{PCR} \le 511$$

where $\lfloor x \rfloor$ stands for rounding down to the nearest integer value.

The coding scheme has been designed so that any peak cell rate value is always less than its corresponding 1/peak emission interval value.

The negotiated PCR_{sig} value supported by signalling will be rounded up to the nearest ATM layer PCR value within the list of values specified for conformance testing; this may be done using the following formulae as long as $PCR_{sig} > 0$:

$$m_{PCR} = \left\lceil \log_2\left(\frac{PCR_{sig}}{1023}\right) + 9 \right\rceil$$

$$k_{PCR} = \left\lceil \frac{PCR_{sig}}{2^{m_{PCR}-9}} - 512 \right\rceil$$

where $\lceil x \rceil$ stands for rounding up to the nearest integer value.

5.4.1.3 Cell delay variation tolerance specification for peak cell rate

The CDV tolerance referring to the user data component can be declared either explicitly (e.g., by conveying the value in the signalling message on a connection basis) or implicitly. Implicit declaration is achieved by specifying the characteristics of the CDV at a given interface (e.g., UNI or INI) on a subscription basis or by means of mutual agreements between operator and user or between operators.

Characterization of the CDV tolerance at a given interface shall take into account the functions available at that interface. For the time being two extreme cases have been identified:

– Stringent requirement on CDV tolerance: a connection request should not be denied solely on the basis of a CDV tolerance requirement if this CDV tolerance requirement is less than or equal to τ_{PCR} where τ_{PCR} is given by:

$$\frac{\tau_{PCR}}{\Delta} = \max\left[\frac{T_{PCR}}{\Delta}, \alpha\left(1 - \frac{\Delta}{T_{PCR}}\right)\right]$$

where:

T_{PCR} is the peak emission interval of the connection (expressed in seconds)

Δ is the cell transmission time (in seconds) at the interface link speed

α is a dimensionless coefficient; the suggested value is $\alpha = 80$

– Loose requirements on CDV tolerance: a large amount of CDV can be tolerated. In this case, only the specification of the maximum value of CDV tolerance τ_{MAX} that can be allocated to a connection is envisaged. τ_{MAX} shall be intended as the maximum amount of CDV that can be tolerated on the user data cell stream. τ_{MAX} is not specified in this Recommendation.

Between these two extreme cases, there are intermediate cases that could pertain at network interfaces and for which a default rule, based on PEI, could be specified.

The cases above do not preclude any operator to support different values of CDV tolerance, that can be specified at subscription times or by mutual agreements; in particular a given value of τ_{PCR} might be specified for all connections at the interface. In addition, CDV tolerance may be conveyed via signalling on a per connection basis.

The subset of the generic list specified for time intervals that will be used to select values for τ_{PCR} will be coded as follows:

$$\tau_{PCR} = 2^{e_{PCR}-32} \cdot 2^{9} \cdot \left(1 + \frac{w_{PCR} \cdot 2^{5}}{2^{10}}\right) \text{seconds}$$

$$0 \le e_{PCR} \le 31$$

$$0 \le w_{PCR} \le 31$$

This coding scheme should be used to support declaration of τ_{PCR} by signalling or management means.

Values of τ_{PCR} that will effectively be used out of that generic list is left to operators' decision.

5.4.2 Sustainable cell rate

The Sustainable Cell Rate (SCR) together with a parameter characterizing the maximum burst size at the peak cell rate (Intrinsic Burst Tolerance, IBT) are intended to describe VBR sources and allow for statistical multiplexing of traffic flows from such sources.

The definition of the sustainable cell rate (Λ_{SCR}) and intrinsic burst tolerance (τ_{IBT}) uses the reference algorithm in Annex A, referred to as the Generic Cell Rate Algorithm (GCRA). Intrinsic burst tolerance is supported by signalling in terms of a Maximum Burst Size (MBS).

The sustainable cell rate parameter set as negotiated and agreed upon at connection establishment or subsequently modified and conveyed by signalling shall be the same along a given ATM connection. The CDV tolerance τ'_{SCR} associated to the sustainable cell rate parameter set may be different at different interfaces along that ATM connection. Whether τ'_{SCR} is the same for all components of an ATM connection at a given interface is for further study.

5.4.2.1 Sustainable cell rate for a VPC/VCC

Location

– at the physical layer SAP for an equivalent terminal representing the VPC/VCC (see Figure 5); or

– equivalently at the Transmission Path Termination (TPT) for the reference configuration representing the VPC/VCC, see Figure 4.

Event

Request to send an ATM_PDU.

Definition

The sustainable cell rate denoted as Λ_{SCR}, and the intrinsic burst tolerance, denoted as τ_{IBT}, of an ATM connection are defined by the GCRA(T_{SCR}, τ_{IBT}) based on the arrivals of the basic event above. Λ_{SCR} is the inverse of T_{SCR}.

Sustainable cell rate and intrinsic burst tolerance belong to the ATM traffic descriptor.

For conformance definition at the UNI/INI, a tolerance τ'_{SCR} has to be added to the intrinsic burst tolerance τ_{IBT}. τ'_{SCR} accounts for the CDV introduced by multiplexing schemes at the cell level and at the burst level. One upper bound for τ'_{SCR} is the difference between the longest and the shortest cell transfer delays between the source and the UNI/INI of that connection. Also, τ'_{SCR} may be chosen to be a small quantile, e.g., 10^{-9}, of the possible delay variation.

When the peak cell rate is complemented by the sustainable cell rate parameter set (T_{SCR} and τ_{IBT}), the source traffic descriptor contains the peak cell rate, the sustainable cell rate and the intrinsic burst tolerance traffic parameters. In addition, the traffic contract should provide for the cell delay variation tolerance parameters τ_{PCR} (related to the peak cell rate) and τ'_{SCR} (related to the sustainable cell rate).

When the peak cell rate is complemented by the sustainable cell rate for an ATM connection, T_{SCR} is always larger than T_{PCR} (Λ_{SCR} smaller than Λ_{PCR}).

5.4.2.2 Specification of the sustainable cell rate and intrinsic burst tolerance

Values for T_{SCR} will use the same subset of the generic list of values and the same coding as that one specified for T_{PCR}; this includes translation rules from Λ_{SCR} and from signalling (see 5.4.1.2).

The intrinsic burst tolerance will use the same subset of the generic list of values and the same coding as the one specified for τ_{PCR} (see 5.4.1.3). Translation from the maximum burst size supported by signalling will use the following rule:

$$\tau_{IBT} = \lceil (MBS - 1)(T_{SCR} - T_{PCR}) \rceil \text{seconds}$$

where $\lceil x \rceil$ stands for the first value above x out of the generic list of values.

If the user has the knowledge of τ_{IBT} rather than of the maximum burst size, then the following rule applies:

$$MBS = 1 + \left\lfloor \frac{\tau_{IBT}}{T_{SCR} - T_{PCR}} \right\rfloor \text{cells}$$

where $\lfloor x \rfloor$ stands for rounding down to the nearest integer value.

Values of MBS or τ_{IBT} that will effectively be selected are left to operators' decision. However, MBS values declared by signalling should not result in a value for τ_{IBT} exceeding the maximum value achievable by the coding scheme for τ_{IBT}.

5.4.2.3 Cell delay variation tolerance specification for sustainable cell rate

The same coding scheme applies to τ'_{SCR} as to τ_{PCR}, see 5.4.1.3.

This coding scheme should be used to support declaration of τ'_{SCR} by signalling or management means.

Values of τ'_{SCR} that will effectively be used out of that generic list is left to operators' decision.

5.4.3 Specification of other traffic parameters

In addition to the peak cell rate (see 5.4.1) and the sustainable cell rate/intrinsic burst tolerance (see 5.4.2), the following other traffic parameters are used in this Recommendation.

– Minimum Cell Rate (MCR): A lower bound on the allowed cell rate for ABR sources; it is specified on a per-connection basis. (ABR.)
– Initial Allowed Cell Rate (IACR): An upper bound on the allowed cell rate for ABR sources at initialization.
– Rate Decrease Factor (RDF): The parameter that controls the decrease in the cell transmission rate for ABR sources.

– Rate Increase Factor (RIF): The parameter that controls the increase in the cell transmission rate for ABR sources.

– For each forward RM cell, N_{RM} is the maximum number of in-rate cells (including this particular RM cell) an ABR source may send.

– Fixed Round Trip Time (FRTT): An estimate of the minimum round trip time for the connection. (ABR.)

– Transient Buffer Exposure (TBE): The number of cells that the network would like to limit the source to sending during start-up periods, before the first RM cell returns. (ABR.)

– Minimum Cell Rate (MCR): For a GFR connection it is used (in conjunction with other parameters) to quantify the lower bound on the number of cells to which the committed QoS applies. It is specified on a per-connection basis.

– Maximum Frame Size (MFS): The maximum number of user generated cells in a frame that may be sent on a GFR connection.

5.4.4 Traffic characteristics relevant to ATCs

Table 1 lists traffic characteristics, including source traffic descriptor and CDV tolerances, that are relevant to each ATM transfer capability as defined in clause 6. An X in the table indicates that the traffic characteristic is relevant to the corresponding ATC.

Table 1/I.371 – Traffic characteristics relevant to ATCs

ATC reference	Parameter reference	DBR	SBR1	SBR2, SBR3	ABT/DT, ABT/IT	ABR	GFR
		Clause 6.4	Clause 6.5	Clause 6.5	Clause 6.6	Clause 6.7	Clause 6.8
PCR(0+1)	5.4.1	X	X	X	X	X (Note 4)	X
$\tau_{PCR}(0+1)$	5.4.1	X	X	X	X		X
SCR(0)	5.4.2			X			
$\tau_{IBT}(0)$	5.4.2			X			X (Note 9)
$\tau'_{SCR}(0)$	5.4.2			X			
SCR(0+1)	5.4.2		X		X (Note 3)		
$\tau_{IBT}(0+1)$	5.4.2		X		X (Note 3)		
$\tau'_{SCR}(0+1)$	5.4.2		X		X (Note 3)		
MCR(0)	5.4.3, 6.7.2					X (Note 6)	X
τ_1	6.7.5					X (Note 5)	
τ_2	6.7.5					(Note 7)	
τ_3	6.7.5					(Note 7)	
IACR(0)	6.7.2					X	
FRTT	6.7.3					X	
TBE	6.7.3					X	
RDF	6.7.3					X	
RIF	6.7.3					X	
Cell tagging	5.3.4			(Note 2)			(Note 2)
$\tau_{MCR}(0)$	6.8.2						X
Frame tagging	6.8						(Note 10)
MFS	6.8.2						X

Table 1/I.371 – Traffic characteristics relevant to ATCs

ATC reference	Parameter reference	DBR	SBR1	SBR2, SBR3	ABT/DT, ABT/IT	ABR	GFR
		Clause 6.4	Clause 6.5	Clause 6.5	Clause 6.6	Clause 6.7	Clause 6.8
PCR(RM), τ_{PCR} (RM)	5.4				X (Note 8)		
PCR(OAM), τ_{PCR} (OAM)	5.4	X (Note 1)			X (Note 1)		

NOTE 1 – Separate declaration of user OAM traffic characteristics is only possible for the DBR and ABT capabilities, and this separate declaration is optional (see 6.4). If this option is not chosen, user OAM cells are aggregated to user data cells when characterizing the offered traffic.

NOTE 2 – Cell tagging (see 7.2.3.6) applies to SBR3 (see 6.5.2) and GFR2 (see 6.8.1) only. SBR2 and SBR3 as well as GFR1 and GFR2 are identical except for the application of cell tagging.

NOTE 3 – Declaration of an (SCR, IBT) parameter set for ABT capabilities is optional. If this parameter set and a QoS class with specified CLR objectives is negotiated, there is a QoS commitment at the block level. If this parameter set is not negotiated, SCR is assumed to be 0 and there is no QoS commitments at the block level (see 6.6).

NOTE 4 – User generated traffic should be transmitted with the CLP bit set to 0, although user data traffic is specified by means of a PCR(0+1). Some user RM cells may be transmitted with the CLP bit set to 1 (see 6.7).

NOTE 5 – The same CDVT parameter τ_1 applies to ACR(0) rates ranging from PCR(0+1) to MCR(0).

NOTE 6 – User generated traffic should be transmitted with CLP set to 0. The MCR may be set to 0.

NOTE 7 – Negotiation of τ_2 and τ_3 is for further study.

NOTE 8 – It is necessary to define a default value of τ_{RM} at each standardized interface that is valid for any ABT connection.

NOTE 9 – The value of τ_{IBT} is derived from the parameters MBS, PCR and MCR (see 6.8.3.3).

NOTE 10 – Frame tagging applies to GFR2 only.

6 ATM transfer capabilities

6.1 General

6.1.1 Definition and requirements

An ATM transfer capability is intended to support an ATM layer service model and associated QoS through a set of ATM layer traffic parameters and procedures. The use of ATCs has both a user's perspective, wherein an ATC is seen as suitable for a given set of applications, and a network operator's perspective, wherein an ATC may provide gain through statistical multiplexing. An ATM transfer capability may include the specification of primitives to comply with, and of traffic control information to be exchanged through standardized interfaces.

In order for network providers to be able to make QoS commitments, traffic conformance has to be defined and is specified in this Recommendation at standardized interfaces (UNI, INI). With such conformance definitions, a network provider can make QoS commitments to part of the submitted traffic depending on the conformance of this traffic (see ITU-T Rec. I.356). There may be more than one QoS class for a given ATC (see ITU-T Rec. I.356).

Given that a user can commit to submitting cells in conformance to traffic descriptors in addition to the PCR, the reason a user would choose a service based on an ATM transfer capability other than the deterministic bit rate transfer capability is the potential of incurring a lower cost from the network provider. The specifics of such cost savings is outside of the scope of this Recommendation.

It is mandatory that the ATM transfer capability used on a given ATM connection, among these that are made available by the network, be implicitly or explicitly declared at connection establishment.

Once an ATM connection is established, the agreed upon ATC is the same at all standardized interfaces along the connection (see 5.3.1). However, it is a network operator's choice on how to support a given ATM transfer capability, provided it is complying with specifications at standardized interfaces.

There is not a one-to-one correspondence between services or service classes (e.g., broadband bearer service categories) and ATM transfer capabilities that may be used. For example, an upper layer data service such as FMBS may make use of a DBR, an SBR, an ABR or an ABT transfer capability. As a result, the requested ATM transfer capability as supported by signalling should not be checked by CAC against any other information than the information contained in the traffic contract (i.e., QoS Class, source traffic descriptor and associated CDV tolerances). DBR is the default ATM transfer capability.

A given ATM communication uses the same ATM transfer capability for both directions. Use of different transfer capabilities for the two connections of a communication raises issues relating for example to OAM and resource management cells or to routing and is currently not specified in this Recommendation. This also applies to multicast connections.

6.1.2 ATCs and multiplexing

Multiplexing and cell scheduling policies within a network element (see ITU-T Rec. E.736) are essential to realize QoS commitments, in particular to support more combinations of ATC and QoS class in a network, and in addition to optimize the utilization of network resources. These policies should take into account the ATCs and QoS classes to be supported. Once these criteria are met, the operation of these policies is operator specific. The implementation is outside the scope of this Recommendation.

Multiplexing of VCCs into a VPC raises the issue of maintaining the QoS of each multiplexed VCC. VCCs with the same ATC may be multiplexed into a DBR VPC. A DBR VCC or VPC with QoS class 1 may be used to emulate a VCC or VPC with any different ATC. With the exception of these two cases, the following is not currently addressed in this Recommendation:

– multiplexing VCCs with different ATCs and/or QoS classes into a single VPC (e.g., ABR VCCs and ABT VCCs within an SBR VPC);
– multiplexing VCCs in a VPC with the same or a different ATC (e.g., an ABR VPC carrying ABR VCCs);
– the emulation of one ATC by an alternative ATC (e.g., the use of SBR to transport an ABR service).

6.2 High level description of ATM transfer capabilities

An ATM Transfer Capability (ATC) specifies a set of ATM layer parameters and procedures that is intended to support an ATM layer service model and a range of associated QoS classes. Each individual ATC is further specified in terms of a service model, a traffic descriptor, specific procedures if relevant, a conformance definition and associated QoS commitments. Open-loop controlled ATCs (DBR and SBR) and closed-loop controlled ATCs (ABT and ABR) are specified as follows.

6.2.1 Deterministic transfer capability – DBR

The DBR transfer capability is intended to be used to meet the requirements of CBR traffic and therefore to provide for QoS commitments in terms of cell loss ratio, cell transfer delay and cell delay variation suitable for such traffic. However, DBR is not restricted to CBR applications and may be used in combination with looser QoS requirements, including unspecified requirements as indicated in ITU-T Rec. I.356.

DBR is solely based on the peak cell rate PCR(0+1) for the aggregate CLP = 0 and CLP = 1 cell flow, user-generated OAM cells being either aggregated or separately handled. Conformance definition for DBR is specified by one or two applications of the GCRA, depending on how user OAM cells are handled. Neither selective cell discard (see 7.2.4) nor cell tagging (see 7.2.3.6) apply to DBR.

For a complete specification of the DBR ATC, refer to 6.4.

6.2.2 Statistical bit rate transfer capability – SBR

The SBR transfer capability uses the sustainable cell rate and intrinsic burst tolerance in addition to the peak cell rate and is suitable for applications where there exists a prior knowledge of traffic characteristics beyond the peak cell rate, from which the network may obtain a statistical gain. QoS commitments are in terms of cell loss ratio. There may or may not be QoS commitments on delay.

There are three variants of SBR, depending on which parameter set is used in addition to the PCR(0+1). In the three cases, conformance to PCR(0+1) is specified by a GCRA(T_{PCR}, τ_{PCR}). SBR type 1 handles cells irrespective of the CLP bit value. SBR types 2 or 3 can be used for applications that can distinguish more loss-sensitive information (CLP = 0 cells) from less loss-sensitive information (CLP = 1 cells).

SBR type 1 uses SCR(0+1) and τ_{IBT}(0+1). Conformance to SCR(0+1) and τ_{IBT}(0+1) is specified by a GCRA(T_{SCR}, τ_{SCR}). QoS commitments are on CLP = 0 + 1 cells, both for cell loss ratio and optionally delay. Neither selective cell discard (see 7.2.4) nor cell tagging (see 7.2.3.6) apply to SBR type 1.

SBR type 2 and 3 use SCR(0) and τ_{IBT}(0). Conformance to SCR(0) and τ_{IBT}(0) is specified by a GCRA(T_{SCR}, τ_{SCR}). QoS commitments in terms of cell loss ratio are on CLP = 0 cells. Cell loss ratio for CLP = 0 + 1 cells is unspecified. There may be a QoS commitment concerning delay, and if there is, it applies to the CLP = 0 + 1 cell flow. Selective cell discard (see 7.2.4) applies to both SBR types 2 and 3. Cell tagging (see 7.2.3.6) only applies to SBR type 3.

For a complete specification of the SBR ATCs, refer to 6.5.

6.2.3 ATM block transfer – ABT

The ABT transfer capability is intended for applications that may adapt their instantaneous peak cell rate on a per-block basis. An ATM block is a group of cells delimited by RM cells. ABT uses static parameters declared at connection set-up and dynamic parameters renegotiable on an ATM block basis via resource management procedures using RM cells.

Static parameters are PCR(0+1), SCR(0+1) and associated tolerances. Dynamic parameters are peak cell rate for an ATM block: block cell rate BCR(0+1), and associated tolerance. PCR(0+1) specifies the maximum BCR(0+1) that may be negotiated via RM procedures for the connection. User-generated OAM cells may be aggregated or separately handled. SCR(0+1) specifies a longer term average behaviour of the connection; it is optional and may be set to 0.

There are two variants of ABT. In ABT/DT (delayed transmission), the source can start transmitting an ATM block only after having received a positive acknowledgement from the network by means of an RM cell. In ABT/IT (immediate transmission), the source starts transmitting user data cells immediately after the request RM cell; the ATM block is transferred as a whole if resources requested for that ATM block are available in the network; otherwise, it is discarded. In both cases, the BCR request may be elastic, in which case the network may choose to select a BCR smaller than the one requested by the source.

In ABT/DT, QoS commitments at the cell level are in terms of cell loss ratio, cell transfer delay and cell delay variation within an ATM block. Conformance definition at the cell level is specified within a block by one or two applications of the dynamic generic cell rate algorithm DGCRA, the variables of which are updated according to information conveyed by RM cells. If an SCR is specified, QoS commitments at the ATM block level are in terms of maximum delay for a BCR request to succeed.

In ABT/IT, QoS commitments at the cell level are in terms of cell loss ratio within an ATM block, assuming the BCR request is accepted along the entire connection. QoS commitments on delays within an ATM block only pertain when the elastic mode is not used. As for ABT/DT, conformance definition at the cell level is specified within a block by one or two applications of the DGCRA. If an SCR is specified, QoS commitments at the ATM block level are in terms of block loss ratio. In this respect, ABT/IT implements frame discard.

Selective cell discard (see 7.2.4) and cell tagging (see 7.2.3.6) do not apply to ABT.

Clause 6.6 describes the ABT/DT and ABT/IT service models, specifies the ABT RM cell format and the types of messages exchanged at standardized interfaces. ABT conformance is defined in 6.6.1.4 and 6.6.2.4.

6.2.4 Available Bit Rate – ABR

The ABR transfer capability is intended to support elastic applications that may adapt to the instantaneous bandwidth available within the network and do not have strict delay requirements. In such a case, the network may share the available resources between connections supporting such applications. ABR uses static parameters declared at connection set-up and dynamic parameters renegotiable via resource management procedures based on RM cells.

Static parameters are peak cell rate PCR(0+1), minimum cell rate MCR(0) and initial allowed cell rate IACR(0). User data cells and user OAM cells have the CLP bit set to 0. Dynamic parameters conveyed by RM cells are Explicit Cell Rate (ECR), Congestion Indication (CI), No-increase Indication (NI) and queue length. The value of the Allowed Cell Rate ACR(0) to the source is derived from these parameters and ranges between those of the MCR and the PCR.

In ABR, the user regularly polls the network for the currently available bandwidth by sending RM cells conveying a requested rate to the network. There are two modes of operation: explicit rate mode and binary mode. In the explicit rate mode, the network regularly returns to the source the ECR, from which the source derives its ACR. In the binary mode, the network may also return binary indicators; the source should use the binary indicators to compute its ACR.

Clause 6.7 describes the ABR service model, specifies the ABR RM cell format and the types of messages exchanged at standardized interfaces.

Clause 6.7.5 specifies the conformance definition for ABR in the explicit rate mode only. Reference source and destination behaviours to network indications are provided for both the explicit rate mode and the binary mode in Appendix VII.

In ABR, QoS commitments are in terms of cell loss ratio for CLP = 0 cells. In the binary mode, no QoS commitments can be made, but QoS indications in terms of cell loss ratio may be provided to connections respecting source and destination reference behaviours.

6.2.5 Guaranteed Frame Rate – GFR

The GFR transfer capability provides a minimum cell rate (MCR) for loss tolerant, non-real time applications with the expectation of transmitting data in excess of the MCR. It is assumed that the user generated data cells are organized in the form of frames that are delineated at the ATM layer. The network does not provide feedback to the user concerning instantaneous available network resources.

Traffic parameters are PCR(0+1), MCR(0), a maximum burst size MBS(0), a maximum frame size MFS(0+1) and tolerances associated with PCR(0+1) and MCR(0). A GFR cell is conforming if it conforms to the PCR(0+1), if it conforms to the maximum frame size and if it conforms to the homogeneous setting of the CLP bit of cells within the same frame. A GFR frame is conforming if all its cells are conforming and if the frame conforms to the Frame based GCRA, F-GCRA(T,τ), with parameters $T = 1/\text{MCR}$ and $\tau = \tau_{IBT} + \tau_{MCR}$, where $\tau_{IBT} = (\text{MBS} - 1) \times (1/\text{MCR} - 1/\text{PCR})$. By sending a frame with all CLP = 1 cells, the user indicates to the network that such a frame is of lesser importance than a frame with all CLP = 0 cells on the same GFR connection.

The GFR ATC allows the commitment to transmit the number of cells in conforming frames with the QoS corresponding to the associated QoS class. QoS commitments can only apply to cells in CLP = 0 frames of which all cells are conforming. Furthermore, with GFR the network attempts to deliver complete frames in excess of the minimum cell rate commitment, if all cells in the frame are conforming and provided that sufficient resources are available.

There are two variants of GFR: GFR1 and GFR2. In GFR1, tagging is not applicable. In GFR2, the network may apply frame tagging to non-conforming frames. Frame tagging means that the CLP bit of each cell of the frame is changed to 1. Frames tagged by the network and frames marked as CLP = 1 by the user are treated identically by the network.

For a complete specification of the GFR ATC, refer to 6.8.

6.3 Applicability of ATM transfer capabilities to applications

This clause provides guidelines on the possible use of individual ATM transfer capabilities to the transport of data belonging to a number of example broadband applications. The purpose of this clause is not to create a restrictive correspondence between ATM transfer capabilities and applications; rather, it is intended to provide both examples of how ATM transfer capabilities can be used and to create design targets for these transfer capabilities.

It is expected that ultimately the choice of which ATM transfer capabilities is used to transport application data is influenced by a number of factors, among which:

– Availability of transfer capabilities: not all networks may wish to provide all transfer capabilities described in this Recommendation.

– Actual attainable QoS for each ATM transfer capability: this depends, among other things, on the resource management policy adopted and traffic engineering.

– Capability of the application to cope with degradation of the ATM layer transfer characteristics: for some applications a reduction of the available bandwidth would result in application failure (e.g., circuit emulation); for others this can be translated into acceptable reduction in the quality of one or more components (e.g., coarser image definition or reduced animation rate for video applications).

– Tariffing: the user's choice is determined in part by the tariffs applied to each combination of an ATM transfer capability and a QoS class.

Two examples illustrate the above points:

– The DBR capability in combination with an appropriate QoS class can of course be used, with the appropriate PCR value, to transport data for all broadband applications. The user's choice on whether to use DBR or one of the other ATM transfer capabilities will depend on a number of factors, among them being network tariffing.

– The ABR capability can also be used to transport data belonging to all applications, provided that network resource management offers unused bandwidth so that in effect no limitations will be imposed on the application.

The two cases above should be seen as extreme cases, illustrating the interaction between tariffing, resource management and network traffic engineering. Typically, there will be appropriate mappings between applications and ATM transfer capabilities.

6.4 Deterministic bit rate transfer capability (DBR)

6.4.1 Definition and service model

The deterministic bit rate transfer capability can be used by connections that characterize the traffic with a single parameter, the peak cell rate value. The source may emit cells at or below the negotiated peak cell rate and may also even be silent for periods of time.

The basic commitment made by the network to a user who reserves resources via the DBR capability is that, once the connection is established, the negotiated ATM layer QoS is assured to all cells when all cells are conforming to the relevant conformance tests.

In the DBR capability, the source can emit cells at the peak cell rate at any time and for any duration and the QoS commitments still pertain.

The DBR capability is intended to allow networks to make QoS commitments that enable the network to support CBR applications but is not restricted to these applications; see 6.4.6 on "QoS aspects".

The DBR capability may be used for both VPCs and VCCs. Neither cell tagging (see 7.2.3.6) nor selective cell discard (see 7.2.4) apply to the DBR capability. VC RM cells on a VCC and VP RM cells on a VPC are not used to operate DBR; however, if such cells are present on the connection, they are considered as part of the aggregate CLP = 0 + 1 user data cell flow.

It is recommended that each network element be capable of supporting DBR VPCs with QoS class 1, (see ITU-T Rec. I.356).

NOTE 1 – How to comply with the specified CDV tolerance at the VP level is an implementational or operational option, e.g., by controlling the number of VCCs multiplexed and the load or by shaping the collection of VCCs into the VPC.

NOTE 2 – The above sentence does not imply either that each network operator must offer a DBR VPC service.

Use of DBR with QoS class U

If the unspecified QoS class is selected by the user of the DBR connection, there are no QoS commitments to the connection. In this case, it is a network option how to resource the connection, how to apply intelligent cell scheduling and buffer allocation schemes and how cells are discarded in case of buffer overflow. For example, it may be preferable to discard consecutive cells of a single connection instead of spreading cell discard on a number of connections. This may be done by a traffic control function such as frame discard.

6.4.2 Options regarding source traffic descriptor and conformance definition

DBR allows the use of different source traffic descriptors with either a single peak cell rate or two peak cell rates, one for user data cells and one for user OAM cells. The different options for the DBR source traffic descriptor are given in 6.4.3.

DBR allows the operator to apply different conformance definitions. Clause 6.4.5 describes the aggregate and the separate conformance definition. If the source traffic descriptor with a single peak cell rate is used, the aggregate conformance definition applies. If the source traffic descriptor with separate peak cell rates is used, the operator applies either the separate or the aggregate conformance definition, taking into account the specific conditions stated in 6.4.5 regarding shaping.

Depending on the conformance definition used, one or two CDV tolerances are needed; depending on the source traffic descriptor used, the CDV tolerance applied may need to be calculated. Clause 6.4.4 provides the details related to CDV tolerances.

The conformance definition applied has some impact on the QoS commitments for a DBR connection. This impact is described in 6.4.6 on QoS commitments. The conformance definition applied also has some impact on the UPC/NPC function. This impact is described in 6.4.5 on UPC/NPC mechanisms.

6.4.3 Source traffic descriptor

DBR uses peak cell rate as specified in 5.4.1.

At connection establishment or at subscription time, the user and the network agree upon one of the three subsequent source traffic descriptors:

i) Two peak cell rates, namely, PCR for user data cells with the associated peak emission interval T_{PCR}(data) and PCR for end-to-end user OAM cells with the associated peak emission interval T_{PCR}(OAM).

ii) One peak cell rate, namely, PCR for user generated cells with the associated peak emission interval T_{PCR}(agg).

The peak cell rate values shall not be renegotiated at standardized interfaces by means of RM procedures during the lifetime of the connection, but could be renegotiated via signalling or via network management procedures.

Use of source traffic descriptor i)

If the user intends to transmit user OAM cells and does not allow aggregate shaping of user data and user OAM, source traffic descriptor i) shall apply.

When source traffic descriptor i) is used, the PCR of the user OAM cells can be specified by declaring the nominal period n of the forward performance monitoring cell stream (see ITU-T Rec. I.610); in this case, the following default rules apply for computing default T_{PCR}(OAM), T_{PCR}(agg):

– when the operator uses separate conformance definition (see 6.4.5): $T_{PCR}(\text{OAM}) = n \cdot T_{PCR}(\text{data})$;

– when the operator uses aggregate conformance definition (see 6.4.5): $T_{PCR}(agg) = \frac{n}{n+1} T_{PCR}(data)$.

These default rules assume that OAM user generated cell flows only consist of forward performance monitoring. The case when other user OAM cells are generated by the user is for further study. The rules for computing the traffic parameters from the relevant information conveyed by the present signalling protocol (see ITU-T Rec. Q.2931) are given in Appendix II.

Use of source traffic descriptor ii)

The use of source traffic descriptor ii) by the user does not imply that user OAM cells are precluded. If any user OAM cells are present, they are aggregated with the user data cells in the descriptor T_{PCR}(agg). T_{PCR}(agg) is computed as: $T_{PCR}(agg) = \frac{1}{\frac{1}{T_{PCR}(data)} + \frac{1}{T_{PCR}(OAM)}}$.

When source traffic descriptor ii) is used, when the user wants to send user OAM cells and knows the values of T_{PCR}(data) and T_{PCR}(OAM), then the user could derive the value for T_{PCR}(agg) by the same formula as given above.

6.4.4 CDV tolerances

DBR uses CDV tolerance as specified in 5.4.1.

For each negotiated PCR and at each interface where cell conformance applies, there corresponds a CDV tolerance value that accounts for cell delay variation that may affect the respective cell flows of the connection (see 5.4.1).

The renegotiation of the PCR (see 6.4.3) may imply also the modification of the associated CDV tolerance.

CDV tolerances for source traffic descriptor i) and separate conformance

When source traffic descriptor i) is used and separate conformance definition is applied, the following values of CDV tolerance are required:

– the CDV tolerance corresponding to the PCR of user data CLP = 0 + 1 cells τ_{PCR}(data);

– the CDV tolerance corresponding to the PCR of end-to-end user OAM cells τ_{PCR} (OAM).

In this case, τ_{PCR}(data) can either be conveyed via signalling or assigned on a per subscription basis (see 5.4.1.3). On the contrary, τ_{PCR}(OAM) is always computed on the basis of the PCR of the end-to-end user OAM cells as $\tau_{PCR}(\text{OAM}) = T_{PCR}(\text{OAM})$.

The rule for computing the CDV tolerance τ_{PCR}(OAM) from the relevant information conveyed by the present signalling protocol (see ITU-T Rec. Q.2931) is given in Appendix II.

CDV tolerance for source traffic descriptor i) and aggregate conformance

When source traffic descriptor i) is used and aggregate conformance definition is applied, only the value of CDV tolerance corresponding to the PCR of user generated CLP = 0 + 1 cells, τ_{PCR}(agg), is required. In this case, τ_{PCR}(agg) is computed as $\tau_{PCR}(\text{agg}) = T_{PCR}(\text{agg}) + \tau_{PCR}(\text{data})$, where τ_{PCR}(data) is either conveyed via signalling or assigned on a per subscription basis and T_{PCR}(agg) is computed as given in 6.4.3.

The rule for computing the CDV tolerance τ_{PCR}(agg) from the relevant information conveyed by the present signalling protocol (see ITU-T Rec. Q.2931) is given in Appendix II.

CDV tolerance for source traffic descriptor ii)

When source traffic descriptor ii) is used, only the value of CDV tolerance corresponding to the PCR of user generated CLP = 0 + 1 cells, τ_{PCR}(agg), is required. In this case, τ_{PCR}(agg) can either be conveyed via signalling or assigned on a per subscription basis. Note that when source traffic descriptor ii) is used, when the user wants to send user OAM cells and knows the values of T_{PCR}(data), τ_{PCR}(data), and T_{PCR}(OAM), then the user could derive the value for τ_{PCR}(agg) by the same formula as given for source traffic descriptor i) and aggregate conformance.

The rule for computing the CDV tolerance τ_{PCR}(agg) from the relevant information conveyed by the present signalling protocol (see ITU-T Rec. Q.2931) is given in Appendix II.

6.4.5 Conformance definition

The conformance definition to the peak cell rate ($\Lambda_{PCR} = 1/T_{PCR}$) of a cell stream requires that the CDV tolerance τ_{PCR} allocated to the upstream portion of the connection be specified (see 5.3.1, 5.4.1). These two parameters complete the specification of the generic cell rate algorithm given in Annex A.

The number of non-conforming cells of a cell stream is defined as the number of cells either of the two equivalent versions of the GCRA would determine to be non-conforming.

Conformance of a cell flow according to this definition can be estimated by the 1-point measurement process in ITU-T Rec. I.356.

Conformance is always specified regardless of the CLP bit value.

If the user intends to transmit user OAM cells, and if the user does not allow aggregate shaping of user data and user OAM and if the network performs PCR shaping actions (see 7.2.7), then the conformance definition is separate for the user OAM component and the user data component. Otherwise, the conformance definition is aggregate.

Separate conformance

If conformance is separate for user data and user OAM cells, the following applies:

– cell conformance of the aggregate CLP = 0 + 1 user data cell flow is checked with respect to the PCR descriptor negotiated for user data CLP = 0 + 1 traffic;

– cell conformance of the user OAM cell flow is checked with respect to the PCR descriptor negotiated for user OAM traffic.

The separate conformance definition for the DBR capability is illustrated in Figure 6-a.

Aggregate conformance

If conformance is aggregate, cell conformance of the aggregate user generated CLP = 0 + 1 cell flow is checked with respect to the PCR descriptor;

– in case of source traffic descriptor ii): to the PCR descriptor negotiated for user generated CLP = 0 + 1 traffic;

– in case of source traffic descriptor i): to the PCR descriptor for the user generated CLP = 0 + 1 traffic that is obtained from the PCR descriptors negotiated for user data CLP = 0 + 1 traffic and for user OAM traffic (see 6.4.3).

The aggregate conformance definition for the DBR capability is illustrated in Figure 6-b.

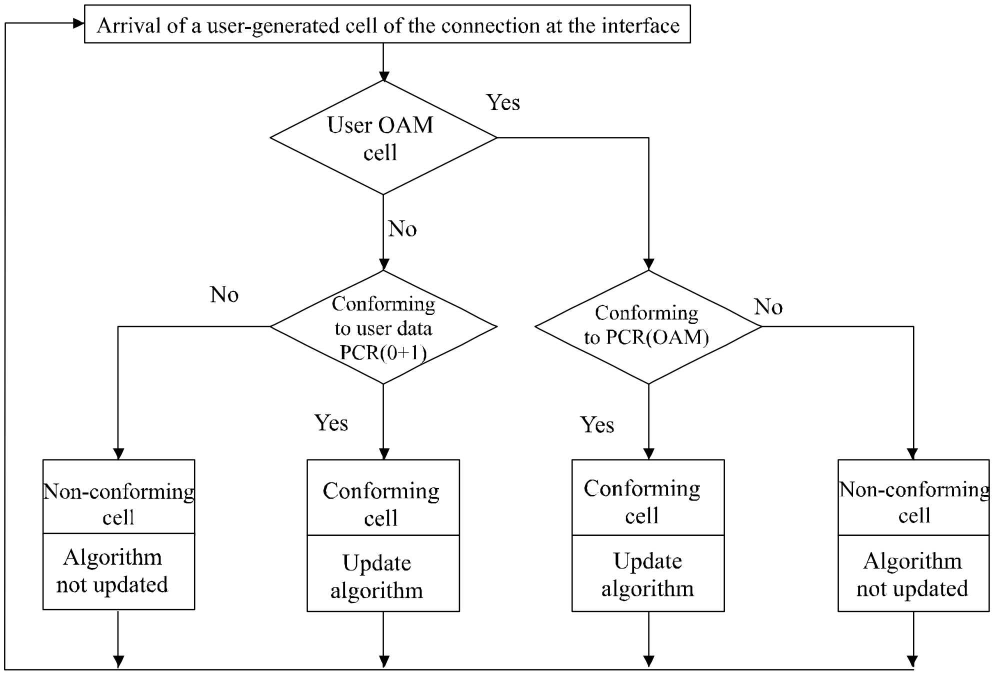

a) Separate conformance definition

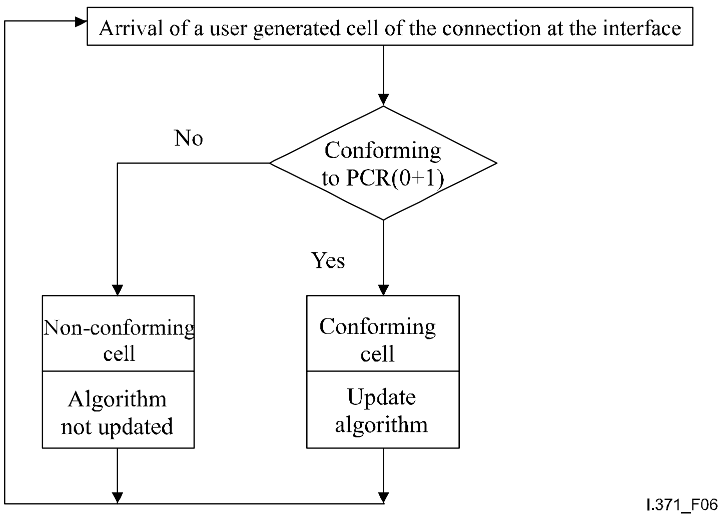

b) Aggregate conformance definition

Figure 6/I.371 – Conformance definition for DBR

UPC/NPC mechanisms

During the connection lifetime, conformance to peak cell rate traffic descriptors may be continuously checked within the network by static UPC/NPC mechanisms, given such UPC/NPC mechanisms are present (see 7.2.3). Where present, UPC/NPC mechanisms applying to user data cells are performed on the user data CLP = 0 + 1 flow, regardless of the CLP bit value.

The conformance definition above does not imply any particular implementation of the UPC/NPC. Furthermore, even if the conformance definition is separate, UPC/NPC can be done on the aggregate cell flow. However, the UPC/NPC should satisfy certain requirements as discussed in 7.2.3.2.

6.4.6 QoS aspects

The QoS class is negotiated at call establishment or on a subscription basis for the connection. The negotiated ATM layer QoS is assured to all cells when all cells are conforming to the relevant conformance tests. The default QoS class for the DBR capability shall be QoS class 1 (see ITU-T Rec. I.356), which includes a specified CLR commitment for the aggregate CLP = 0 + 1 cell flow, regardless of the CLP bit value, and an end-to-end CDV commitment suitable for circuit emulation applications. The support of additional QoS classes with less stringent commitments is a network operator's choice. In particular, QoS class U (see ITU-T Rec. I.356) may be used in combination with DBR. Then, there are no QoS commitments to the connection and it is a network option on how to allocate resources to the connection. However, the network may still achieve some QoS indications via suitable traffic engineering rules. QoS commitments are always specified regardless of the CLP bit value.

If some cells are non-conforming to some of the relevant conformance tests, the network may consider the connection as non-compliant (see 5.3.2). If the network chooses to offer QoS commitments to a connection with some non-conforming cells, the ATM layer QoS is only assured to a volume of cells that is conforming to all relevant conformance tests. This is obtained if the network operator chooses to assure the ATM layer QoS to all cells admitted by the UPC/NPC where the parameters of the UPC/NPC have been set with sufficient margin such that the UPC/NPC does not falsely discard cells.

The following principles apply (see ITU-T Rec. I.356).

QoS commitments for separate user data and user OAM conformance definition

– The ATM layer QoS commitment for user data cells applies to the volume of user data cells equivalent to the volume of conforming user data cells.

– Cell conformance of the user OAM cell stream does not influence the above computation. The ATM layer QoS commitment for the user OAM cell stream, if any, applies to the volume of user OAM cells equivalent to the volume of conforming user OAM cells.

– RM cells, if any, are treated as user data cells.

– When there are non-conforming cells, the method for determining how many cells are assured of the ATM layer QoS is network specific (see ITU-T Rec. I.356).

QoS commitments for aggregate conformance definition

– The ATM layer QoS commitment for user data and user OAM cells applies to the volume of such cells equivalent to the volume of conforming cells.

– User OAM cells and RM cells, if any, are treated as user data cells.

– When there are non-conforming cells, the method for determining how many cells are assured of the ATM layer QoS is network specific (see ITU-T Rec. I.356).

6.5 Statistical bit rate transfer capability (SBR)

6.5.1 Definition and service model

In the Statistical Bit Rate (SBR) transfer capability the end-system uses standardized traffic parameters (SCR/IBT) to describe, in greater detail than just the peak cell rate, the cell flow that will be emitted on the connection.

The SBR capability is suitable for applications where there exists prior knowledge of some traffic characteristics of the application.

The delay performance of the SBR capability may be specified by negotiating a suitable QoS class (see ITU-T Rec. I.356).

The SBR capability may be used for both VPCs and VCCs. Depending on the configuration, the SBR capability allows cell tagging (see 7.2.3.6). Also depending on the configuration, selective cell discard (see 7.2.4) applies. VC RM cells on a VCC and VP RM cells on a VPC are not used to operate SBR; however, such cells that would still be present on the connection are considered as part of the user data cell flows.

6.5.2 Source traffic descriptor and CDV tolerances

SBR uses peak cell rate and associated CDV tolerance as specified in 5.4.1, sustainable cell rate, intrinsic burst tolerance and associated CDV tolerance as specified in 5.4.2.

At connection set-up or at subscription time, the user and the network agree upon the subsequent source traffic descriptor:

PCR and SCR/IBT.

In the SBR capability, the traffic parameters (PCR and SCR/IBT) characterize the negotiated user generated cell flow.

The following configurations of the PCR and SCR/IBT traffic parameters with the CLP bit and cell tagging (see 7.2.3.6) are specified:

1) PCR traffic parameter for the user generated CLP = 0 + 1 cell flow and SCR/IBT traffic parameter for the user generated CLP = 0 + 1 cell flow. Cell tagging (see 7.2.3.6) is not applied.
2) PCR traffic parameter for the user generated CLP = 0 + 1 cell flow and SCR/IBT traffic parameter for the CLP = 0 cell flow; cell tagging (see 7.2.3.6) is not applied.
3) PCR traffic parameter for the CLP = 0 + 1 cell flow and SCR/IBT traffic parameter for the CLP = 0 cell flow; cell tagging (see 7.2.3.6) is applied.

Selective cell discard (see 7.2.4) does not apply to configuration 1. Selective cell discard may apply to configurations 2) and 3).

When configuration 1) is used, the following values of CDV tolerance are required:

– the CDV tolerance corresponding to the PCR of CLP = 0 + 1 user generated cells $\tau_{PCR}(0 + 1)$;
– the CDV tolerance corresponding to the SCR of CLP = 0 + 1 user generated cells $\tau'_{SCR}(0 + 1)$.

When configuration 2) or 3) is used, the following values of CDV tolerance are required:

– the CDV tolerance corresponding to the PCR of CLP = 0 + 1 user generated cells $\tau_{PCR}(0 + 1)$;
– the CDV tolerance corresponding to the SCR of CLP = 0 user generated cells $\tau'_{SCR}(0)$.

In these two cases, both τ_{PCR} and τ'_{SCR} could either be conveyed via signalling or assigned on a per subscription basis.

The present signalling protocol conveys the parameter MBS instead of the parameter IBT. The rule for computing IBT from MBS is given in 5.4.2.2.

The source traffic descriptor shall not be renegotiated at standardized interfaces by means of RM procedures during the lifetime of the connection, but could be renegotiated via signalling or via network management procedures. The renegotiation of the source traffic descriptor may imply also the modification of the associated CDV tolerances.

6.5.3 Conformance definition and QoS commitments

Conformance definition

The conformance definition at an interface is based on two instances of the generic cell rate algorithm (GCRA) (see Annex A) where the two instances operate in a coordinated mode. Coordinated mode means the states of the GCRA are updated if and only if an arriving cell is conforming to all relevant instances of the GCRA (see Annex B for further details). Conformance definitions for the above three configurations are illustrated in Figures 7, 8 and 9 and reference algorithms are detailed in Annex B. In these configurations:

– For traffic-parameter configuration 1, Figure 7, a user data or user OAM cell is conforming when it is conforming to both the PCR(0+1) and SCR(0+1) conformance tests.

– For traffic-parameter configurations 2 and 3, Figures 8 and 9, a CLP = 0 user data or user OAM cell is conforming when it is conforming to both the PCR(0+1) and the SCR(0) conformance tests.

– For traffic-parameter configurations 2 and 3, Figures 8 and 9, a CLP = 1 cell is conforming if it is conforming to the PCR(0+1) conformance test. In this case, the state of the SCR(0) conformance test is left unchanged, and the state of the PCR(0+1) test is updated (this is independent of the cell tagging (see 7.2.3.6)).

– For traffic-parameter configuration 3, Figure 9, a CLP = 0 user data or user OAM cell that is conforming to the PCR(0+1) conformance test but not conforming to the SCR(0) conformance test is virtually tagged, i.e., is considered by the conformance definition as a conforming CLP = 1 cell. In this case, the state of the SCR(0) conformance test is left unchanged, and the state of the PCR(0+1) test is updated.

This conformance definition does not imply any particular implementation of the UPC/NPC.

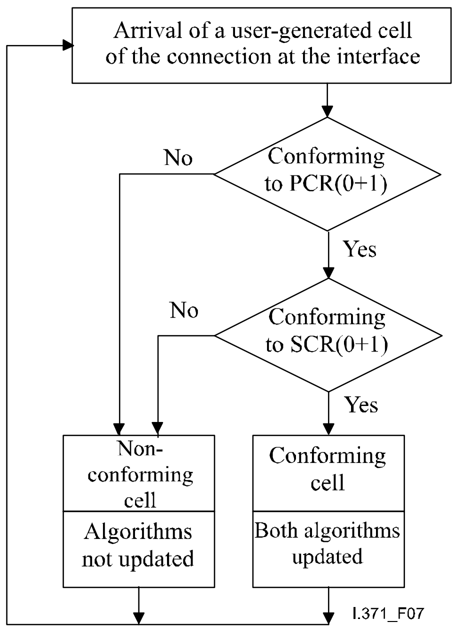

Figure 7/I.371 – Conformance definition for SBR configuration 1

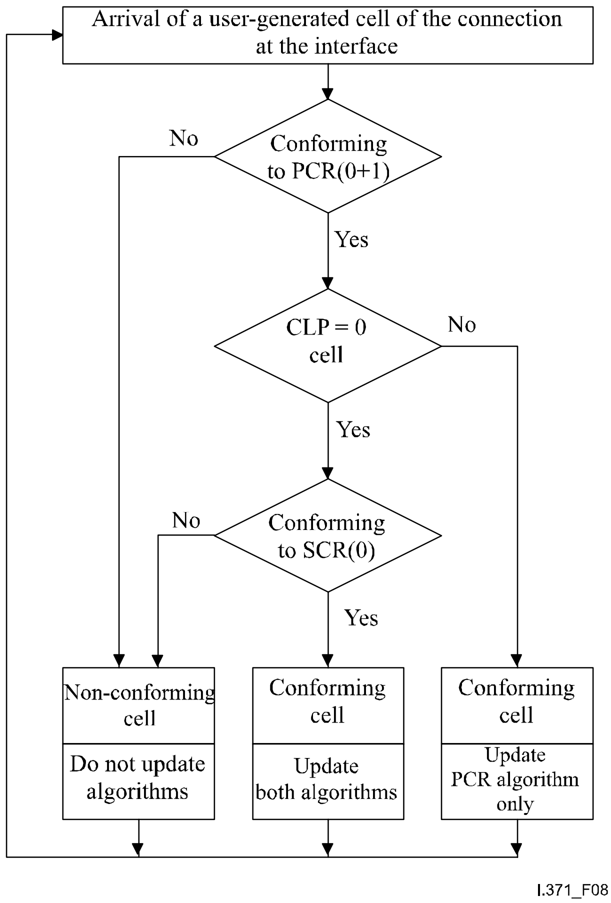

Figure 8/I.371 – Conformance definition for SBR configuration 2

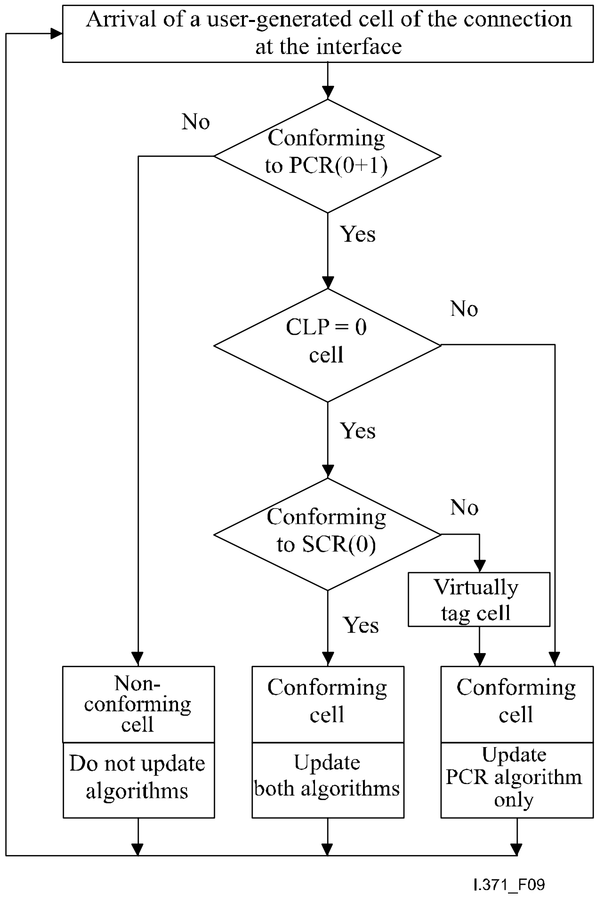

NOTE – The term "virtually tagged" is used because a conformance definition, since it is not a physical device, cannot alter a bit in an ATM cell. However, a UPC/NPC that implements the tagging option would actually change the CLP bit from 0 to 1.

Figure 9/I.371 – Conformance definition for SBR configuration 3 (Note)

QoS aspects

For traffic-parameter configuration 1 above, any QoS commitments apply to the aggregate CLP = 0 + 1 cell flow. Thus, the QoS commitments are independent of the CLP bit. A QoS class with specified CLR objectives for the aggregate cell flow (see ITU-T Rec. I.356) or a QoS class with both specified CLR objectives and delay objectives for the CLP = 0 + 1 cell flow (see ITU-T Rec. I.356) may be negotiated.

For traffic-parameter configurations 2 and 3 above, a QoS class with specified CLR objectives for the CLP = 0 cell flow may be negotiated (see ITU-T Rec. I.356). The QoS commitment on the CLR objective of the aggregate CLP = 0 + 1 cell flow is unspecified. Hence, the QoS commitment on the CLR objective of the CLP = 1 cell flow is also unspecified. Alternatively, a QoS class with specified CLR objectives for the CLP = 0 cell flow and specified delay objectives for the CLP = 0 + 1 cell flow (see ITU-T Rec. I.356) may be negotiated.

If some cells are non-conforming to some of the relevant conformance tests, the network may consider the connection to be non-compliant (see 5.3.2). If the network chooses to offer QoS commitments to a connection with some non-conforming cells, the ATM layer QoS is only assured to a volume of cells that is conforming to all relevant conformance tests. This is obtained if the network operator chooses to assure the ATM layer QoS to all cells admitted by the UPC/NPC where the parameters of the UPC/NPC have been set with sufficient margin such that the UPC/NPC does not falsely discard cells.

The following principles apply:

– For configuration 1, the ATM layer QoS commitment for user-generated cells applies to the volume of such cells equivalent to the volume of conforming cells.

– For configurations 2 and 3, when the user-generated cells are conforming to the PCR(0+1) traffic descriptor, the ATM layer QoS commitment on the CLP = 0 component applies to the volume of CLP = 0 cells that are conforming to the SCR(0) conformance test.

– For configurations 2 and 3, when some user-generated cells are not conforming to the CLP = 0 + 1 conformance test, the ATM layer QoS commitments to the CLP = 0 and CLP = 0 + 1 components are network specific.

6.6 ATM block transfer capabilities (ABT)

An ATM Block Transfer (ABT) capability is an ATM layer mechanism for providing a service where the ATM layer transfer characteristics are negotiated on an ATM block basis. Within an ATM block accepted by the network, the network allocates sufficient resources such that the QoS received by the ATM block is equivalent to QoS received by a DBR connection with the same peak cell rate as the negotiated peak cell rate of the ATM block, referred to as the Block Cell Rate (BCR), and with the same negotiated QoS class.

Specifically, an ATM block is defined as follows.

Definition (ATM block)

An ATM block is a group of cells of an ATM connection delineated by two Resource Management (RM) cells, one before the first cell of the ATM block (leading RM cell) and another after the last cell of the ATM block (trailing RM cell). The exact definition of the RM cells delineating an ATM block depends on the specific usage of RM cells, namely on the ABT capability. The trailing RM cell of an ATM block may be the leading RM cell of the subsequent ATM block (see also Annexes C and D). The BCR of an ATM block is constant over the ATM block duration.

ATM blocks are not necessarily related to upper layer protocols, e.g., at the CS-PDU level (Figure 10).

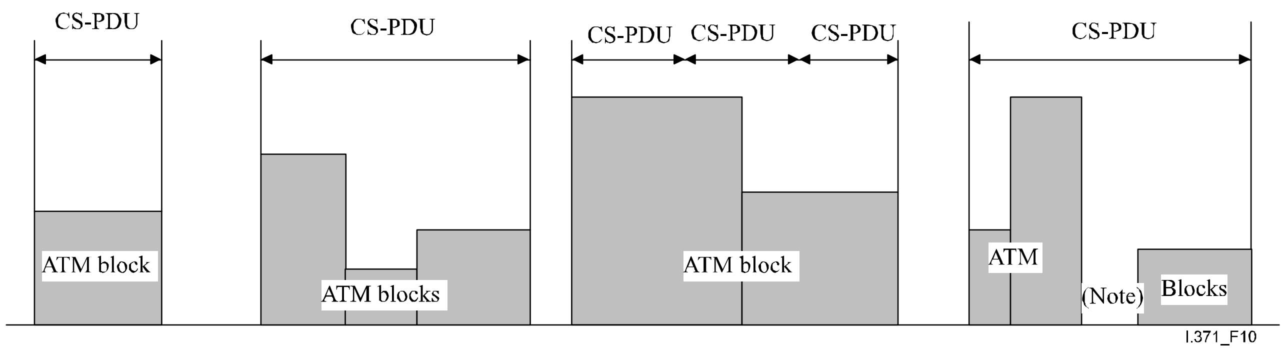

NOTE – ATM block with BCR = 0.

Figure 10/I.371 – Examples of relationship between ATM blocks and CS-PDUs

The ABT capability may be used for both VPCs and VCCs. Specifically, ABT may apply to any cell flow of a VCC or a VPC. If ABT applies to some VCCs within a VPC, the CLP = 0 + 1 cell flow of ABT VCCs share the capacity of the CLP = 0 + 1 cell flow of the VPC that is allocated to ABT. In this case, the VPC is allocated a static bandwidth. Dynamically changing the bandwidth of the VPC via ABT is currently not specified in this Recommendation.

ABT does not support cell tagging (see 7.2.3.6). Selective cell discard functions (see 7.2.4) are currently not used in ABT. For ABT RM cells that are used to delineate ATM blocks, cell sequence integrity of user and ABT RM cells is mandatory all along the connection.

At connection establishment, the connectivity between two users is established by means of two point-to-point, unidirectional connections, but a zero BCR is allocated to user cells. Only point-to-point communication is considered in the present specification of ABT. Use of ABT for point-to-multipoint communication is not specified in this Recommendation.

Also at connection establishment, the user negotiates the following parameters via signalling or by means of network management:

i) the maximum cell rate by specifying the peak cell rates and the CDV tolerances of all the relevant cell flows, namely CLP = 0 + 1 (user OAM included) and user OAM of the connection;

ii) the maximum frequency of BCR renegotiation transactions by specifying the peak cell rates and the CDV tolerances of the ABT RM cell flows in the forward and backward directions;

iii) a sustainable cell rate (SCR/IBT) for the CLP = 0 + 1 cell flow, where SCR may be set equal to 0.

The above parameters are static and it is not specified in this Recommendation that they will be renegotiated during the lifetime of a connection.

Two ABT traffic handling capabilities are defined, namely the ATM block transfer with delayed transmission (ABT/DT) and the ATM block transfer with immediate transmission (ABT/IT).

6.6.1 ABT with delayed transmission (ABT/DT)

6.6.1.1 Definition and service model

In ABT/DT, during the connection lifetime, the BCR of successive ATM blocks is dynamically negotiated with the network. BCR renegotiations may be initiated by either end user on both the forward and the backward direction. For a given direction, only one renegotiation initiated by a given user can be in progress within the network. A BCR modification is achieved by sending a request to the network by means of an ABT/DT RM cell.

Since a user may initiate BCR renegotiations in either direction, two BCR renegotiations initiated by both end users may collide within the network. Let the labels source and destination denote for a given ATM block the entities generating and receiving traffic respectively. In the case of colliding BCR renegotiations, the BCR renegotiation initiated by the destination in the backward direction has priority over the BCR renegotiation initiated by the source in the forward direction. Furthermore, BCR renegotiations initiated by the network have priority over BCR renegotiations initiated by end users (see 6.6.1.4).

A committed bandwidth for the user data and user OAM CLP = 0 + 1 cell flow is specified as follows:

Definition (Committed bandwidth)

A committed bandwidth for the user data and user OAM CLP = 0 + 1 cell flow in a given direction is defined when a sustainable cell rate greater than 0 is specified at connection establishment for this cell flow in the relevant direction and a QoS class with specified CLR objectives is negotiated. The committed bandwidth is related to the amount of reserved resources and its value is equal to that of the specified sustainable cell rate. The long-term average of resources, which can potentially be reserved, is at least as large as the SCR. Moreover, if the amount of traffic is conforming to the specified sustainable cell rate traffic descriptor (see 6.6.1.2), then a new BCR reservation should be accepted by the network within a finite time interval, consistent with QoS commitments (see 6.6.1.4).

In the case that a sustainable cell rate equal to 0 has been specified, the network may accept or may deny BCR renegotiations and does not provide for any commitments on the time to access to network resources (no commitment on the time for a BCR increase request to succeed). Note, however, that the network may achieve some delay objectives (QoS indications) via suitable traffic engineering rules. The cell level QoS commitments are still met by the network as long as the network has not renegotiated the allocated BCR.

6.6.1.2 Source traffic descriptor and CDV tolerances

At connection set-up or at subscription time, the user and the network agree upon a source traffic descriptor including the following traffic parameters:

– maximum cell rate PCR for the user-generated CLP = 0 + 1 cells (including user OAM cells but no RM cells);

– optionally, maximum cell rate PCR_{OAM} for user OAM cells;

– SCR/IBT traffic parameters for user-generated CLP = 0 + 1 cells (not including RM cells); the SCR may be set equal to 0;

– peak renegotiation rate, namely the peak cell rate PCR(RM) of the ABT/DT RM cell flows.

The above parameters are static and it is not specified in this Recommendation that they will be renegotiated during the lifetime of a connection.

In addition to the above source traffic descriptor, the following values of CDV tolerance are required:

– the CDV tolerance corresponding to the CLP = 0 + 1 cell flow PCR;

– the CDV tolerance corresponding to OAM cell flow PCR (when this parameter is included in the source traffic descriptor);

– the CDV tolerance corresponding to SCR/IBT (if SCR ≠ 0);

– the CDV tolerance corresponding to forward and backward RM cell flow PCR.

All the above values can either be conveyed via signalling or assigned on a per subscription basis.

6.6.1.3 Dynamically changing traffic parameters and RM cell format for ABT/DT

During the lifetime of the connection, the values of the following dynamic parameters are renegotiated between the user of the ABT capability and network elements along the connection via RM cells: Block Cell Rate (BCR) for the user data plus user OAM cell flow and user OAM BCR. The BCR may not exceed the PCR negotiated at connection establishment.

Table 2 contains the RM cell format for ABT transfer capabilities.

Table 2/I.371 – RM cell format for ABT

Field	Octet(s)	Bit(s)	Coding
ATM Header (Note 1)	1-5	All	As in ITU-T Rec. I.361
Protocol ID (Note 5)	6	All	2 (ABT/DT) 3 (ABT/IT)
Message type: Direction	7	8	(Note 2)
Message type: Traffic management cell	7	7	
Message type: Congestion indication	7	6	
Message type: Maintenance	7	5	
Message type: Req/Ack	7	4	

Table 2/I.371 – RM cell format for ABT

Field	Octet(s)	Bit(s)	Coding
Message type: Elastic/rigid	7	3	
Message type: Reserved	7	1-2	Clause 8.1
CLP = 0 + 1 BCR (user data + user OAM cells)	8-9	All	
User OAM BCR	10-11	All	
Reserved	12-13	All	Clause 8.1
Block size	14-17	All	(Note 3)
Sequence number	18-21	All	(Note 4)
Reserved	22-51	All	Clause 8.1
Reserved	52	3-8	Clause 8.1
CRC-10	52	1-2	
	53	All	

NOTE 1 – Only ABT RM cells with CLP bit set to 0 are currently specified in this Recommendation.

NOTE 2 – The direction bit equals 0 for forward RM cells and equals 1 for backward RM cells.

NOTE 3 – Integer valued. Least significant bit is bit 1 of octet 17.

NOTE 4 – Integer valued. Least significant bit is bit 1 of octet 21.

NOTE 5 – ABR and ABT connections use their own RM cells with a specified Protocol ID for traffic management purposes (PID = 1 for ABR, PID = 2 for ABT/DT, PID = 3 for ABT/IT). All other RM cells on the same level (VPC or VCC) with a different PID are non-conforming and may be discarded at the entrance of a network or where the first RM cell processing is done for this connection in the network.

Protocol identifier

ABT/DT RM cells are identified by protocol ID 2.

Message type

Message type is an octet containing six single bit fields plus two reserved bits. The message type field gives the semantic meaning of the ABT/DT RM cell.

Direction: This bit indicates the direction for which the ABT/DT RM cell applies. If DIR = 0, the ABT/DT RM cell applies to the forward direction. If DIR = 1, the ABT/DT RM cell applies to the backward direction.

Traffic management: The traffic management bit distinguishes a normal ABT/DT RM cell used by the user for BCR renegotiation purposes from an ABT/DT RM cell generated by the network for traffic control purposes (traffic management cell). The traffic management bit is set equal to 1 for a traffic management cell and is set equal to 0 otherwise.

Congestion Indication (CI): When the ABT/DT RM cell is involved in a BCR modification, this bit indicates whether the BCR renegotiation has succeeded or failed. If CI = 0, the BCR modification has succeeded; if CI = 1, the BCR modification has failed.

Maintenance: Two types of ABT/DT RM cells are defined in ABT/DT. ABT/DT RM cells used for BCR modifications by the user and possibly by the network are identified by maintenance = 0, and ABT/DT RM cells used for maintenance of ABT/DT procedures are identified by maintenance = 1. Maintenance procedures are for further study.

Req/Ack: This bit indicates whether the ABT/DT RM cell is a request or an acknowledgement message. Specifically, the exact meaning is as follows:

1) when sent by the user and Req/Ack = 0, the RM cell is a BCR modification request;
2) when sent by the user and Req/Ack = 1, the ABT/DT RM cell is an acknowledgement to a request or an acknowledgement of a BCR modification sent by the network;
3) when sent by the network and Req/Ack = 1, the ABT/DT RM cell is an acknowledgement RM cell of a BCR modification;
4) when sent by the network and Req/Ack = 0, the ABT/DT RM cell is the request of a BCR modification.

Elastic/rigid bit: This bit is set equal to 0 by the source to indicate that the network may optionally overwrite the cell rate fields, otherwise, this bit is set equal to 1.

CLP = 0 + 1 block cell rate

This field is used in ABT/DT RM cells for BCR modification, maintenance, etc. (in relation with the message type coding) concerning the CLP = 0 + 1 cell flow (user OAM cells included). In the case of BCR request message sent by the user, the value of this field is equal to the requested BCR. In the case of a BCR allocation message sent by the network, the value of this field is equal to the allocated BCR. In the rigid mode, the allocated BCR to an accepted ATM block will be equal to the requested BCR. In the elastic mode, the allocated BCR will be equal to or less than the requested BCR and should not be less than the SCR if the requested BCR was larger than the SCR.

The requested/allocated BCR is coded according to the 5-bit exponent and 9-bit mantissa coding scheme given in 5.4.1.2. In addition, a bit, referred to as *nz* bit, indicates whether the BCR is null or not. Specifically, a BCR Λ is coded as:

$$\Lambda = \left[2^m \cdot \left(1 + \frac{k}{512} \right) \right] \cdot nz$$

$$0 \le m \le 31 \text{ and } 0 \le k \le 511$$

$$nz \in \{0,1\}$$

The requested/allocated BCR is coded on 16 bits, the most significant bit is reserved, the next bit is the *nz* bit, the next five bits contain the exponent, and the remaining bits the mantissa.

User OAM BCR

Similar to the requested/allocated BCR for CLP = 0 + 1 cell flow but applies to the user OAM cell flow.

Block size and sequence number

These fields are used in ABT/IT (see 6.6.2.3).

6.6.1.4 Conformance definition and QoS commitments for ABT/DT

Conformance for ABT/DT at a standardized interface is defined at the cell level; in addition, conformance at the block level is defined if the sustainable cell rate is larger than zero. The cell level conformance definition includes conformance of RM cells and of the cells within a block with respect to the current block cell rates. The block level conformance definition is tested against the sustainable cell rate.

Both conformance definitions depend on RM cells that pass the interface. General principles for the ABT/DT conformance definition are described in 6.6.1.4.1.

6.6.1.4.1 General principles of conformance definition for ABT/DT

The control messages that define the ATM block for ABT/DT at an interface are described in Annex C.

RM cells delineating ATM blocks in the forward direction are:

1) either BCR decrease RM cells sent by the source (TM = 0); or
2) acknowledgement RM cells sent by the source in response to:
   – a positive acknowledgement sent by the network following a BCR increase request by the source;
   – a BCR modification initiated by the destination or by the network.

It is desirable not to have multiple outstanding BCR negotiations. This can be done by introducing priority levels between BCR negotiations (see Annex E).

A network should not initiate a BCR negotiation while another with the same or a higher priority level is pending.

Conformance for an ABT connection is tested against:

1) the BCR value of the user data CLP = 0 + 1 cell flow and optionally the user OAM cell flow (cell conformance);
2) the sustainable cell rate for the aggregate CLP = 0 + 1 (including user OAM) cell flow of an ABT/DT connection (ATM block conformance).

6.6.1.4.2 Conformance of RM cells

6.6.1.4.2.1 Conformance of user-generated RM cells

Conformance of request RM cells sent by the user is defined at a given interface by a GCRA(T_{RM}, τ_{RM}), where $1/T_{RM}$ is the peak cell rate of the ABT/DT request RM cell flow and τ_{RM} is the associated CDV tolerance.

Conformance of an acknowledgement RM cell sent by the user following a user or a network request is checked against the three following tests:

1) It is the response of the source to either an acknowledgement RM cell or a request RM cell sent by the network to the source (see Annex C).
2) It arrives within a time-out interval after the RM cell sent by the network to the source it responds to has crossed the interface. The time-out value depends on the round-trip time from the interface to the source. This value is either determined by the network operator or, if applying to an INI, negotiated between network operators. It may be specified on a subscription basis or on a per-connection basis.
3) It conveys information (BCR values, Sequence Number, CI bit, etc.) consistent with the message sent by the network. In particular, valid BCR values are BCR values less than or equal to the BCR values conveyed by the RM cell sent by the network to the source.

The processing of non-conforming RM cells is network-operator-specific. If an acknowledgement RM cell sent by the user arrives after the time-out has expired or if the content of such a cell is invalid, the network may not meet QoS commitments. The actions taken by the network under such conditions (e.g., defined recovery procedures) are not specified in this Recommendation.

6.6.1.4.2.2 Conformance of network-generated RM cells

Network-generated RM cells are conforming up to a certain limit fixed by mutual agreement between network operators.

6.6.1.4.3 Dynamic GCRA for ABT/DT

In ABT/DT, cell conformance is tested by a dynamic GCRA for both user data and user OAM cells.

As soon as a BCR greater than 0 is negotiated for the user OAM cell flow, cell conformance is tested separately for user OAM cells. Thus, for an ABT/DT connection, cell conformance is checked:

i) against the BCR dynamically negotiated for the CLP = 0 + 1 cell flow;

ii) against the BCR of the OAM cell flow as soon as the BCR allocated to this user OAM cell flow is greater than 0.

Since the BCR of the cell flows of an ABT/DT connection may vary in time, conformance testing algorithms should take into account the BCR modifications performed by means of some RM cells. Thus, some specific RM cells should be interpreted by these algorithms, namely:

– RM_1 bandwidth decrease RM cells on the forward direction with TM = 0;

– RM_2 acknowledgement RM cells sent on the forward direction (with TM = 0 or 1).

The cell conformance algorithm makes use of the Last Virtual Schedule Time (LVST), which is the scheduled time of the last conforming data cell, instead of the usual Theoretical Arrival Time (TAT).

The cell conformance algorithm is depicted in Figure 11.

The following notation is used in Figure 11:

$\lambda(x)$ current BCR of cell flow x;

$T(x)$ current peak emission interval of the component x that corresponds to BCR $\lambda(x)$;

$\tau(\lambda(x))$ CDV tolerance used to test conformance of cell flow x for the allocated BCR $\lambda(x)$, the function $\tau(\lambda)$ is specified at connection establishment for user data cell flows, a unique value may be specified; for OAM traffic $\tau(\lambda)$ should be consistent with the standardized default rule specified for OAM traffic (see Appendix II); if $\lambda = 0$, t takes a default value;

$\Lambda(x)$ BCR of cell flow x conveyed in a specific ABT/DT RM cell;

$T(\lambda)$ peak emission interval corresponding to the BCR λ in the standardized ATM layer peak cell rate granularity list given in 5.4.1.2; if $\lambda = 0$, T takes a default value equal to the maximum value supported by the network;

$PCR(\Lambda)$ denotes the nearest greater value in the ATM layer peak cell rate granularity list corresponding to rate Λ;

x denotes the CLP = 0 + 1 or OAM cell flow.

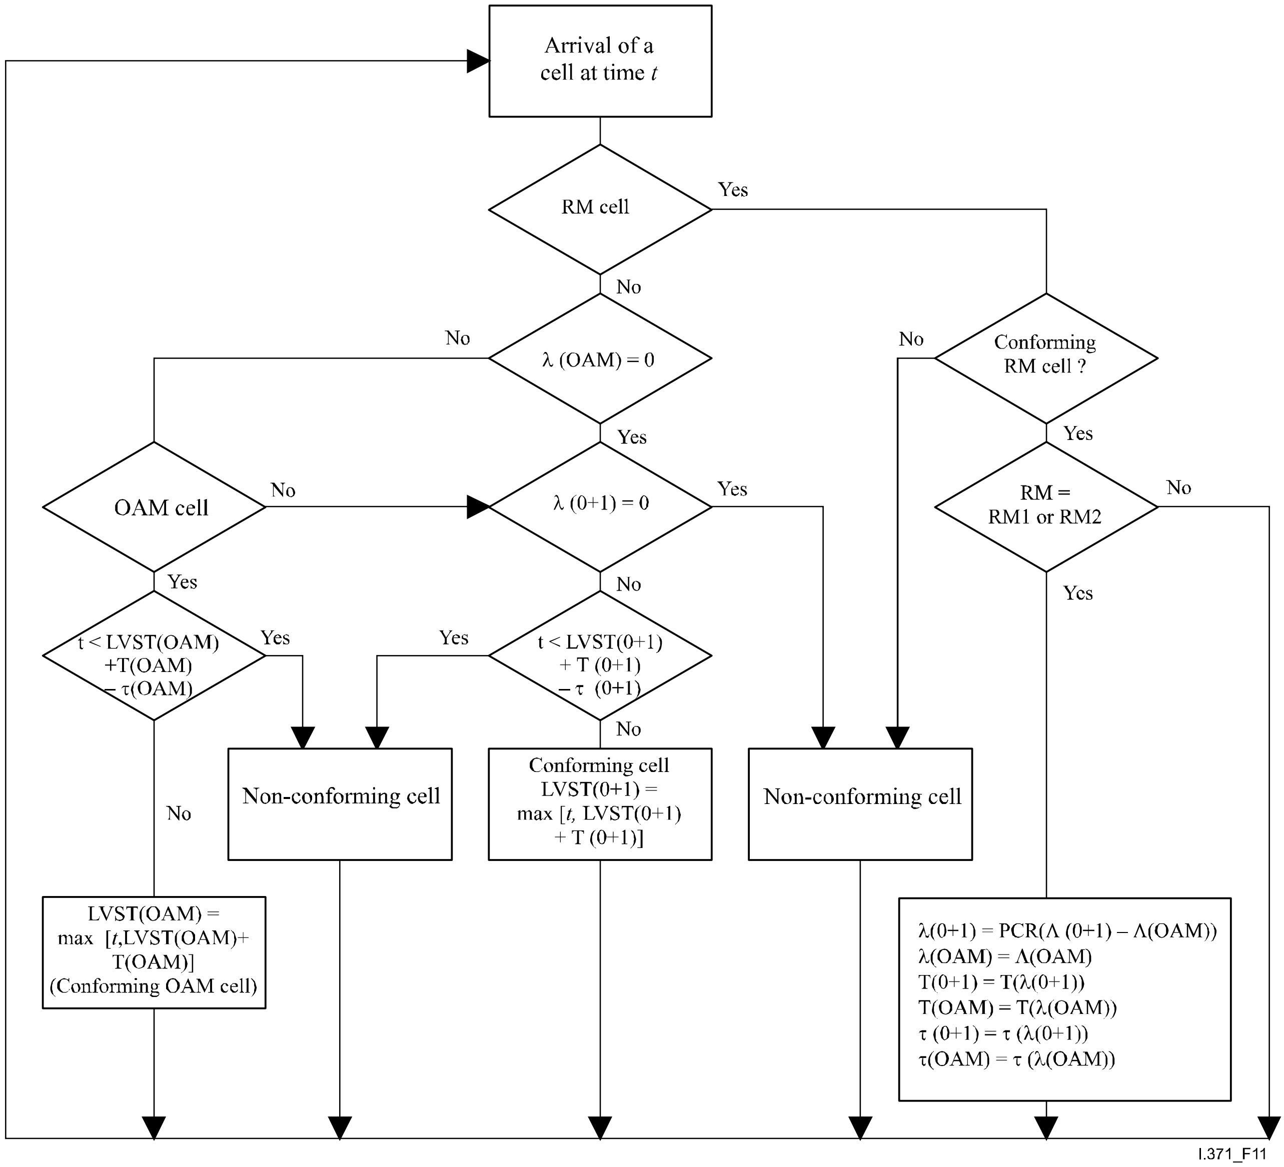

NOTE 1 – LVST(0+1) and LVST(OAM) are initialized to $-\infty$, which is a default value for identifying the first cell of an ATM connection; λ(0+1) and λ(OAM) are initialized to 0.

NOTE 2 – By definition, $\Lambda(0 + 1) \geq \Lambda(OAM)$; otherwise, the peak cell rate renegotiation would be invalid.

Figure 11/I.371 – Cell conformance for an ABT/DT connection

6.6.1.4.4 ATM block conformance for ABT/DT

ATM block conformance is tested against the sustainable cell rate, if greater than 0, that is specified for the CLP = 0 + 1 cell flow. ATM block conformance testing relies on an algorithm which computes a number of credits. ATM blocks are non-conforming when the number of credits is null. Moreover, the ATM block conformance algorithm makes use of a virtual time u defined at the arrival time of a cell as the maximum between this arrival time and the LVST of the last conforming CLP = 0 + 1 data cell, which is computed by the cell level conformance algorithm (see 6.6.1.4.3). More precisely, u = max{LVST,t} where t is the current time.

The sustainable cell rate Λ_{SCR} and the tolerance τ_{SCR} used in this conformance algorithm are those valid at the interface considered and deduced from the sustainable cell rate Λ^0_{SCR} and the maximum burst size MBS^0, negotiated at connection establishment, as (see Appendix VI):

$$\Lambda_{SCR} = \min\left(\Lambda^0{}_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \left(\frac{1}{T_{RM}} + \frac{1}{T'_{RM}}\right), \frac{1}{T}\right)$$

$$\tau_{SCR} = \left(MBS^0 - 1 + \frac{1}{T} \times \tau''_{SCR} \times \left\lfloor 2 + \frac{\tau_{RM}}{T_{RM} - \Delta} + \frac{\tau'_{RM}}{T'_{RM} - \Delta} \right\rfloor\right)(T_{SCR} - T)$$

where:

1) 1/T is the peak cell rate of the connection and T_{SCR} is the emission interval corresponding to Λ_{SCR};

2) it is assumed that forward and backward user request RM cell flows issued by both users of the ABT/DT communication are at the interface considered conforming to GCRA(T_{RM}, τ_{RM}) and GCRA(T'_{RM}, τ'_{RM}), respectively;

3) τ''_M is the difference between the maximum and the minimum (or equivalently remote quantiles) of the virtual transfer delays for RM cells delineating ATM blocks. The virtual transfer delay for an RM cell delineating an ATM block is defined as the difference between the time when the RM cell is transmitted at the PHY-SAP of the equivalent terminal and the virtual time u when it is received at the interface;

4) Δ is the cell transmission time (in seconds) at the interface link speed.

ATM block conformance is tested by taking into account the volume of reserved resources. The ATM block conformance testing algorithm is depicted in Figure 12. The principles of ATM block conformance are as follows (see Figure 13):

– Specific RM cells are RM cells of RM_1 and RM_2 types defined above.

– ATM block conformance is tested slot by slot by comparing the current time t with a non-conformance date t_{exp}.

– The non-conformance date is relevant only when the allocated BCR λ is greater than the sustainable cell rate $\Lambda_{SCR} = 1/T_{SCR}$.

– The non-conformance time t_{exp} is computed by using a variable X, the sustainable cell rate Λ_{SCR}, and the BCR λ allocated to the cell flow.

– X is updated at each specific RM cell arrival time and represents the number of credits for the new BCR reservation (X is computed by using the BCR allocated to the previous BCR reservation).

– Upon reception of a forward bandwidth acknowledgement RM cell with the Traffic Management and Maintenance bits set equal to 1 and 0, respectively, the number X of credits is reset to 0. This is intended to realign ATM block conformance algorithms when a policing procedure is run.

– The non-conformance time t_{exp} and variable X are computed from the maximum of the current time and the Last Virtual Schedule Time (LVST) of the cell flow (LVST is computed by the cell conformance test for the CLP = 0 + 1 cell flow).

The following relations hold:

$$\begin{cases} X = \min\left\{\dfrac{T_{SCR} + \tau_{SCR}}{T_{SCR}}, \left[X - (\lambda - \Lambda_{SCR})(u - t_{LVMT})\right]^{+}\right\} \\ \lambda = \Lambda, t_{\exp} = u + \dfrac{X}{\lambda - \Lambda_{SCR}} \text{ if } \lambda > \Lambda_{SCR} \end{cases}$$

where u is the virtual time, t_{LVMT} the virtual time corresponding to the previous BCR modification, namely Last Virtual Modification Time (LVMT), and $x^{+} = \max\{0,x\}$.

– The non-conformance time is relevant only if $\lambda > \Lambda_{SCR}$; otherwise, the need of the source is less than expected and the ATM block is conforming.

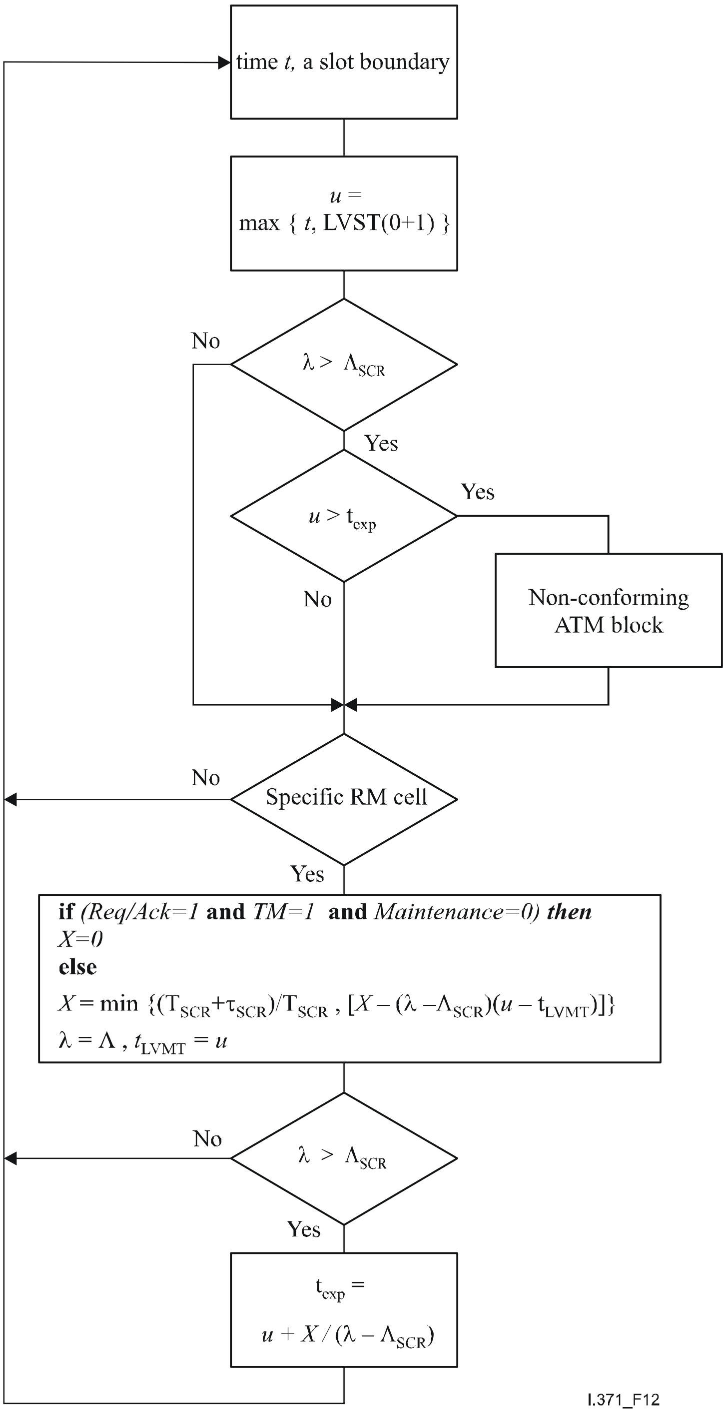

NOTE 1 – t_{LVMT} and u are initialized to $-\infty$, which is a default value for identifying the first cell of an ATM connection, X is initialized to $\frac{T_{SCR} + \tau_{SCR}}{T_{SCR}}$, λ is initialized to 0.

NOTE 2 – BCRs and LVST are those of the cell flow considered, LVST is given by the cell conformance algorithm (see Figure 11).

Figure 12/I.371 – ATM block conformance for a cell flow of an ABT/DT connection

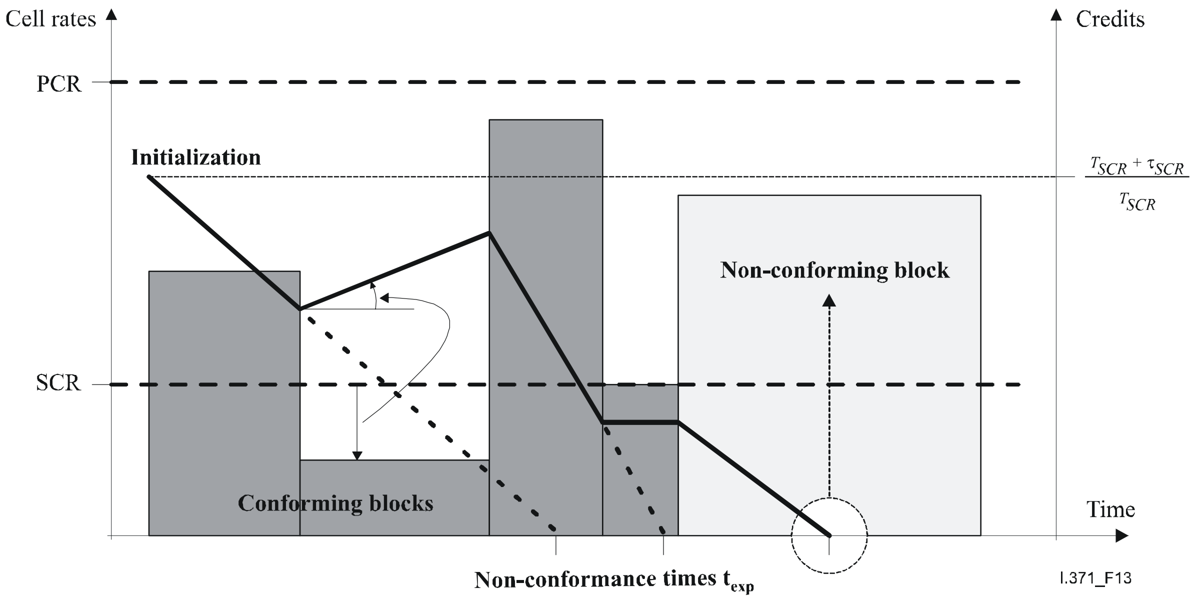

NOTE – This figure is for illustration purposes. The bold solid line represents the instantaneous number of credits. The slope of this curve is equal to the difference between the SCR and the BCR. Computation of t_{exp} for a given ATM block is based on the number of credits available at the ATM block boundary.

Figure 13/I.371 – Examples of evolution of block conformance variables

Loss of RM cells may unduly produce block non-conformance and, in the case of lost RM_1 or RM_2 cells, may require recovery or reinitialization of conformance variables. ATM block conformance algorithms are realigned by policing procedures (see 7.2.3.6).

QoS aspects

At call establishment or at connection subscription time, the user may negotiate a QoS class for the connection. If QoS commitments apply, then they apply both at cell and ATM block level. QoS commitments at the cell level may include end-to-end CDV and CLR objectives.

Cell level QoS commitments

The basic QoS commitment at cell level that is offered to the connection is that, as long as, within an ATM block, the connection is conforming to the BCR negotiated for the ATM block, the cell level QoS commitments pertain. It is expected that, within an ATM block, this QoS will be equivalent to that normally offered to DBR with the same PCR and CDV tolerance and the same negotiated QoS class. Furthermore, cell level QoS is assured to all cells when all cells are conforming to the relevant BCR conformance tests. If some cells are non-conforming to the relevant BCR conformance tests, the network may consider the connection as non-compliant and as a consequence, the network need not meet the QoS commitments. If the network chooses to offer cell level QoS commitments to a connection with non-conforming cells, the ATM layer QoS is only assured to a volume of cells that is conforming to all relevant BCR conformance tests.

Block level QoS commitments

In the case when a sustainable cell rate greater than 0 has been specified at connection establishment and a QoS class with specified CLR objectives is negotiated, the commitments the network can make to the connection at the ATM block level is that, as long as ATM blocks are conforming to the sustainable cell rate traffic descriptor, a new BCR reservation should be accepted by the network within finite time limits. These limits are part of the QoS class negotiated at call establishment. In the case of ATM block non-conformance or if the sustainable cell rate is set equal to 0, the network does not assure any block level QoS commitment. Furthermore, in that case, the network may initiate a BCR renegotiation. The cell level QoS is in any case assured as long as the network has not renegotiated the allocated BCR.

Resource management in the case when an SCR is set equal to 0 is performed on the basis of suitable traffic engineering rules. The network may thus achieve some block level delay objectives (QoS indications on the time of access to network resources). These QoS indications are not contractual between the user and the network but indicative only.

NOTE – The delay experienced by an application using ABT/DT is not constant over the connection lifetime.

6.6.2 ABT with immediate transmission (ABT/IT)

6.6.2.1 Definition and service model

In ABT/IT, the user transmits ATM blocks without positive acknowledgement from the network. As a result, ABT/IT ATM blocks may be discarded by the network if sufficient network resources are not available. The ATM block loss probability can be made small by reserving resources via an SCR. As with ABT/DT, there is a PCR which is negotiated at call establishment. For each ATM block, there is a BCR associated with that block. If the ATM block is marked as "elastic" (elastic/rigid bit = 0), then a network element can buffer the ATM block, reduce the ATM block BCR, and transmit the ATM block at a new BCR.

During the connection lifetime, ATM blocks are directly transmitted onto the network by the traffic source. The leading ABT/IT RM cell of an ATM block, which may also be the trailing ABT/IT RM cell of the previous ATM block, requests network resources computed on the basis of the BCRs of the different cell flows of the ABT/IT connection conveyed by this ABT/IT RM cell. The trailing ABT/IT RM cell, which may also be the leading ABT/IT RM cell of the subsequent ATM block, releases network resources or requests resources for the subsequent ATM block. If resources needed for the transfer of an ATM block are available within the network, the ATM block is transferred; otherwise, the ATM block is discarded.

A committed bandwidth for the user data and user OAM CLP = 0 + 1 cell flow is specified as follows.

Definition (Committed bandwidth)

A committed bandwidth for the user data and user OAM CLP = 0 + 1 cell flow in a given direction is defined when a sustainable cell rate greater than 0 is specified at connection establishment for this cell flow in the relevant direction and a QoS class with specified CLR objectives is negotiated. The committed bandwidth is related to the amount of reserved resources and its value is equal to that of the specified sustainable cell rate. The long-term average of resources, which can be potentially reserved, is at least as large as the committed bandwidth. Moreover, if the traffic is conforming to the specified sustainable cell rate traffic descriptor negotiated for this cell flow (see 6.6.2.4), then the probability that an ATM block transfer fails (i.e., the ATM block is discarded in the network) is lower than a given threshold.

No committed bandwidth is assured by the network to the cell flow if a sustainable cell rate equal to 0 is specified for this cell flow or if QoS class U has been negotiated. In these cases, ABT/IT does not provide for any commitments on the success of the transfer of an ATM block and there is no guarantee on the ATM block discarding probability. Note, however, that the network may try to achieve some objectives for this probability (QoS indications) without requiring any sustainable cell rate specification.

6.6.2.2 Source traffic descriptor and CDV tolerances

The source traffic descriptors and the relevant CDV tolerances for ABT/IT are identical to those for ABT/DT (see 6.6.1.2).

6.6.2.3 Dynamically changing parameters and RM cell format for ABT/IT

The dynamically changing parameters and RM cell format for ABT/IT are identical to those for ABT/DT and are given in Table 2 except that ABT/IT RM cells are identified by the protocol identifier 3.

Message type and CLP = 0 + 1 BCR

The message type fields (direction, traffic management, congestion indication, maintenance, req/ack, elastic/rigid) and the CLP = 0 + 1 BCR and OAM BCR fields are as in ABT/DT RM cells (see 6.6.1.3).

Block size

The block size field conveys the length of the ATM block expressed in cells. The value carried in this field is indicative only and may be used by specific implementations. It is not used in the conformance definition.

Sequence number

A source may optionally make use of the sequence number field by incrementing the value of the sequence number (modulo 2^{32}) in each subsequent ABT/IT RM cells. In the case it does, the following applies:

– the sequence number is always present in ABT/IT RM cells;

– it is always incremented by one in ABT/IT RM cells.

A source that does not make use of the SN field sets its value to 0.

No other entity than the source is allowed to modify the sequence number.

When the destination sends an ABT/IT RM cell in response to an ABT/IT RM cell generated by the source, the sequence number is copied unchanged into that cell. Any other RM cell generated by the network or the destination must have the sequence number set equal to 0.

6.6.2.4 Conformance definition and QoS commitments for ABT/IT

Conformance for ABT/IT at a standardized interface is defined at the cell level; in addition, conformance at the block level is defined if the sustainable cell rate is larger than zero. The cell level conformance definition includes conformance of RM cells and of the cells within a block with respect to the current block cell rates. The block level conformance definition is tested against the sustainable cell rate. Both conformance definitions depend on RM cells that pass that interface.

6.6.2.4.1 Cell conformance for ABT/IT

Cell conformance for ABT/IT is identical to cell conformance for ABT/DT except that:

- RM cells delineating ATM blocks are:
  1) either block cell rate modification request RM cells (Traffic Management = 0) sent by the source; or
  2) acknowledgement RM cells with Traffic management = 1 sent by the source on the forward direction in response to a BCR modification initiated by the network;
- the user should send only forward request RM cell. Backward request RM cells are non-conforming.

The specific RM cells to take into account in the conformance definition are then:

– RM_1: conforming bandwidth increase or decrease request RM cells sent by the source (TM = 0);

– RM_2: acknowledgement RM cells (TM = 1) sent by the source on the forward direction in response to a BCR negotiation initiated by the network.

6.6.2.4.2 ATM block conformance for ABT/IT

The ATM block level conformance algorithm for ABT/IT is identical to that for ABT/DT (given by Figure 12), except that the specific RM cells to take into account are conforming bandwidth increase or decrease request RM cells sent by the source (TM = 0) and acknowledgement RM cells (TM = 1) sent by the source on the forward direction. Furthermore, the sustainable cell rate Λ_{SCR} and the tolerance τ_{SCR}, taken into account in the ATM block conformance definition, are those valid at the interface considered and deduced from the sustainable cell rate Λ^0_{SCR} and the maximum burst size MBS^0 negotiated at connection establishment (see Appendix VI) as:

$$\Lambda'_{SCR} = \min\left(\Lambda^0{}_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \frac{1}{T_{RM}}, \frac{1}{T}\right)$$

$$\tau_{SCR} = \left(MBS^0 - 1 + \frac{1}{T} \times \tau''_{SCR} \times \left\lfloor 1 + \frac{\tau_{RM}}{T_{RM} - \Delta} \right\rfloor\right)(T_{SCR} - T)$$

where the notation of 6.6.1.4.3 is used.

Loss of RM cells may unduly produce block non-conformance and, in the case of lost RM_1 or RM_2 cells, may require recovery or reinitialization of conformance variables. ATM block conformance algorithms are resynchronized when a policing action is initiated by a network along the ABT/IT communication.

QoS aspects

At call establishment or at connection subscription time, the user may negotiate a QoS class for the connection. If QoS commitments apply, then they apply both at cell and ATM block level. QoS commitments at the cell level may include end-to-end CDV and CLR objectives.

Cell level QoS commitments

The cell level QoS is composed of CLR and end-to-end CDV commitments. The basic CLR commitment at the cell level that is offered to the connection is that, once the BCR has been accepted within the network for an ATM block by means of ABT/IT RM cells, the cells of that ATM block are transferred with a CLR equivalent to that offered to a DBR connection with the same PCR and CDV tolerance and the same negotiated QoS class. Moreover, if an ATM block has the elastic bit equal to 1 (i.e., no shaping intended to reduce the BCR of the ATM block is allowed within the network), the cells of the ATM block are transferred with an end-to-end CDV equivalent to that offered to the connection using a combination of DBR and a QoS class with specified CDV objectives as above. If, by means of the elastic/rigid bit set to 0, the user accepts that an ATM block be shaped within the network at a BCR that is strictly smaller than the BCR requested for the block, then the end-to-end CDV objectives are unspecified.

The cell level QoS is assured to all cells when all cells are conforming to the relevant BCR conformance tests. If some cells are non-conforming to the relevant BCR conformance tests, the network may consider the connection as non-compliant and as a consequence, the network need not meet the QoS commitments. If the network chooses to offer cell level QoS commitments to a connection with non-conforming cells, the ATM layer QoS is only assured to a volume of cells that is conforming to all relevant BCR conformance tests.

In ABT/DT, ATM blocks are negotiated (require an implicit or an explicit acknowledgement by the network) whereas in ABT/IT, ATM blocks are accepted or discarded.

Block level QoS commitments

In the case when a sustainable cell rate greater than 0 has been specified at connection establishment and a QoS class with specified CLR objectives is negotiated, the commitments the network can make at the block level to the connection is that, as long as ATM blocks are conforming to the sustainable cell rate traffic descriptor, a new BCR reservation should be accepted by the network with a specified block discard probability, which is a block level QoS characteristic of the connection using ABT/IT. In the case of ATM block non-conformance or if the sustainable cell rate is set equal to 0, the network does not assure any block level QoS commitment. Furthermore, in that case, the network may initiate a BCR renegotiation. The cell level QoS is in any case assured as long as the network has not renegotiated the allocated BCR.

The block level QoS commitments in ABT/IT differs from the one in ABT/DT:

– In ABT/IT, ATM blocks may be discarded.
– In ABT/DT, the delay experienced by the application is larger than the one experienced by using ABT/IT due to the BCR negotiation phase in ABT/DT.

Resource management in the case when SCR is set equal to 0 is performed on the basis of suitable traffic engineering rules. The network may thus achieve some block level QoS objectives (QoS indications on block discard probability). These objectives are not contractual between the user and the network but indicative only (QoS indications).

6.7 Available bit rate transfer capability (ABR)

Many applications have the ability to reduce their information transfer rate if the network requires them to do so. Likewise they may wish to increase their information transfer rate if there is extra bandwidth available within the network. There may not only be static traffic parameters but also dynamic traffic parameters because the users are willing to accept unreserved bandwidth. To support traffic from such sources in an ATM network, an ATM transfer capability is defined which will be termed Available Bit Rate (ABR).

6.7.1 Definition and service model

ABR is an ATM transfer capability where the limiting ATM layer transfer characteristics provided by the network may change subsequent to connection establishment. It is expected that a user that adapts its traffic to the changing ATM layer transfer characteristics will experience a low Cell Loss Ratio (CLR). Cell delay variation and cell transfer delay are not controlled. The ABR capability is not intended to support CBR applications.

The user adapts to the changing ATM layer transfer characteristics upon receiving feedback from the network. Due to cell transfer delay, this feedback reflects the status of the network at some time prior to the instant when the user receives it. So even when the user adapts correctly to the feedback, the network may still have to provide some buffering to enable low cell loss operation of ABR.

User actions and its responses to the feedback from the network together with the feedback from the network constitute a control loop on the ABR connection.

A user will specify a maximum required bandwidth to the network on establishment of an ABR connection. The maximum required bandwidth is negotiated between user and network and between user and user at connection establishment. A minimum usable bandwidth (also referred to as the Minimum Cell Rate or MCR) shall be specified on a per-connection basis, but it may be specified as zero. The bandwidth available from the network may become as small as the minimum usable bandwidth. The maximum required bandwidth (also referred to as the Peak Cell Rate or PCR) and the MCR are defined by the GCRA. The value of the PCR and of the MCR can be different on forward and backward connections.

An ABR capability may apply to VCCs or to VPCs. If ABR applies to some VCCs within a VPC, the ABR VCCs share the capacity of the VPC that is allocated to ABR.

In ABR, user data cells have the CLP bit set to zero. Cell tagging (see 7.2.3.6) is currently not supported by ABR. It is for further study whether ABR may also use CLP = 1 user data cells and whether cell tagging may be applied.

The bandwidth available on an ABR connection is the sum of an MCR, which could be 0, and of a variable cell rate that results from sharing the available bandwidth among ABR connections via a defined allocation policy. A defined allocation policy means that the allocation that the user receives above the minimum cell rate is not only determined by what the user requests or submits, but also by network policy. Defined allocation policies are not subject to standardization. However, network stability requires that for a given configuration of ABR user bandwidth requests, the allocation policy should support convergence to a stable allocation of bandwidth within the network.

There is a quantitative commitment made by the network that the user can continuously send cells at the MCR and the committed CLR pertains.

A source is not precluded from sending at a rate less than MCR when MCR greater than 0 is negotiated. The MCR agreed between the end systems and the network(s) carrying the connection may range from 0 to the maximum value supported by the network(s). This maximum value may be 0. Although the network commits to support MCR, a source may receive indications to reduce its rate below the MCR. If a source receives such indication and its rate is above the MCR, it should reduce its rate to the MCR. Likewise, if a source receives such indication and its rate is at or below the MCR, the source need make no change to its rate.

For providing feedback from the network to the source, each ABR connection uses Resource Management (RM) cells. User RM cells that are included with the user data cells in the current allowed cell rate, called in-rate RM cells, have the CLP bit set equal to 0. RM cells that are excluded from the current allowed cell rate, called out-of-rate RM cells, have the CLP bit set equal to 1. Out-of-rate RM cells may be generated by a source or a network element during periods when the generation of in-rate RM cells is not suitable. It is expected that on each ABR connection, CLP = 1 RM cells are not systematically discarded. The procedures and restrictions regarding the generation of out-of-rate RM cells is currently not specified in this Recommendation.

Network elements and ABR destinations may insert RM cells for Backward Explicit Congestion Notification (BECN). Such cells are called BECN RM cells. The BECN bit of a BECN RM cell is set to 1. A BECN RM cell has CLP = 0. The purpose of a BECN cell is to notify the source to decrease or to not increase its rate.

Network assurances

For ABR the network does not offer quantitative assurances for bandwidth above the minimum usable bandwidth. In this case there are two types of assurances that the network can make: relative and procedural.

Relative assurances include, e.g., for connections that share the same path, no connection shall be arbitrarily discriminated against and no connection shall be arbitrarily favoured, although resources may be allocated according to a defined policy. Note that this relative assurance cannot depend on assumptions about how other sources behave.

Procedural assurances mean that, if all cells of the source are conforming to the conformance definition, then the network offers QoS commitments (see 6.7.5).

Flow control model for ABR

ABR flow control currently occurs between a sending terminal (source) and a receiving terminal (destination). Sources and destinations are connected via one connection for each direction. For a bidirectional ABR connection, each of the terminals is both a source and a destination. For the sake of simplicity, only the information flow from the source to the destination with its associated RM flows is considered. The forward direction is the direction from the source to the destination, and the backward direction is the direction from the destination to the source. For the information flow from the source to the destination, there is a control loop consisting of two RM flows, one in the forward direction and one in the backward direction (Figure 14).

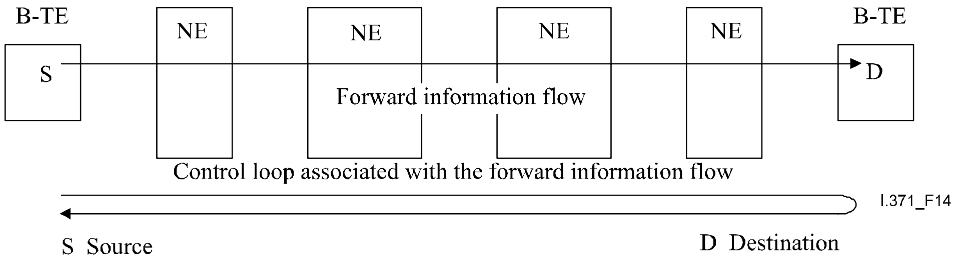

Figure 14/I.371 – Example of a source to destination ABR control loop

Recovery from error conditions due to RM cell loss is implementation specific.

Segmentation of the ABR control loop

An ABR connection may be segmented at selected network elements into two or more separately controlled ABR segments. Each ABR control segment (except the first) is sourced by a virtual source. A virtual source assumes the behaviour of an ABR source endpoint. Backward RM cells received by a virtual source are removed from the control loop.

Each ABR control segment (except the last) is terminated by a virtual destination. A virtual destination assumes the behaviour of an ABR destination endpoint.

Figure 15 illustrates an ABR virtual connection which incorporates segmentation.

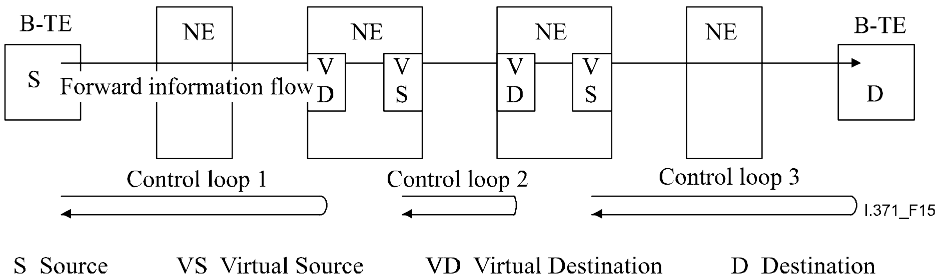

Figure 15/I.371 – Example of a segmented virtual connection

The coupling between two adjacent control segments associated with an ABR connection (e.g., within a network element, or across a group of network elements whose operation is not specified) is implementation specific. The effect of such configurations on the performance of the ABR service requires further study.

6.7.2 Source traffic descriptor and CDV tolerances

There are three types of traffic parameters associated with ABR:

1) traffic parameters which are in the source traffic descriptor, are negotiated at call establishment and cannot be changed by RM procedures. These traffic parameters are taken into account in the ABR conformance definition;
2) traffic parameters which are not in the source traffic descriptor, which are negotiated at call establishment and cannot be changed by RM procedures (see 6.7.3);

3) dynamically changing traffic parameters which may be changed by RM procedures (see 6.7.4).

At call establishment, the user negotiates a source traffic descriptor for the connections in each direction that consists of:

– A PCR(0+1) for the aggregate user generated cell flow, which is the maximum cell rate requested by the application (see 5.4.1).
– An MCR(0) traffic parameter for the user generated CLP = 0 flow which is the minimum cell rate requested by the application. MCR has location, basic event and coding identical to the PCR (see 5.4.1).
– Initial Allowed Cell Rate for the user generated CLP = 0 cell flow: IACR(0). IACR has location, basic event and coding identical to the PCR (see 5.4.1). In particular, IACR is positive. Also, IACR is less than or equal to PCR.

In the ABR transfer capability, the user OAM cell flow and the user RM cell flow are aggregated with the user data cell flow in the traffic parameters and in the conformance definition. When a user OAM component is included in the user cell flow, the user should select source traffic parameters, in particular MCR, to also accommodate the user OAM component.

In the ABR transfer capability, the operation of the control loop requires that RM cells are carried in the backward direction (see 4.1). The user should negotiate source traffic parameters PCR(0+1), MCR(0) and IACR(0) for the connection in the backward direction to also accommodate the traffic related to the operation of the control loop.

With respect to conformance, CDVT values are required at each standardized interface for the following traffic parameters: PCR(0+1), MCR(0) and IACR(0), see 6.7.5. For all three parameters the CDV tolerance τ_1 associated with the ACR is used. In addition to the above source traffic descriptors, the following values of tolerances are required (see 6.7.5.1):

– the CDV tolerance τ_1 to be associated with the dynamically changing value of ACR;
– the upper bound to the round-trip feedback delay between the interface and the source: τ_2;
– the lower bound to the round-trip feedback delay between the interface and the source: τ_3.

The CDV tolerance τ_1 can be conveyed via signalling or assigned on a per subscription basis. The same coding scheme that applies to τ_{PCR} applies to τ_1 (see 5.4.1.3). It is for further study whether the tolerance parameters τ_2 and τ_3 may be conveyed by signalling or shall only be set on a per subscription basis.

6.7.3 Additional ABR traffic parameters used in the reference behaviour

The following ABR connection parameters are used in the source reference behaviours but are not used in the conformance definition.

– Rate Decrease Factor (RDF) is the parameter that controls the decrease in the cell transmission rate. It is assigned by the network.
– Rate Increase Factor (RIF) is the parameter that controls the increase in the cell transmission upon receipt of an RM cell. It is assigned by the network.
– For each forward RM cell, N_{RM} is the maximum number of in-rate cells (including this particular RM cell) an ABR source may send.

The present signalling protocol conveys also the parameters FRTT and TBE. These two parameters are defined as follows:

– The Fixed Round Trip Time (FRTT) is an estimate of the minimum round trip time for the connection. The value of FRTT is assigned to a connection and provided to the user. This value is an indication by the network and not a commitment. The assignment of that value can be made by network management procedures or by means of signalling.

– Transient Buffer Exposure (TBE) is the number of cells that the network would like to limit the source to sending during start-up periods, before the first RM cell returns. It is assigned by the network.

6.7.4 Dynamically changing traffic parameters and RM cell format for ABR

The values of the following dynamic parameters are determined by network elements along the connection and are forwarded to the user of the ABR capability via RM cells: explicit cell rate (ECR), CI and NI, and QueueLength. The Allowed Cell Rate (ACR) is the maximum allowed rate that the source derives from the above feedback from the network. The ACR ranges between those of the MCR(0) and the PCR(0+1).

NOTE – EFCI may be used as a feedback parameter by the application.

Table 3 contains the RM cell format for the ABR transfer capability.

Table 3/I.371 – Format for the RM cell that supports ABR

Field	Octet(s)	Bit(s)	Coding
Header	1-5	All	As in ITU-T Rec. I.361
Protocol ID (Note 7)	6	All	1
Message type: Direction	7	8	(Note 1)
Message type: BECN indication	7	7	
Message type: Congestion indication	7	6	(Note 2)
Message type: No-increase	7	5	(Note 3)
Message type: Reserved	7	1-4	Clause 8.1
Explicit Cell Rate (ECR)	8-9	All	(Note 4)
Current Cell Rate (CCR)	10-11	All	(Note 4)
Minimum Cell Rate (MCR)	12-13	All	(Note 4)
Queue length	14-17	All	(Note 5)
Sequence number	18-21	All	(Note 6)
Reserved	22-51	All	Clause 8.1
Reserved	52	3-8	Clause 8.1
CRC-10	52	1-2	
	53	All	

Table 3/I.371 – Format for the RM cell that supports ABR

NOTE 1 – The direction bit equals 0 for forward RM cells and equals 1 for backward RM cells.
NOTE 2 – The congestion indication bit equals 1 to indicate congestion and equals 0 otherwise.
NOTE 3 – The no-increase bit equals 1 to indicate that no rate increase by the source is allowed and equals 0 otherwise.
NOTE 4 – The coding is based on the 5-bit exponent and 9-bit mantissa coding used for the peak cell rate. The specifics are given in 6.7.4.1 below.
NOTE 5 – Integer valued. Least significant bit is bit 1 of octet 17.
NOTE 6 – Integer valued. Least significant bit is bit 1 of octet 21.
NOTE 7 – ABR and ABT connections use their own RM cells with a specified Protocol ID for traffic management purposes (PID = 1 for ABR, PID = 2 for ABT/DT, PID = 3 for ABT/IT). All other RM cells on the same level (VPC or VCC) with a different PID are non-conforming and may be discarded at the entrance of a network or where the first RM cell processing is done for this connection in the network.

6.7.4.1 Details on the fields

Protocol ID

ABR RM cells are identified by protocol ID 1.

Message type

Message type is an octet containing four single bit fields plus four reserved bits. The interpretation of the defined bits is as follows:

Direction: This bit distinguishes RM cells travelling from a source ("forward" cells) to a destination from RM cells travelling from a destination ("backwards" cells) to a source.

BECN indication: This bit distinguishes a normal RM cell generated by a source and looped back by the destination from a RM cell generated by an intermediate congested switch or an ABR destination. The BECN indication is set to 1 in the BECN RM cell.

Congestion Indication (CI): This bit indicates congestion or impending congestion in the forward path.

No-increase (NI): This bit, when used in combination with the CI bit, can indicate to the source that it should continue to send at its current rate, which is desirable when the network is in a steady state or to avoid unnecessary oscillations. In particular, when CI = 0 and NI = 0 then the source can increase the sending rate, and when CI = 0 and NI = 1, the allowed cell rate is not increased.

Explicit Cell Rate (ECR)

This field is set by the source to a value that is at most the negotiated peak cell rate and may be reduced by an intermediate switch to notify the source of the Allowed Cell Rate (ACR) at that switch. The ECR value received by a source will explicitly determine the maximum cell rate of the source. The coding is the 14-bit binary floating point representation used for the peak cell rate in 5.4.1.2, which employs a 5-bit exponent, m and a 9-bit mantissa, k, plus a 1-bit nz field as described below:

$$ECR = \left[2^m \cdot \left(1 + \frac{k}{512}\right)\right] \cdot nz \text{ cells per second}$$

$$0 \le m \le 31 \text{ and } 0 \le k \le 511$$

$$nz = \{0,1\}$$

The following are the bit positions within a 16-bit word: the most significant bit is reserved; the next bit contains the value of *nz*; the next five bits contain the value of *m*; the remaining nine bits contain the value of *k*.

Current Cell Rate (CCR)

The CCR field contains the allowed source cell rate that was in effect at the time the RM cell was transmitted by the source. The information in this field can optionally be used in the computation of the value of the Explicit Cell Rate (ECR) field described above. The same encoding and format used for the ECR field is adopted for the CCR field.

Minimum Cell Rate (MCR)

This field contains the minimum cell rate, which is determined at connection establishment. The information in this field can optionally be used in the computation of the value of the Explicit Cell Rate field which is described above. The same encoding and format used for the ECR field is adopted for the MCR field.

Queue length

The QueueLength parameter is optionally supported by network elements. It represents the maximum number of cells currently queued for this connection among those network elements supporting this parameter. Any given network element writes into this field the maximum of the current value of the field and of the number of cells of the given connection queued at this network element. It is set to zero by the source.

If the network element does not know the "number of cells of the given connection queued at the given buffer at that network element", then the network element leaves the value of the field unchanged.

Sequence number

A source/virtual source may optionally make use of the sequence number field by incrementing its value by one (modulo 2^{32}) in each subsequent forward RM cell emitted by the sender. In the case it does, the following applies:

– the sequence number is always present in RM cells;

– it is always incremented by one in forward RM cells by the sender.

A source that does not make use of the SN field sets its value to 0.

No other entity than the source may modify the sequence number.

When the destination makes a backward RM cell from a forward RM cell, the sequence number is copied unchanged into the backward RM cell. Any backward RM cell not generated from a forward RM cell must have the sequence number set to zero and the BECN bit set to 1.

6.7.5 Conformance definition and QoS commitments for ABR

The following conformance definition applies to the cell flow consisting of user-generated cells and in-rate RM (CLP = 0) cells, excluding BECN RM cells.

NOTE – Although RM cells with CLP = 0, including BECN cells, are included in the current allowed cell rate (according to 6.7.1), BECN cells are excluded from the flow tested by the conformance definition. It follows that, where equipment emit BECN cells as part of the current allowed cell rate flow, these BECN cells will not cause any loss of conformance.

Conformance of BECN cells is determined by mutual agreement between source/networks. A policer could still police the aggregate CLP = 0 cell flow by setting margins on the policed rate.

Out-of-rate (CLP = 1) user data cells are non-conforming. Definition of conformance for out-of-rate (CLP = 1) RM cells is not addressed in this Recommendation.

The concepts of compliance of an ABR connection and conformance of individual cells on that connection define the conditions under which a network operator is responsible for supporting QoS objectives for the connection. Conformance applies to cells as they are tested upon their arrival at the UNI or inter-network interface. Each then is either conforming or non-conforming. Based in part on the results of the conformance test, a network operator will designate a connection as either compliant or non-compliant.

If some cells are non-conforming to some of the relevant conformance tests, the network may consider the connection as non-compliant (see 5.3.2). If the network chooses to offer QoS commitments to a connection with some non-conforming cells, the ATM layer QoS is only assured to a volume of cells that is conforming to all relevant conformance tests. The precise definition of a compliant ABR connection is left to the network operator. Any definition of a compliant ABR connection shall find a connection compliant if all cells on the connection are conforming and if the RM cells on the connection satisfy the requirements, if any, of the mechanism implemented by the network operator(s).

For compliant connections at the UNI or inter-network interface, 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.

A source receives feedback information from backward RM cells. Feedback may include information in the Explicit Cell Rate (ECR) field, the queue length field, the Congestion Indicator (CI) bit, the No Increase (NI) bit of each backward RM cell on the companion backward connection. A source that behaves as described in Appendix VII would be conforming.

Checking the values of the CCR and MCR cell fields is not part of the ABR conformance definition.

Note that in the ABR Capability, a source is not required to send RM cells. However, if there is not a flow of user-generated, backward RM cells, and if the network wishes to convey feedback to the user, the network can make use of the capability of itself generating backward (BECN) RM cells (see 6.7.4.1).

6.7.5.1 Definitions of ABR delays used in the conformance definition

The algorithm that defines conformance at an interface should take into account the delays between when a new rate is known at the interface and when cells arrive to the interface that have been emitted by the source after the new rate is known by the source. These delays are variable.

The characteristics of traffic received at the UNI or inter-network interface on a given ABR connection depend critically on the delays between that interface and the source (or virtual source) that generates the traffic. The delays most relevant to the characteristics of a flow received at the interface are defined relative to the transmission times of each cell by the traffic source. Note that the source may have cells queued for transmission. The next cell to be transmitted would (nominally) be scheduled for transmission according to the reciprocal of the current ACR. While waiting, a backward RM cell could arrive and a new ACR determined. The source could plausibly leave the scheduled transmission time of the lead cell unchanged, or could update the scheduled transmission time according to the new ACR. In the context of the conformance definition, it is assumed that the source could choose the alternative that yields the earlier transmission time. Hence, a transmission time for a cell is called an Ideal Transmission Time (ITT) if the difference between itself and the transmission time for the previous cell on the connection is greater than or equal to the minimum of:

a) the inverse of the ACR in effect immediately after the transmission time of the first of the two cells; and

b) the inverse of the ACR in effect immediately before the transmission time of the second of the two cells.

The transmission time for the first cell on the connection is automatically an ITT.

Two delays, t_1 and t_2, are particularly relevant to traffic characteristics at an interface:

- The delay t_1 denotes the time from a cell's transmission time by the traffic source to its receipt at the interface in question.
- The delay t_2 denotes the sum of:
  1) the delay from the departure at the interface in question of a backward RM cell on the backward connection to the receipt of the RM cell by the traffic source; and
  2) the delay from the next transmission time of a cell on the forward connection (following the receipt of the RM cell by the traffic source) to the arrival at the interface in question of said cell.

Hence, t_1 is the one-way transfer delay from the source to the interface and t_2 is the round-trip feedback delay between the interface and the source, excluding the residual of the inter-cell interval between successive transmission times.

The delays t_1 and t_2 vary during the course of the session. Let τ_1 be an upper bound on the variation of t_1 and let τ_2 and τ_3 be upper and lower bounds respectively of t_2.

The parameters τ_1, τ_2 and τ_3 are specified at the given interface for the given connection. (Note that for simplicity, τ_3 could be set to zero, though with the consequence of a less tight conformance definition.) The conformance definition in 6.7.5.3 makes use of these parameters, as well as the ACRs determined by backward RM cells on the companion backward connection.

6.7.5.2 Requirements on the ABR conformance definition

The ABR conformance definition must satisfy the following design constraints relative to parameters τ_1, τ_2 and τ_3 as specified for the connection, and delays t_1 and t_2:

1) The conformance definition shall identify each cell as either conforming or non-conforming.
2) The conformance definition shall be testable at an interface.
3) If an MCR > 0 has been negotiated, the conformance definition shall find all CLP = 0 cells that are not BECN cells on a connection conforming if these cells conform to GCRA(mcr^{-1}, τ_1).
4) The conformance definition used at an interface shall find a cell non-conforming only if its arrival time there and those of the preceding conforming cells on the connection could not have resulted from the ideal transmission times of an ABR source, and delays t_1 and t_2 for the connection satisfying $\tau_3 \le t_2 \le \tau_2$ and $\max(t_1) - \min(t_1) \le \tau_1$. In determining whether a cell is conforming, it can be assumed that the inter-cell interval between that cell and the previous cell on the connection:
   i) shall account for feedback conveyed in backward RM cells transmitted across the interface on the backward connection more than τ_2 before that previous cell; and
   ii) shall not account for feedback conveyed in backward RM cells transmitted across the interface on the backward connection less than τ_3 before that previous cell.

6.7.5.3 ABR conformance algorithm

6.7.5.3.1 Dynamic Generic Cell Rate Algorithm (DGCRA) for ABR

The conformance definition is based on the dynamic GCRA. The dynamic GCRA (DGCRA) is an extension of the GCRA defined in Annex A. The DGCRA differs from the GCRA in that the increment T changes with time, as determined by ABR feedback information conveyed on the corresponding backward connection.

The DGCRA checks the conformance of CLP = 0 cells on the ABR connection, excluding the BECN RM cells.

Let T(k) denote the increment that pertains for the kth cell on the connection that is tested by the DGCRA. The tolerance τ_1, which accommodates jitter or bursts, is a constant that does not depend on k.

At the arrival time $t_a(k)$ of the kth cell, the DGCRA first calculates T(k) (see 6.7.5.3.2) and then checks the cell's conformance and updates its own Last Virtual Scheduling Time (LVST) as follows:

Initialize:

LVST = $t_a(1)$, T_{old} = T(1)

At each arrival time $t_a(k)$ of a cell for k ≥ 2:

if: $t_a(k) \geq LVST + \min(T(k), T_{old}) - \tau_1$, # cell is conforming

then set $LVST = \max(t_a(k), LVST + \min(T(k), T_{old}))$

else: # cell is non-conforming

do not update algorithm state.

$T_{old} = T(k)$

In the special case where T(k) = T (a constant) for all k, the above algorithm is equivalent to GCRA(T, τ_1). The term "min(T(k), T_{old})" accounts for the option of the source to reschedule or not reschedule the lead cell queued for transmission when new feedback is received.

The selection of T(k) depends on two additional delay parameters τ_2 and τ_3 for the connection. The interval T(k) must satisfy the constraints that:

- T(0) = the reciprocal of the initial value of the ACR;
- $T(k) \geq \frac{1}{PCR}$ for $k \geq 1$, where PCR is the peak cell rate for the connection;
- if $MCR > 0, T(k) \leq \frac{1}{MCR}$ for $k \geq 1$, where MCR is the minimum cell rate for the connection;
- T(k) ≤ 1 s for k ≥ 1

The sequence {T(k), k ≥ 1} of increments, which are successively used at arrival times {$t_a(k)$, k ≥ 1} of cells at the interface, depends on the feedback information in the backward RM cell sent across the interface at departure times {$t_b(j)$, j ≥ 1} on the backward connection (see 6.7.5.3.2). Each backward RM cell determines an allowed cell rate that could apply to some future cells on the forward direction.

Taking into account other events pertinent to the connection is currently not specified.

Note that it is possible that this rate never actually applies to any cells on the forward direction as no cells may be transmitted in the interval when the conformance definition would be using this rate.

Thus, we call these computed rates "Potential Allowed Cell Rates" (PACRs). Let PACR(j) be the potential allowed cell rate as determined at the interface by the backward RM cell sent across the interface at departure time $t_b(j)$.

In mode 1 (explicit rate mode), the ECR field is the only field in the pertinent backward RM cells (see 6.7.5.3.2 for a definition of the set of pertinent backward RM cells) that is used in the calculation of T(k). Conformance to mode 1 is specified in this Recommendation.

In mode 2 (binary mode), the determination of T(k) may also make use of QueueLength, CI, and NI fields. Mode 2 is under study and may depend on the further specification of the source reference behaviour.

The DGCRA defers mapping increases in the sequence {PACR(j)} into the increments {T(k)} until after a lag τ_1, and defers mapping decreases in {PACR(j)} into {T(k)} until after a lag τ_2. This accommodates the behaviour of a connection that requires at least a time τ_3 and at most a time $\tau_2 > \tau_3$ to affect the instructed changes in the rate at which cells arrive at the interface.

6.7.5.3.2 Algorithm for determination of T(k) in explicit mode

This clause presents the reference algorithm for determining the sequence of increments $\{T(k), k \geq 1\}$ for the case of explicit rate mode. At the arrival time of the *kth* cell at the interface in the forward direction at time $t_a(k)$, the algorithm first determines the allowed cell rate *ACR(k)*; then *T(k)* is set to the reciprocal of *ACR(k)*. If ACR happens to be computed to a value less than 1 cell/s, *T(k)* is set to 1 s. Thus, even if ACR is computed to be < 1 cell/s, the algorithm may identify all cells conforming if the rate is not higher than 1 cell/s. The increment for the DGCRA at each cell arrival in the forward direction is thus determined.

The time and rate values used in the algorithm are:

– $t_a(k)$ is the arrival time of the *kth* cell in the forward direction;
– *PACR(j)* is the potential allowed cell rate as determined at the interface by the backward RM cell sent across the interface at departure time $t_b(j)$;
– *ACR* is the value of an auxiliary variable ACR computed at $t_a(k)$.

Only the values *PACR(j)* of the ECR field carried in pertinent RM cells are used. Pertinent RM cells are backward RM cells with correct CRC-10 in the EDC field (see 8.1) that are either non-BECN cells, or BECN cells with ECR < $PACR(j-1)$ for $j > 1$.

Initialize: $t_a(0) = 0$.

At each arrival time $t_a(k)$ of a cell for $k \geq 1$,

a) if the set of indices j of backward RM cells such that $0 < t_b(j) \leq t_a(k) - \tau_2$ is non-empty with j_{max}, being the largest element then:

   set $ACR = PACR(j_{max})$,

   otherwise:

   set *ACR* = IACR;

b) if the set of indices j of backward RM cells such that $t_a(k) - \tau_2 < t_b(j) \leq t_a(k) - \tau_3$ is non-empty, then:

   set $ACR = max(ACR, PACR_{max})$, where $PACR_{max}$ is the largest explicit cell rate *PACR(j)* for j in the set;

c) $ACR = max(ACR, MCR); ACR = min(ACR, PCR);$

d) If *ACR* > 1 cell/s, then $T(k) = 1/ACR$; else $T(k) = 1$ s.

Clarification of the algorithm

The rate change induced by a backward RM cell departing from the interface (on the backward connection) at time t_b can be observed on the forward connection at some later time t_a such that $t_b + \tau_3 < t_a < t_b + \tau_2$.

If a cell arrives on the forward connection at time t_a, one can conclude the following:

a) The rate change due to backward cells that passed at times t_b such that $t_a \geq t_b + \tau_2$ will have taken place at the interface by time t_a. If there are such RM cells, the most recent one is selected with the index $j_{\max}$ and the corresponding Potential Allowed Cell Rate (PACR) is used. If no such backward RM cells are present, the initial value for the allowed cell rate is used. This is reflected in part a) of the algorithm.

b) Backward RM cells that passed at times t_b such that $t_b + \tau_3 < t_a < t_b + \tau_2$ may also have an impact on the rate at time t_a. To be on the safe side, the maximum of these rates is selected. Then, select the higher value of this operation and the result from the first step. This is reflected in part b) of the algorithm.

c) The resulting value should be between MCR and PCR.

d) Finally, if ACR > 1 cell/s, then *T(k)* is set to the reciprocal of *ACR*; else *T(k)* is set to 1 second.

NOTE – This algorithm may require a large number of PACR values to be stored in its implementation. It may be approximated with a less complex and less stringent implementation with a lower number of PACR values stored. Appendix VIII proposes an algorithm which limits the number of stored PACRs to 2 and is always less stringent than the ABR conformance algorithm.

QoS aspects

QoS commitments for ABR are in terms of CLR for CLP = 0 cells. There are no commitments on CDV or cell transfer delay.

For a user whose traffic conforms to the conformance definition above, the QoS commitment on cell loss ratio pertains. In particular, the user can transmit at the MCR at any time and still get the QoS commitment.

QoS commitments pertain to user generated cells with CLP = 0 and conformance is also checked on the aggregate user generated cell flow with CLP = 0. However, for user RM cells with CLP = 0, it is recognized that there may be technical limits on the amount of these RM cells that can be processed by a given network element.

ABR capability without using a conformance definition

The ABR capability may also be used in a mode where the conformance definition in this clause is not used. In this mode, QoS indications assuming source, destination, and network element reference behaviours are provided (under study). In this mode, the network may achieve such QoS indications on cell loss ratios via suitable traffic engineering and operating rules.

Even when the conformance definition of this clause is not used, the network operator still should give feedback in accordance with the defined allocation policy and may perform network specific enforcement of resource allocations.

6.8 Guaranteed frame rate transfer capability (GFR)

Some users have traffic characteristics that make it difficult to determine traffic parameters required by existing ATM transfer capabilities. Often such users are also not able to react to explicit feedback from the ATM network. In addition, user data is often organized in frames and may tolerate loss. For sources of such user data it may be sufficient to get a low cell loss commitment that applies to a minimum cell rate and to expect that some of the frames in excess of the minimum cell rate are delivered. To support such traffic in an ATM network, an ATM transfer capability is defined which is termed Guaranteed Frame Rate (GFR).

6.8.1 Definition and service model

The Guaranteed Frame Rate (GFR) ATM transfer capability is intended to support non-real-time applications. The GFR ATC requires that the user data cells are organized in the form of frames that can be delineated at the ATM layer. The GFR applies to ATM connections that delineate frames using the AUU indication. Any other delineation method, e.g., the use of RM cells, is for further study. The GFR ATC only applies to VCCs, because the AUU indication is not a reliable frame delineation at the VP sublayer.

In the GFR ATC, the user can send a frame either unmarked or marked. By marking a frame, the user indicates that the frame is of lesser importance than an unmarked frame on that particular GFR connection. An unmarked frame has all cells of the frame set to CLP = 0, a marked frame has all cells in the frame set to CLP = 1. Frames sent by the user should have all cells with the same setting of the CLP bit. QoS commitments do not apply to cells in frames with mixed setting of the CLP bit nor to cells in marked frames.

The GFR ATC uses a Minimum Cell Rate (MCR), in conjunction with a given Maximum Frame Size (MFS) and a given Maximum Burst Size (MBS). The MFS and MBS are both expressed in cells. QoS commitments apply only if MCR is greater than zero. It is a network option whether the GFR ATC supports MCRs greater than zero.

In addition to MCR, MBS and MFS, a PCR for the user generated CLP = 0 + 1 cells is defined for the GFR ATC. PCR is always larger than MCR.

The following two examples describe the commitment that the user of a GFR connection will receive:

– If MCR > 0 and if the user sends unmarked frames that do not exceed the Maximum Frame Size and the user sends them at a constant rate that is less than or equal to the MCR, then the commitment is that all these frames are delivered across the network according to the QoS class.
– If MCR > 0 and if the user sends unmarked frames that do not exceed the Maximum Frame Size and the user has not sent cells for a long time and the user sends them in a burst with a length that does not exceed the Maximum Burst Size and at a rate that does not exceed PCR, then the commitment is that all these frames are delivered across the network according to the QoS class.

The GFR ATC also allows the user to send in excess of the negotiated MCR, but traffic that exceeds MCR will only be delivered within the limits of available resources.

The GFR ATC does not provide explicit ATM layer feedback to the source regarding the current level of network congestion. Instead, the level of congestion is derived by the higher layer protocols from the delivery or discard of frames of the connection. VC RM cells on a VCC are not used to operate GFR; however, such cells that would still be present on the connection are considered as part of the user data cell flow. Support of OAM cells within a GFR connection is not precluded. See Appendix XIV on OAM support for GFR connections.

The service model distinguishes between frames of which all cells are conforming and frames of which not all cells are conforming. GFR cell conformance is defined in 6.8.3.1.

– For a frame of which all cells are conforming, the network shall attempt to either deliver all the cells or deliver none of the frame's cells. However, if the network delivers only part of such a frame it shall attempt to deliver also the last cell of that frame. For the frames of which all cells are conforming, the ratio of the number of cells in partially delivered frames to the number of cells in all these frames should not be higher than the number of cells in MFS times the committed CLR of QoS class 2, irrespective of the QoS class with which the connection is associated.

NOTE – Australia maintains its technical reservation on the GFR service model. The reservation reflects Australia's position that GFR delivers partial frames with a probability not higher than the CLR associated with QoS class 1.

– For a frame of which not all cells are conforming, there are no commitments or expectations on the network's delivery of the frame. However, if the network delivers part of such a frame, it should attempt to deliver also the last cell of that frame.

There are two versions of GFR, GFR1 and GFR2. They differ with respect to the treatment of the CLP bit of non-conforming frames:

– GFR1: The network conveys the CLP bit transparently. Tagging is not applicable.

– GFR2: The network may apply frame tagging, by tagging all cells of a frame that does not pass the F-GCRA frame test (see 6.8.3.2).

6.8.2 Source traffic descriptor and CDV tolerance

The user and the network agree upon a source traffic descriptor with the following traffic parameters:

– a Peak Cell Rate PCR(0+1) for the user generated CLP = 0 + 1 cells and the associated CDV tolerance $\tau_{PCR}(0+1)$;

– a Minimum Cell Rate MCR(0) for user generated CLP = 0 cells and the associated CDV tolerance $\tau_{MCR}(0)$. If MCR > 0, then it has location, basic events, and coding identical to the PCR (see 5.4.1);

– a Maximum Frame Size MFS(0+1) expressed in cells;

– a Maximum Burst Size MBS(0) for user generated CLP = 0 cells expressed in cells. MBS should be greater than or equal to MFS.

All the above values can either be conveyed via signalling or assigned on a per subscription basis.

When selecting values for the traffic descriptor and CDV tolerances, the phenomenon described in 6.8.3.3 under "Parameter selection for minimum throughput" is to be taken into account.

6.8.3 Conformance definition and QoS commitments

6.8.3.1 Cell conformance

A GFR user generated cell is conforming if all of the following three conditions are met:

– The cell conforms to the GCRA(1/PCR, τ_{PCR}) test for CLP = 0 + 1 cells.

– The cell is either the last cell of the frame or the number of cells in this frame up to and including this cell is less than MFS.

– The CLP bit of the cell has the same value as the CLP bit of the first cell of the frame.

The GCRA test is applied to every cell and the GCRA is updated (incrementing by T = 1/PCR) when the cell conforms to the GCRA(1/PCR, τ_{PCR}) test.

See 6.8.3.4 regarding UPC/NPC actions.

6.8.3.2 Frame conformance and F-GCRA(T,τ)

This clause defines frame conformance for GFR1 and GFR2. A frame is conforming if all cells of the frame are conforming (see 6.8.3.1) and if the frame passes the Frame based Generic Cell Rate Algorithm F-GCRA as described below.

The F-GCRA uses the negotiated value of a cell rate 1/T, assuming that a tolerance τ is allowed.

The variables of the F-GCRA are as follows:

– t_a denotes the arrival time of the latest cell at a standardized interface;

– *X* denotes the value of the leaky bucket counter, as in the continuous-state leaky bucket algorithm;

– *LIT* denotes the Last Incrementing Time;

– *X_1* and *LIT_1* denote the values of the parameters *X* and *LIT* at the end of the last frame whose first cell was a CLP = 0 cell. The parameters *LIT_1* and *X_1* are used so that the F-GCRA is not updated for a CLP = 0 frame of which all cells are conforming but that did not pass the frame test. It is updated for all other frames that start with a CLP = 0 cell;

– *Frame_test_passed* denotes a connection specific variable that stores the frame test result;

– *Frame_tagging* denotes a connection specific variable that is only used in GFR2. It stores the frame tagging status. If frame tagging for GFR is implemented, then this status information could be used to change the CLP bit from 0 to 1;

– *X'* is an auxiliary variable.

Initialization of the F-GCRA variables:

– At the time of arrival t_a of the first cell of the connection to cross the given interface, $X = X\_1 = 0$ and $LIT = LIT\_1 = t_a$.

– The initial values of frame_test_passed and frame_tagging are irrelevant.

The F-GCRA is defined as follows:

The algorithm below has three parts. Part 1 is executed before part 2, and part 2 is executed before part 3.

Part 1: At the arrival of the *first* cell of a frame at a given interface T_B or inter-network interface, on the ATM connection.

GFR1

```
if (CLP = 1)
    then frame_test_passed = false
    else
        X' = X – (t_a – LIT)
        if (X' > τ)
            then frame_test_passed = false
            else frame_test_passed = true
```

GFR2

```
if (CLP = 1)
    then frame_test_passed = false;
        frame_tagging = false
    else
        X' = X – (t_a – LIT)
        if (X' > τ)
            then frame_test_passed = false;
                frame_tagging = true
            else frame_test_passed = true;
                frame_tagging = false
```

Part 2: At the arrival of *each* cell of a frame whose first cell was a CLP = 0 cell.

GFR1 and GFR2

```
X' = X – (t_a – LIT)
X = max(0, X') + T
LIT = t_a
```

Part 3: At the arrival of the *last* cell of a frame whose first cell was a CLP = 0 cell.

GFR1 and GFR2

```
if (frame contained a non-conforming cell) or (frame_test_passed = true)
    then   X_1 = X; LIT_1 = LIT
    else   X = X_1; LIT = LIT_1
```

NOTE – The reader is referred to Appendix XI for an algorithm called CF-GCRA. This algorithm is less exact than F-GCRA in testing frame conformance, but could form the basis of simple implementations, provided tolerances are set to sufficiently large values.

6.8.3.3 QoS commitments

QoS commitments are the same for GFR1 and GFR2.

QoS commitments apply only to connections with MCR greater than zero. Therefore it is assumed that MCR is greater than zero in the following. There are no commitments on CDV or cell transfer delay.

The GFR ATC provides a QoS commitment in terms of a cell loss ratio according to the associated QoS class for the number of cells in conforming frames (see 6.8.3.2), where at a standardized interface the F-GCRA(T,τ) is applied with the parameters $T = 1/\text{MCR}$ and $\tau = \tau_{IBT} + \tau_{MCR}$, and where $\tau_{IBT} = (\text{MBS} - 1)\cdot(1/\text{MCR} - 1/\text{PCR})$.

Additional procedural commitments

In addition to the QoS commitments, the GFR ATC includes the procedural commitment that, when there are sufficient resources available, some CLP = 0 frames of which all cells are conforming yet are failing the F-GCRA test and some CLP = 1 frames of which all cells are conforming will be delivered. To deliver these frames in excess of the QoS commitments, a network-specific policy is applied to allocate a share of the available resources to each GFR connection involved. Network-specific policies are not subject to standardization. In such a network policy, the network could, for example, take the CLP status of the frames into account by discarding CLP = 1 frames in preference to CLP = 0 frames on that particular GFR connection.

There are no commitments for frames of which not all cells are conforming and the network is allowed to discard any cell of these frames. However, if the network delivers part of such a frame, it should attempt to deliver also the last cell of that frame. If some cells on a GFR connection are non-conforming, then the network may consider the GFR connection as non-compliant, see 5.3.2.

Parameter selection for minimum throughput

The F-GCRA may show a phenomenon similar to the phenomenon for the GCRA as described in Appendix III. Under certain conditions, and when CLP = 0 frames of which all cells are conforming, arrive at the F-GCRA(1/MCR, $\tau_{IBT} + \tau_{MCR}$) with a cell rate greater than MCR, the cell rate of *conforming frames* can be less than MCR. It can be shown that this phenomenon is not present if $\tau_{IBT} + \tau_{MCR} \geq \text{MFS/MCR}$.

6.8.3.4 UPC/NPC actions

During the connection lifetime, cell conformance may be continuously checked within the network by static UPC/NPC mechanisms, given such UPC/NPC mechanisms are present (see 7.2.3). The conformance definition does not imply any particular implementation of the UPC/NPC.

For a frame of which not all cells are conforming, the network is allowed to discard any cell of the frame, e.g., to discard isolated cells or to perform frame tail discard. For a frame of which all cells except the last one are conforming, it may be desirable to retain that last cell and to update the GCRA, even if that cell did not pass the GCRA(1/PCR,τ_{PCR}) test.

7 Functions for traffic control and congestion control

7.1 Introduction

Generic traffic control and congestion control functions are defined as the set of actions respectively taken by the network in all the relevant network elements to avoid congestion conditions or to minimize congestion effects and to avoid the congestion state spreading once congestion has occurred.

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

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

7.1.1 Traffic control and congestion control functions

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 Recommendation.

Traffic control functions

i) Network resource management (7.2.1).

ii) Connection admission control (7.2.2).

iii) Usage/network parameter control (7.2.3).

iv) Discard Priority control (7.2.4).

v) Frame discard (7.2.5).

vi) Scheduling control (7.2.6).

vii) Traffic shaping (7.2.7).

viii) Fast resource management (7.2.8).

Congestion control functions

ix) Discard Priority Control (7.3.1).

x) Explicit forward congestion indication (7.3.2).

xi) Frame discard (7.3.4).

xii) Scheduling control (7.3.5).

Additional control functions may be used. Possible useful functions that require further study to determine details are:

xiii) Connection admission control that reacts to and takes account of the measured load on the network.

xiv) Variation of usage monitored parameters by the network. For example, reduction of the amount of network resources (e.g., allowed cell rate) made available to the user (e.g., as in ABR or using GFC as described in ITU-T Rec. I.361).

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

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

Different levels of QoS commitments may be provided on ATM connections by proper routing, traffic shaping, discard priority control and resource allocation.

7.2 Traffic control functions

7.2.1 Use of virtual paths for network resource management

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/NPC can be applied to the traffic aggregate;
– aggregate network capabilities such that the NPC 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 operators' decision.

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 CRF(VC)s (connection-related functions at the VC level, e.g., a VC multiplexer, a VC switch, see Figure 16).

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.

Conversely, when VCCs within a VPC require a range of QoS, the VPC performance objective should be set suitably for the most demanding VCC carried.

Combining common routing and priority control may be used by connection admission control for services requiring a number of VCCs with low differential delays and different cell loss ratios (e.g., multimedia services).

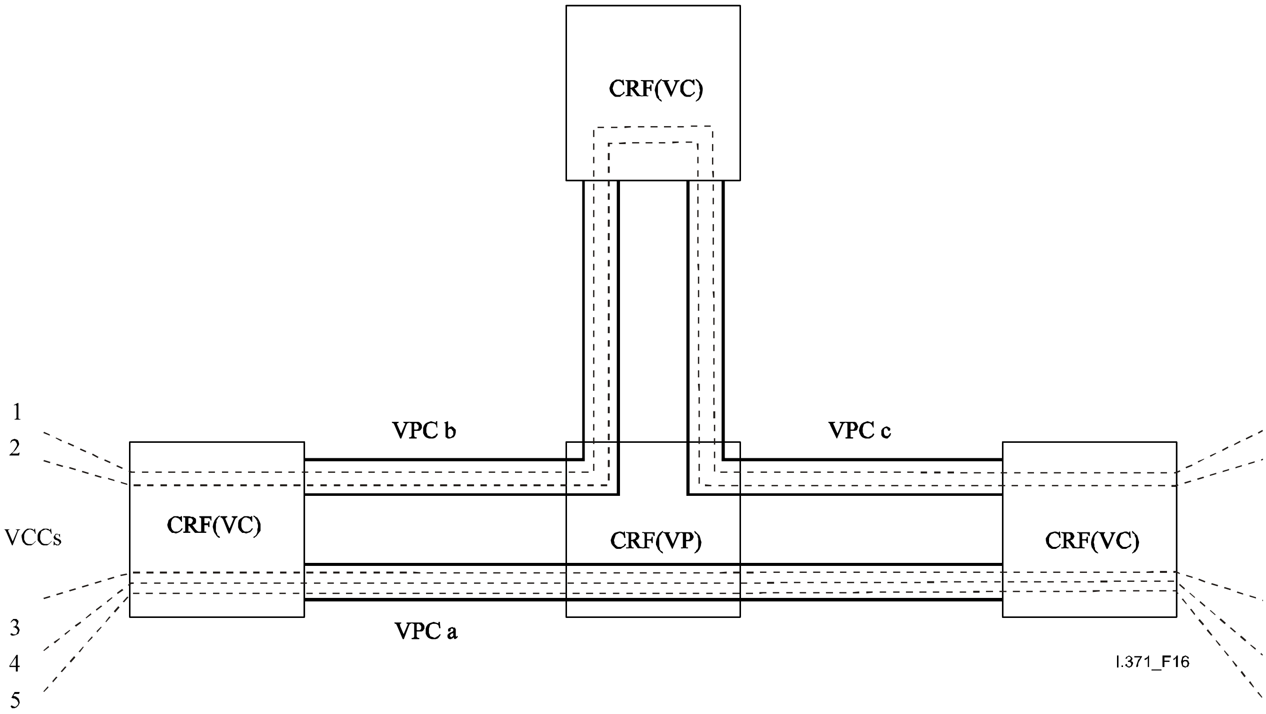

CRF(VC) Connection-related functions at the VC level

NOTE 1 – 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 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.

NOTE 2 – VCCs 3, 4 and 5 experience similar network performances in terms of cell transfer delay and cell delay variation if handled similarly by CRF(VC)s, whilst possibly providing for two different cell loss ratios by using the CLP bit.

NOTE 3 – On a user-to-user VPC, the QoS experienced by individual VCCs depends on CEQ traffic handling capabilities.

Figure 16/I.371 – Mapping cell loss ratios for virtual channel connections and virtual path connections

On the basis of the applications of VPCs contained in ITU-T Rec. I.311, namely:

Case A User-user application: the VPC extends between a pair of T_B reference points.

Case B User-network application: the VPC extends between a T_B reference point and a network node.

Case 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 cases B and C: the network is aware of the QoS of the VCCs carried within the VPC and has to accommodate them. However, the setting of CDV tolerances requires further study.

Statistical multiplexing of VC links within a VPC where the sum of the PCRs of all VC links may exceed the virtual path connection PCR is only possible when all virtual channel links within the virtual path connection can tolerate the QoS that results from statistical multiplexing.

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.

7.2.2 Connection admission control

Connection admission control is defined as the set of actions taken by the network at the call establishment phase (or during call renegotiation phase) in order to establish whether a virtual channel connection or a virtual path connection can be accepted or rejected.

In a B-ISDN environment, a call can require more than one connection (e.g., for multimedia or multiparty services such as videotelephony or videoconferencing). In this case, connection admission control procedures should be performed for each virtual channel connection or virtual path connection.

The user will negotiate the traffic characteristics of the ATM connections with the network at connection establishment via signalling or network management procedures. These characteristics may be renegotiated during the lifetime of the call at the request of the user, using either signalling or network management procedures. The network may limit the frequency of such renegotiations. This is outside the scope of this Recommendation.

In the case of permanent or reserved service (e.g., using a permanent virtual path connection or a permanent virtual channel connection), the traffic characteristics are indicated or renegotiated with an appropriate procedure, either off-line (e.g., at subscription) or on-line via management.

Depending on the selected ATC (see clause 6), dynamic modifications of traffic characteristics may also use ATM layer resource management procedures (see 7.2.8).

On the basis of connection admission control in an ATM network, a connection request for a given call is accepted only when sufficient resources are available to establish the connection through the whole network, to comply with the required quality of service (QoS) and to maintain the agreed QoS of existing connections. This applies as well to renegotiation of connection parameters within a given call.

Different strategies of network resource allocation may be applied for CLP = 0 and CLP = 1 traffic flows. 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 whilst still meeting the performance objectives.

Resource allocation schemes are left to network operators' decision.

The connection establishment procedures will enable CAC to derive at least the following information (see traffic contract, 5.3):

– required ATM layer transfer capability;
– source traffic descriptors;
– CDV tolerances;
– required QoS class.

Connection admission control makes use of this information to determine:

– whether the connection can be accepted or not;
– traffic parameters needed by Usage/Network Parameter Control (UPC/NPC);
– path selection and allocation of network resources.

For a given ATM connection, a user indicates an ATM transfer capability from the transfer capabilities supported by the network. The corresponding standardized traffic parameters included in the source traffic descriptor in addition to the peak cell rate should enable the network operator to implement a connection admission control policy that attains a statistical multiplexing gain, as compared with a CAC policy that statically allocates resources on the sole basis of a peak cell rate to the connections, whilst still meeting QoS commitments to the connection. Many such CAC policies are possible, and the choice is at the discretion of the network operator.

For a single ATM connection, a user indicates a QoS class from the QoS classes which the network supports. Priority control using the CLP bit allows at most to specify two different cell loss ratio objectives for an ATM connection (see 5.3.3 and 7.2.4). The role of priority control in connection admission control is outside the scope of this Recommendation. Delay sensitivity is part of the required QoS. Specific QoS classes are within the scope of ITU-T Rec. I.356.

7.2.3 Usage parameter control and network parameter control

Usage Parameter Control (UPC) and Network Parameter Control (NPC) perform similar functionalities at different interfaces: the UPC function is performed at the user-network interface, whereas the NPC function is performed at the Inter-Network Interface (INI).

The use of a UPC function is recommended, and the use of an NPC function is a network option. Whether or not the operator chooses to use the NPC function, the network-edge-to-network-edge performance objectives still need to be met if the connection complies with the traffic contract (see 5.3.2).

7.2.3.1 UPC/NPC functions

Usage/network parameter control is defined as the set of actions taken by the network to monitor and control that the traffic contract is respected in terms of traffic offered and validity of the ATM connection, at the user access and the network access, respectively. Their main purpose is to protect network resources from malicious as well as unintentional misbehaviour which can affect the QoS of other already established connections by detecting violations of negotiated parameters and procedures and taking appropriate actions.

Connection monitoring encompasses all connections crossing the UNI or inter-network interface. Usage parameter control and network parameter control apply to user VCCs/VPCs, signalling and meta-signalling virtual channels. Additional cell flows may be inserted by the network for its own purposes (e.g., segment OAM cell flows, RM cell flows) and may pass through a UPC or NPC. When such flows are inserted, the network operator must not allow the inserted cell flow to disrupt the QoS commitments it has made to the user's connection. This is not covered by this Recommendation.

The monitoring task for usage parameter control and network 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 assigned), and checking whether the traffic entering the network from active VCCs is such that parameters agreed upon are not violated; this action should be performed at the ingress of a CRF(VC), before any ATM layer multiplexing or switching occurs at the VC sub-layer.

   NOTE – The CDV due to multiplexing at the VP sub-layer needs to be taken into account.

2) Checking the validity of VPI (i.e., whether or not VPI values are assigned), and checking whether the traffic entering the network from active VPCs is such that parameters agreed upon are not violated; this action should be performed at the ingress of a CRF(VP), before any ATM layer multiplexing or switching occurs.

7.2.3.2 UPC/NPC requirements

The need for and the definition of a standardized UPC/NPC algorithm requires further study. A number of desirable features of the UPC/NPC algorithm can be identified as follows:

– capability of detecting any illegal traffic situation;
– selectivity over the range of checked parameters (i.e., the algorithm could determine whether the user behaviour is within an acceptance region);

– rapid response time to parameter violations;
– simplicity of implementation.

There are two sets of requirements relating to the UPC/NPC:
– those which relate to the quality of service impairments the UPC/NPC might directly cause to the user cell flow;
– those which relate to the resource the operator should allocate to a given VPC/VCC and the way the network intends to protect those resources against misbehaviour from the user side or another network (due to fault conditions or maliciousness).

Two performance parameters have been identified. They have to be considered when assessing the performances of UPC/NPC mechanisms. Methods for evaluating UPC/NPC performance and the need to standardize these methods are outside the scope of this Recommendation.
– Response time: the time to detect a given situation that involves non-conforming cells on a VPC/VCC under given reference conditions.
– Transparency: for the same set of reference conditions, the accuracy with which the UPC/NPC initiates appropriate control actions on a cell stream in which some cells are non-conforming and avoids inappropriate control actions on a stream of conforming cells.

A specific UPC/NPC mechanism may commit errors by taking excessive policing actions, i.e., declaring a volume of non-conforming cells larger than the volume of cells that do not conform with the traffic contract. It can also fail to take sufficient policing actions on a cell stream in which some cells are non-conforming.

Excessive actions of the UPC/NPC on any connection are part of the overall network performance degradation and should remain of a very low probability. Quantification of this probability is within the scope of ITU-T Rec. I.356. Safety margins may be provisioned depending upon the UPC/NPC algorithm to limit the degradation introduced by the UPC/NPC.

Policing actions performed on the excess traffic in case of traffic contract violation are not to be included in the network performance degradation allocated to the UPC/NPC.

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

Since cell sequence integrity is maintained on any ATM connection, the UPC/NPC including its optional cell tagging action (see 7.2.3.6) must operate as a single server using First-In-First-Out (FIFO) service discipline for each ATM connection.

7.2.3.2.1 Performance of cell level UPC/NPC

A method to determine the ratio of non-conforming cells to a negotiated cell rate at a given interface is defined in ITU-T Rec. I.356. A 1-point measurement process computes the ratio γ_M between the number of cells exceeding the traffic contract and the total number of submitted cells.

The following applies when a single GCRA is used for conformance definition. Other cases are under study (see ITU-T Rec. I.356).

An ideal UPC/NPC implementing the 1-point measurement process on a cell flow would just take policing actions on cells so that the ratio γ_P of the number of cells on which enforcement action is taken (tagging or discarding) over the number of processed cells approximates γ_M. Although the ideal UPC/NPC allows for a cell-based decision, it is not possible to predict which particular cells of a connection will suffer from the policing action from this ideal UPC/NPC. This is because of measurement phasing, that is the decisions taken by the measurement process depend on which cell the process begins with and on the initial values of the state variables of the process. The measurement of γ_M could provide an estimate of the degree of non-conformance of the user to the negotiated rate. The measurement of cell non-conformance ratio γ_M may be used for arbitration between user and network at the UNI and between two network portions at the INI in case of conflict. Such measurements may be carried out on-demand, either in-service on connections established prior or after the on-demand request, or out-of-service on any flow that emulates the normal operation of the user.

The transparency of a UPC/NPC 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/NPC is taking less policing action than a measurement process would do. A negative difference means that policing actions are unduly taken by the UPC/NPC.

The exact way of measuring the transparency of a given mechanism for cell level UPC/NPC and its dependence on time requires further study.

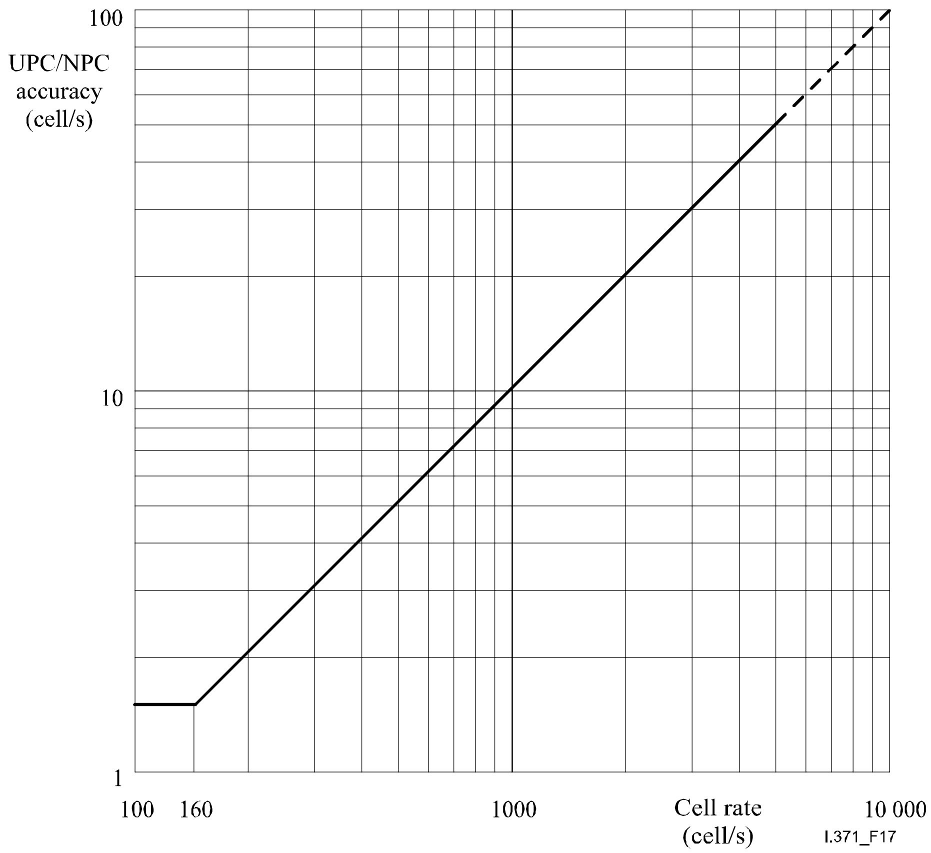

Figure 17/I.371 – UPC/NPC accuracy requirement

One UPC/NPC requirement relating to accuracy is the following: for cell rate control (either peak or sustainable cell rate), the UPC/NPC should be capable of coding a cell rate at most $\delta = 1\%$ larger than the cell rate used in the cell conformance definition. This requirement pertains for cell rates as low as 160 cell/s. For cell rates between 100 cell/s and 160 cell/s, the accuracy is 1.6 cell/s (which is 1% of 160 cell/s) (Figure 17). UPC/NPC accuracy for cell rates in the range of 1 cell/s to 100 cell/s is under study.

The above performance requirement is a requirement on the capability of the UPC/NPC. A network operator is not required to set the parameters of the UPC/NPC to be within the margin given by δ.

The relationship between these accuracy requirements and performance monitoring is addressed in Appendix III.

7.2.3.3 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 can be identified as shown in Figure 18:

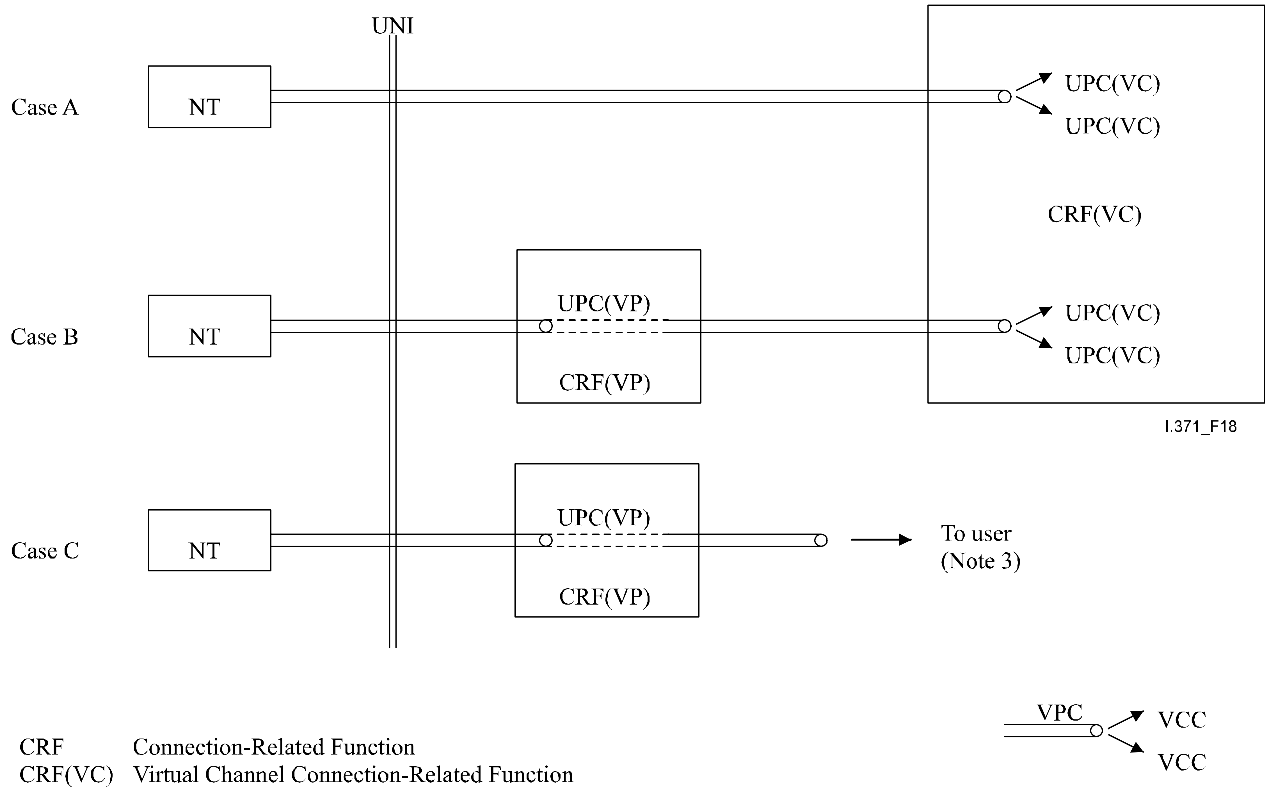

CRF Connection-Related Function
CRF(VC) Virtual Channel Connection-Related Function
CRF(VP) Virtual Path Connection-Related Function
NT Network Termination
UPC Usage Parameter Control

NOTE 1 – In cases B and C, the VPI value does identify a negotiated VPC.
NOTE 2 – In case B, CRF(VP) and CRF(VC) may belong to different network operators.
NOTE 3 – A user may be an upper layer server within the network.

Figure 18/I.371 – Location of the usage parameter control functions

In the following cases, CRF(VC) stands for Virtual Channel Connection-Related Function, and CRF(VP) stands for Virtual Path Connection-Related Function. A CRF(VC) or a CRF(VP) may respectively be a VC or VP concentrator.

Case A (Figure 18): User connected directly to CRF(VC)

Usage parameter control is performed within the CRF(VC) on VCCs (action 1, 7.2.3.1).

Case B (Figure 18): User connected to CRF(VC) via CRF(VP)

Usage parameter control is performed within the CRF(VP) on VPCs only (action 2, 7.2.3.1) and within the CRF(VC) on VCCs only (action 1, 7.2.3.1).

Case C (Figure 18): User connected to user or to another network provider via CRF(VP)

Usage parameter control is performed within the CRF(VP) on VPCs only (action 2, 7.2.3.1).

In Case B, the user may negotiate and need to respect two traffic contracts:

– at the VP level with the CRF(VP) (when the user requests the operator of the CRF(VP) for a VP connection);

– at the VC level with the CRF(VC) (when the user requests the operator of the CRF(VC) for a VC connection).

Implication of this is under study.

NOTE – In case B, the traffic characteristics that are controlled in the UPC(VC) depend not only on those that are negotiated by the user, but also on cell transfer characteristics within the Virtual Path Connection-Related Functions CRFs(VP). The modification of the traffic characteristics is not the user's responsibility.

7.2.3.4 NPC location

Network parameter control is performed on VCCs or VPCs at the point where they are first processed in a network after having crossed an inter-network interface. Three possibilities can be identified as shown in Figure 19.

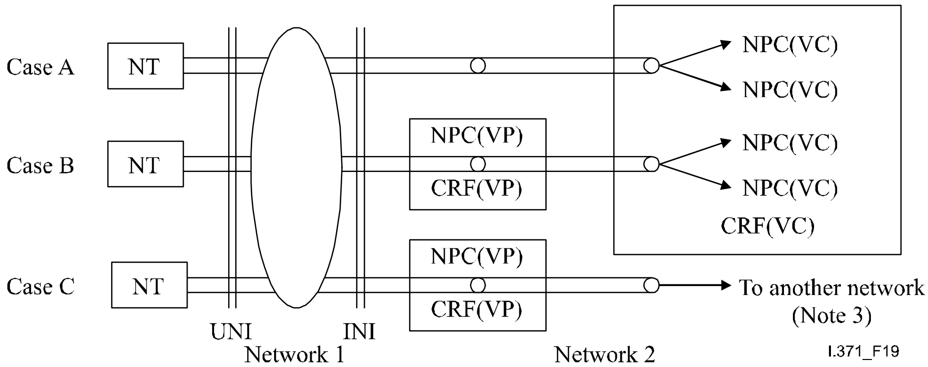

CRF Connection-Related Function
CRF(VC) Virtual Channel Connection-Related Function
CRF(VP) Virtual Path Connection-Related Function
INI Inter-Network Termination
NPC Network Parameter Control
UNI User-Network interface

NOTE 1 – In case A, the VPI value does not identify a negotiated VPC.
NOTE 2 – In cases B and C, the VPI value does not identify a negotiated VPC.
NOTE 3 – Another network may be a user.

Figure 19/I.371 – Location of the network parameter control functions

In the following cases, CRF(VC) [resp. CRF(VP)] stands for virtual channel connection-(resp. virtual path connection) related functions.

Case A (Figure 19): Originating network connected directly to CRF(VC)

NPC is performed within the CRF(VC) (action 1, 7.2.3.1).

Case B (Figure 19): Originating network connected to the CRF(VC) via the CRF(VP)

NPC is performed within the CRF(VP) on VPCs only (action 2, 7.2.3.1) and within the CRF(VC) on VCCs only (action 1, 7.2.3.1).

Case C (Figure 19): Originating network connected to another network via CRF(VP)

NPC is performed within the CRF(VP) on VPCs only (action 2, 7.2.3.1).

In Case B, the user may negotiate and need to respect two traffic contracts:

– at the VP level with the CRF(VP) [when the user requests the operator of the CRF(VP) for a VP connection];

– at the VC level with the CRF(VC) [when the user requests the operator of the CRF(VC) for a VC connection].

Implication of this is under study.

NOTE – In case B, the traffic characteristics that are controlled in the NPC(VC) depend not only on those that are negotiated by the user, but also on cell transfer characteristics within the Virtual Path Connection-Related Functions CRFs(VP). The modification of the traffic characteristics is not the user's responsibility.

7.2.3.5 Traffic parameters subject to control at the UPC/NPC

For each ATM transfer capability, traffic parameters which may be subject to control are those included in the source traffic descriptor and possibly other dynamic parameters specific to a given ATC (see clauses 5 and 6). Whether all these parameters or a subset are subject to control depends upon CAC, ATC and UPC/NPC mechanism. The peak cell rate should never be exceeded by any connection.

7.2.3.6 UPC/NPC actions

The UPC/NPC is intended to control the traffic offered by an ATM connection to enforce the negotiated traffic contract. The objective is that a user will never be able to exceed the traffic contract beyond a level of non-conformance which is operator specific.

The specific control actions to be taken depend on the access network configuration and on the negotiated ATM transfer capability. If this capability allows the renegotiation of traffic parameters by fast resource management function using RM cells, the UPC/NPC mechanism should be dynamic, i.e., be able to dynamically modify its parameters using the information carried by the RM cells. UPC/NPC and traffic shaping functions may be combined in some implementations, in which case cell re-scheduling actions will result from this association.

UPC/NPC actions at the cell level

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

a) cell passing;

b) cell tagging for specific ATCs (see 5.3.4). Cell tagging operates on CLP = 0 cells only, by overwriting the CLP bit to 1. Network elements shall not change the value of the CLP bit, except possibly if the ATM transfer capability (SBR3, GFR2) specifies that tagging applies;

c) cell discarding.

Cell passing is performed on cells that are identified by a UPC/NPC as conforming. Cell tagging and cell discarding are performed on cells that are identified by a UPC/NPC as non-conforming to at least one element of the traffic contract.

Additional UPC/NPC actions

For ABT connections (see 6.6), an action of the UPC/NPC controlling block level conformance may be:

– Initiating a modification of the resources allocated to the connection.

– Discarding all the remaining cells in the ATM block (frame discard, see 7.2.5).

Besides the above actions at the cell level and at the ATM block level, as an option, the UPC/NPC may initiate the release of the connection.

7.2.3.7 Relationship between UPC/NPC, cell loss priority and network performance

A cell identified as non-conforming by the UPC/NPC function performed on the aggregate CLP = 0 + 1 flow is discarded.

When cell tagging (see 7.2.3.6) applies to an ATM connection, CLP = 0 cells identified by the UPC/NPC function performed on CLP = 0 flow as non-conforming are converted to CLP = 1 cells and merged with the user-submitted CLP = 1 traffic flow before the CLP = 0 + 1 traffic flow is checked. If cell tagging (see 7.2.3.6) is not applied to a connection, then cells that are identified by the UPC/NPC as non-conforming to at least one element of the traffic contract are discarded.

When the CLP capability is used by an ATM connection (see SBR configurations 2 and 3, 6.5.3) and some cells do not conform to the CLP = 0 + 1 traffic parameter and corresponding cell delay variation tolerance, the UPC/NPC function performed on the aggregate flow may discard CLP = 0 cells which were not considered in excess by the UPC/NPC function performed on the CLP = 0 cell stream. This is not part of the network performance degradation.

7.2.3.8 ATM layer management functions associated with traffic control

Some examples of ATM layer management functions associated with traffic control functions are given below.

In the case that the UPC/NPC encounters levels of non-conformance beyond a network operator's specific threshold, then indications may be produced by traffic control. These indications may initiate other enforcement actions such as:

– control actions at short time-scales, e.g., indication of an excessive level of non-conformance to the user;
– initiating a renegotiation of the amount of resources allocated to the connection (see 7.2.3.6);
– connection release.

Indications due to non-conformance at standardized interfaces detected by UPC/NPC should not propagate through the network.

With the exception of the renegotiation attempt, the functions addressed above are currently not specified in this Recommendation.

7.2.4 Discard priority control

In the context of traffic control, discard priority control mechanisms may be applied within a network element, primarily in order to protect, as long as possible, traffic flows for which the network has negotiated QoS commitments.

Selective cell discard is a cell discard priority control function which consists of discarding CLP = 1 cells (which are either submitted as lower priority CLP = 1 cells by the user or tagged (see 7.2.3.6) by UPC/NPC) instead of discarding CLP = 0 cells of connections for which QoS commitments pertain.

Selective cell discard may be performed by network elements while still meeting QoS commitments.

The applicability of selective cell discard depends on the ATM transfer capability (see clause 6).

Another form of discard priority control is to discard cells from a connection with QoS class U in order not to discard cells on a different connection with QoS class 1.

7.2.5 Frame discard

If a network needs to discard cells in order to avoid getting into a congested state, it may be desirable to discard consecutive cells on a given connection. This is especially true for connections supporting applications where the information is organized in frames and where each frame results in more than one ATM cell. For such applications, a cell loss results in a corrupted frame that may have to be retransmitted.

Frame cell sequence

A frame cell sequence is a sequence of user generated cells on a given connection. Two methods are currently recognized in this Recommendation to delineate a frame cell sequence at the ATM layer in order to perform frame discard: AUU based and RM-cell based frame cell sequence delineation.

If the user wishes to take advantage of frame discard mechanisms at the ATM layer, the user should ensure that the frame delineation at the ATM layer (i.e., through AUU indication or RM cells), corresponds to frame delineation at a higher layer.

AUU based frame cell sequence delineation

For a VCC, the AUU indication may be used (e.g., as specified for AAL 5) to define a frame cell sequence as follows:

– A frame cell sequence starts with the first user generated cell on the connection or with a user generated cell following a cell with the AUU indication set.

– A frame cell sequence terminates with a user cell that has the AUU indication set.

NOTE 1 – The applicability of AUU based frame delineation on a VPC is currently not specified and requires further study.

NOTE 2 – If the last cell transmitted on a connection before connection termination does not have the AUU indication set, then the corresponding frame cell sequence has not been properly terminated.

RM-cell based frame cell sequence delineation

For a VPC or VCC with the ABT/IT ATC, a frame cell sequence is an ATM block as defined in 6.6. The RM cells that delineate the frame cell sequence (see 6.6.2.4.1) are not part of the frame cell sequence.

Whole frame discard

The function "whole frame discard" is defined as discarding, on the given ATM connection, each user generated cell in a frame cell sequence from the start of the sequence up to and including the user cell terminating the sequence, while preserving frame delineation.

– If AUU-based frame cell sequence delineation is used, the terminating cell is part of the frame cell sequence to be discarded.

– If RM cell-based frame cell sequence delineation is used, the RM cells delineating an ATM block are not part of the frame cell sequence to be discarded.

Frame tail discard

The function "frame tail discard" is defined as, on the given ATM connection, not discarding one or more user generated cells from the start of a frame cell sequence followed by the discard of each user generated cell in a frame cell sequence up to the user generated cell that terminates the sequence, while preserving frame delineation.

– If AUU-based frame delineation is used, the terminating cell, though part of the frame cell sequence, is not discarded to preserve frame cell sequence delineation.

– If RM cell-based frame delineation is used, the RM cells delineating an ATM block are not part of the frame cell sequence to be discarded.

Frame discard application

For connections which have negotiated ABT/IT as the ATC, whole frame discard applies.

Frame discard (whole frame discard and/or frame tail discard) may be applied to connections for which QoS class U has been negotiated, regardless of the ATC.

The application of frame discard for connections with other ATCs or QoS classes is not specified in this Recommendation. Applying frame discard in these cases may lead to the network not meeting the committed QoS. More specifically, the following two cases are noted:

– Application of frame discard to connections for which either QoS class 1 or QoS class 2 has been negotiated. Frame discard as an appropriate UPC/NPC function (see 7.2.3.6) applies only to ABT/IT. Applying frame discard in other cases, for example in case of impending congestion in a network element, may result in the network not meeting the committed QoS.
– Application of frame discard to connections for which QoS class 3 has been negotiated. The application of frame discard to CLP = 1 cells leads to the desired outcome, only if the user marks all cells of any given cell frame sequence consistently either as CLP = 1 or as CLP = 0. In other cases, the network may result in the network not meeting the committed QoS for the CLP = 0 cell flow on the connection.

When QoS commitments are associated with a connection which recognizes the notion of a frame (e.g., ABT), a suitable frame level conformance definition needs to be provided.

NOTE 3 – Performing frame discard on an improperly terminated frame cell sequence is undefined and the corresponding implementation dependent actions may impact the QoS delivered to the connection.

7.2.6 Scheduling control

Cell scheduling mechanisms are implemented in equipment in order to solve contention for transmission. A cell scheduling mechanism determines in which order to transmit cells when several cells are queuing for transmission in a given cell slot time. Cell scheduling mechanisms include:

– time priorities set between different queues;
– weighted fair queuing mechanisms used between different queues.

Scheduling mechanisms may also be applied to achieve a defined allocation policy.

7.2.7 Traffic shaping

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, in order to achieve better network efficiency whilst meeting the QoS objectives or to ensure conformance at a subsequent interface. Traffic shaping must maintain cell sequence integrity on an ATM connection. Shaping modifies traffic characteristics of a cell flow with the consequence of increasing the mean cell transfer delay.

Examples of traffic shaping are peak cell rate reduction, burst length limiting, reduction of CDV by suitably spacing cells in time and queue service schemes.

It is a network operator's choice to determine whether and where traffic shaping is performed. As an example, a network operator may choose to perform traffic shaping in conjunction with suitable UPC/NPC functions.

It is an operator's option to perform traffic shaping on separate or aggregate cell flows.

As a consequence, any ATM connection may be subject to traffic shaping.

The options available to the network operator/service provider are the following:

a) *No shaping*
   – Dimension the network in order to accommodate any flow of conforming cells at the network ingress whilst ensuring conformance at the network egress without any shaping function.

b) *Shaping*

- Dimension and operate the network so that any flow of conforming cells at the ingress is conveyed by the network or network segment whilst meeting QoS objectives and apply output shaping the traffic in order to meet conformance tests at the egress.
- Shape the traffic at the ingress of the network or network segment and allocate resources according to the traffic characteristics achieved by shaping, whilst meeting QoS objectives and subsequent conformance tests at the network or network segment egress.

Traffic shaping may also be used within the customer equipment or at the source in order to ensure that the cells generated by the source or at the UNI are conforming to the negotiated traffic contract relevant to the ATC that is used (see clause 6).

7.2.8 Fast Resource Management (FRM)

Fast resource management functions operate on the time scale of the round trip propagation delay of the ATM connection. Both the ABT (see 6.6) and the ABR (see 6.7) ATM transfer capabilities make use of fast resource management functions to dynamically allocate resources to connections using theses capabilities. Potential other fast resource management functions are for further study.

Fast resource management functions use resource management cells which are described in 8.1.

7.3 Congestion control functions

For some traffic, adaptive rate control facilities at the ATM layer or above may be used. The following congestion control functions at the ATM layer are identified in this Recommendation.

7.3.1 Discard priority control

In the context of congestion control, discard priority control mechanisms as defined in 7.2.4 may be applied within a congested network element, primarily in order to protect, as long as possible, traffic flows for which the network has negotiated QoS commitments.

In particular, selective cell discard of CLP = 1 cells as defined in 7.2.4 may be applied within a congested network element.

7.3.2 Explicit forward congestion indication

The EFCI is a congestion notification mechanism which may be used to assist the network in avoidance of and recovery from a congested state. Since the use of this mechanism by the CEQ is optional, the network operator should not rely on this mechanism to control congestion.

A network element that declares itself in an impending congested state or in a congested state may set an explicit forward congestion indication in the cell header of data cells of relevant connections, so that this indication may be examined by the destination CEQ. For example, the source and destination CEQs may use this indication to implement protocols that adaptively lower the cell rate of the connection during congestion. A network element that does not declare itself in an impending congested state or a congested state will not modify the value of this indication.

A network element may declare itself in a congested or impending congested state under certain conditions. The conditions and the mechanism to detect such states are implementation specific and not subject to standardization. The mechanism by which the congestion indication is used by the higher layer protocols in the CEQ is outside the scope of this Recommendation.

The impact of explicit forward congestion indication on the traffic control and congestion control functions requires further study.

7.3.3 Reaction to UPC/NPC 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 during the call establishment phase. To cope with these situations, specific procedures of the management plane should be designed (e.g., in order to isolate the faulty link).

7.3.4 Frame discard

If a network needs to discard cells in order to recover from a congested state, it may be desirable to discard consecutive cells on a given connection. This is especially true for connections supporting applications where the information is organized in frames and where each frame results in more than one ATM cell. For such applications, a cell loss results in a corrupted frame that may have to be retransmitted.

The definitions and statements regarding the applicability of frame discard to ATM connections in 7.2.5 also apply to frame discard as a congestion control function.

7.3.5 Scheduling control

In the context of congestion control, cell scheduling mechanisms as described in 7.2.6 may be applied within a congested network element, primarily in order to protect, as long as possible, traffic flows for which the network has negotiated QoS commitments.

Moreover, scheduling mechanisms may also be applied to achieve a defined allocation policy.

7.4 Traffic control interworking functions

ATM traffic control functions and procedures are defined according to the objective to integrate services at the ATM layer and achieve network performance objectives compatible with service integration. Traffic functions used by other bearer services may be considered. However, there is no commitment to use such functions for ATM traffic and congestion control.

7.4.1 Traffic control interworking with FMBS

Traffic control functions may be performed at the ingress of each subsequent network according to its specific parameters, whether network or service interworking is considered.

The following reference configurations (Figure 20) apply to traffic control in case of FMBS – B-ISDN interworking. It is left to network operators' decision whether these traffic control functions are actually present or not at the InterWorking Functions (IWFs).

Note that within the B-ISDN network indicated in Figure 20, there may be multiple network operators involved. The inter-operation between network operators within a B-ISDN is not addressed in this clause.

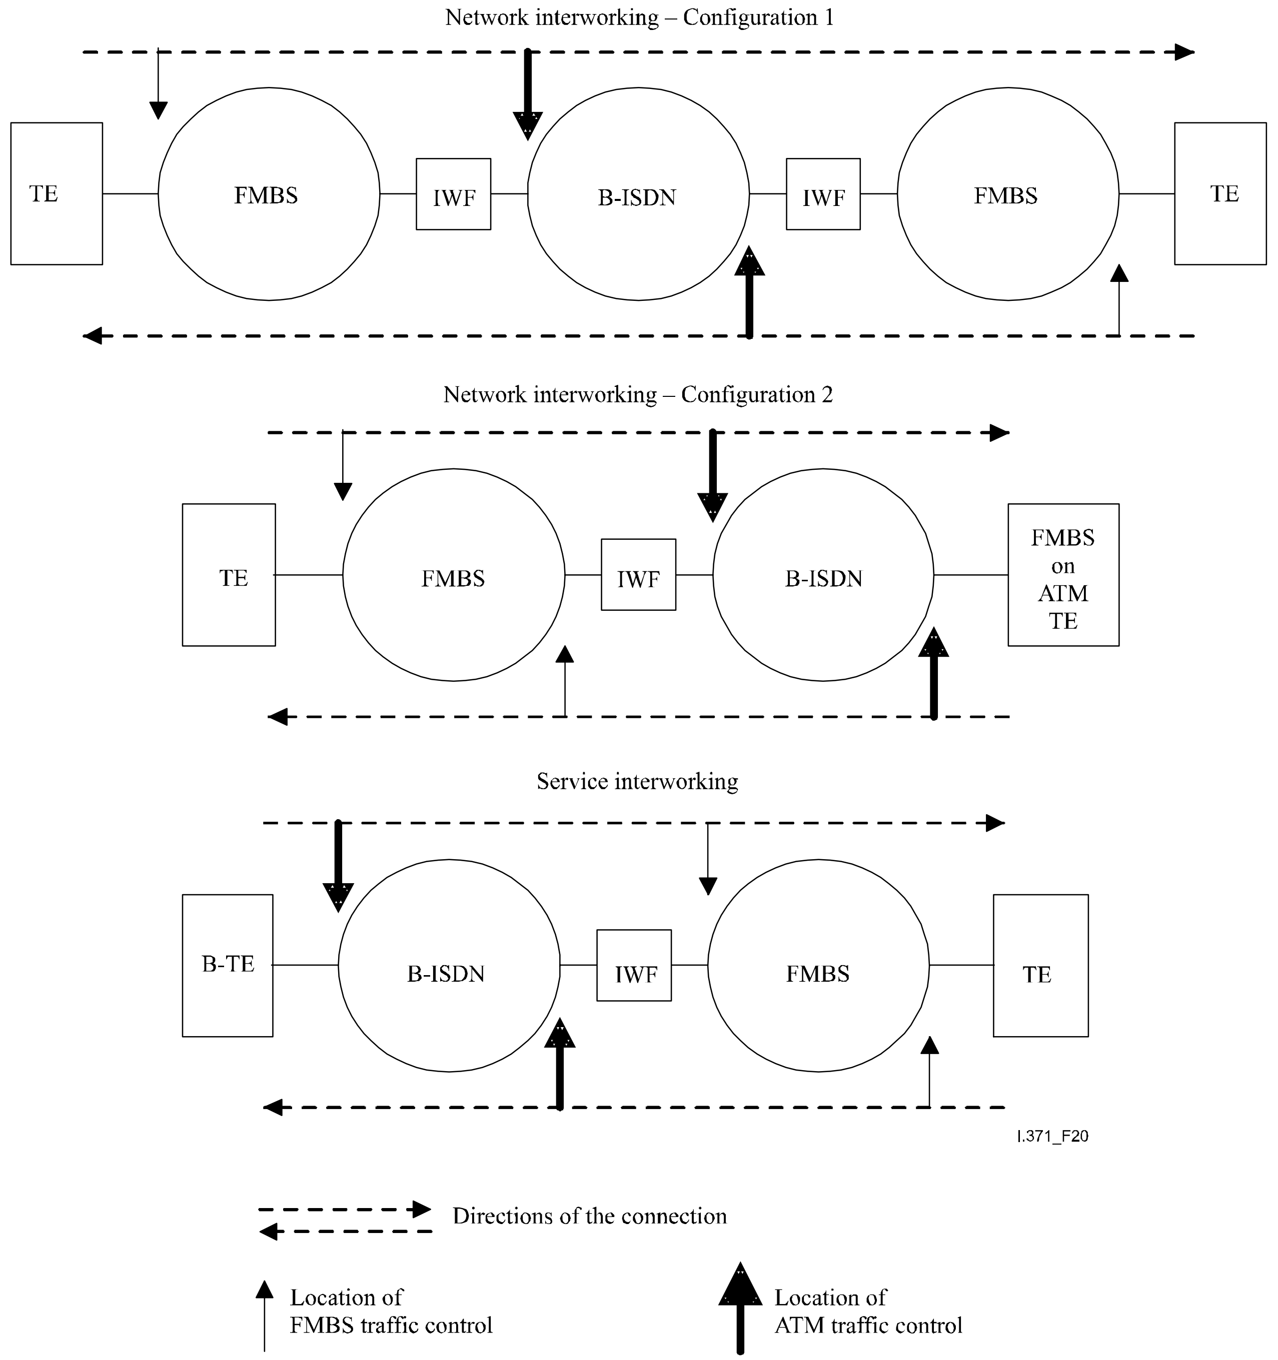

Figure 20/I.371 – Reference configurations for traffic control interworking between FMBS and B-ISDN

For interworking case 2 on Figure 20, there are two traffic contracts applying to the FMBS on ATM terminal. The implications on conformance definition are under study.

8 Procedures for traffic control and congestion control

8.1 Resource management cell format

Resource management functions that have to operate on the time-scale of round trip propagation delays of an ATM connection may need ATM layer management procedures to use resource management cells associated to that ATM connection.

The ATM layer RM cells contain fields common to all types of RM cells (see Figure 21) as well as specific fields for each type of RM cell.

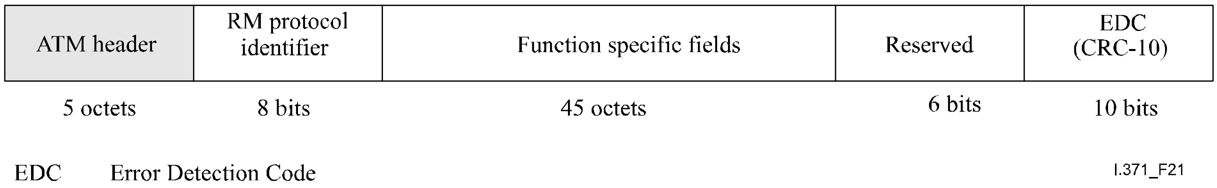

Figure 21/I.371 – Ressource management cell format

The coding principles for unused common and specific fields are:

– unused RM cell information field octets are coded 0110 1010 (6AH);

– unused RM cell information field bits (incomplete octets) are coded all zero.

The unused octets and unused bits are not to be checked by the receiver for conformance to this coding rule.

Further enhancements to this Recommendation should ensure that equipment supporting lower versions has no compatibility problems related to the content of RM cells. That is, functions and encodings of defined fields shall not be redefined in the future.

However, unused fields and unused code-points may be defined in future releases of this Recommendation and are therefore reserved.

For the purpose of this Recommendation, the leftmost bit is the most significant bit and is transmitted first.

See ITU-T Rec. I.361 for the coding of the ATM header for VC and VP resource management cells.

The EDC field carries a CRC-10 error detection code computed over the RM cell information field excluding the EDC field. See ITU-T Rec. I.610 for the encoding procedure.

VC-RM cells should be excluded from VC performance monitoring functions. VC-RM cells should be included in VP-performance monitoring functions. VP-RM cells should be excluded from VP performance monitoring functions.

Protocol IDs 248 to 255 are reserved for network specific use. RM cells identified by these protocol IDs should only cross standardized interfaces by bilateral agreements.

8.2 RM cell error conditions

In case the information in an RM cell is used, the Cyclic Redundancy Check (CRC) shall be performed to determine whether the RM cell payload is errored or not. An errored RM cell is an RM cell that contains a CRC-10 code in the EDC field that does not match the fields it protects.

The information contained in the errored RM cell is to be excluded from normal processing.

Two actions may be performed on RM cells that are detected as being errored:

– the errored RM cell is forwarded (if applicable) unmodified as errored RM cell;

– the errored RM cell is discarded.

Annex A

Generic Cell Rate Algorithm GCRA(T,τ)

This annex provides the reference algorithm that is used in 5.4 to define the cell conformance of a cell stream to the negotiated value of a cell rate $\Lambda = 1/T$, assuming that a tolerance τ is allocated. T and τ are in units of time.

The reference algorithm is described in Figure A.1. Two equivalent versions of this algorithm are shown: the virtual scheduling algorithm and the continuous state leaky bucket algorithm.

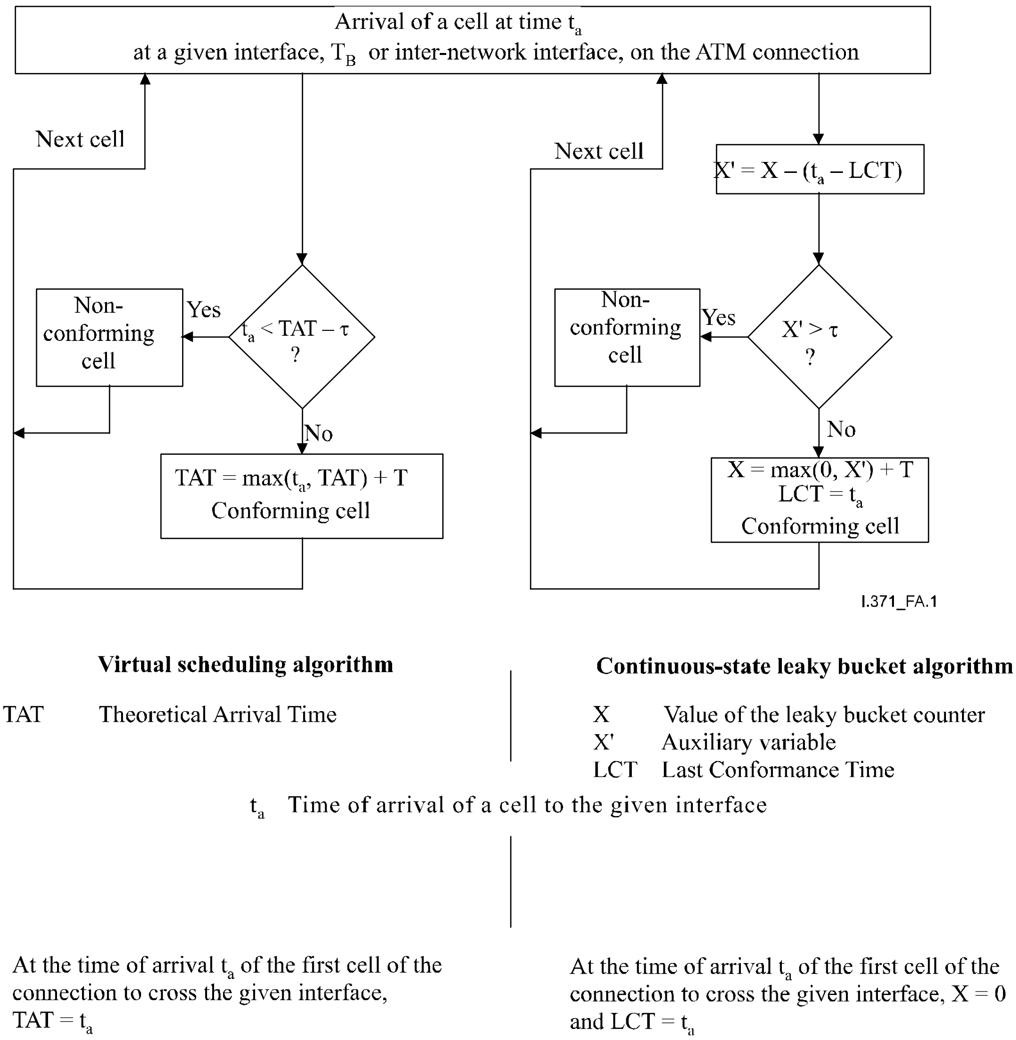

Figure A.1/I.371 – Equivalent versions of the generic cell rate algorithm

A.1 Virtual Scheduling Algorithm (VSA)

The virtual scheduling algorithm updates a Theoretical Arrival Time (TAT), which is the "nominal" arrival time of the cell assuming cells are sent equally spaced at an emission interval T corresponding to the cell rate Λ when the source is active. If the actual arrival time of a cell is not "too early" relative to the TAT and to tolerance τ associated to the cell rate, i.e., if the actual arrival time is after $(TAT - \tau)$, then the cell is conforming; otherwise, the cell is non-conforming.

Tracing the steps of the virtual scheduling algorithm (Figure A.1), 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 kth cell $t_a(k)$ is before the current value of TAT minus the tolerance τ, then the cell is non-conforming and TAT is unchanged. If the cell arrival time $t_a(k)$ is greater than or equal to $(TAT - \tau)$ but less than TAT, then the cell is conforming and TAT is increased by the increment T. Lastly, if the cell arrival time is greater than TAT, then the cell is conforming and TAT is updated to $[t_a(k) + T]$.

A.2 Continuous-state leaky bucket algorithm

The continuous-state leaky bucket 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 T 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 τ, then the cell is conforming; otherwise, the cell is non-conforming. The capacity of the bucket (the upper bound of the counter) is $(T + \tau)$.

Tracing the steps of the continuous-state leaky bucket algorithm (Figure A.1), at the arrival time of the first cell $t_a(1)$, the content of the bucket X is set to zero and the last conformance time LCT is set to $t_a(1)$. At the arrival time of the kth cell $t_a(k)$, first the content of the bucket is provisionally updated to the value X', which equals the content of the bucket after the arrival of the last conforming cell X minus the amount the bucket has drained since that arrival, $[t_a(k) - LCT]$. If X' is less than or equal to the limit value τ, then the cell is conforming and the bucket content X is set to X' (or to 0 if X' is negative) plus the increment T, and the last conformance time LCT is set to the current time $t_a(k)$. If X' is greater than the limit value τ, then the cell is non-conforming and the values of X and LCT are unchanged.

NOTE – The two algorithms presented in Figure A.1 are equivalent in the sense that, for any sequence of cell arrival times $\{t_a(k), k \geq 1\}$, 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 two algorithms have been executed, $TAT = X + LCT$.

Annex B

Application of the GCRA to SBR conformance definition

Figures B.1, B.2 and B.3 respectively show the reference algorithm that results from two instances of the GCRA operated in a coordinated mode. Note that in these figures, T_{SCR} and T_{PCR} are respectively the inverse of the SCR and of the PCR, and the parameters τ_{SCR} and τ_{PCR} are respectively $\tau_{IBT}+\tau'_{SCR}$ and the CDV tolerance τ_{PCR}, the values of the tolerance parameters that pertain at the given interface. These reference algorithms determine the conformance of cells at the given interface.

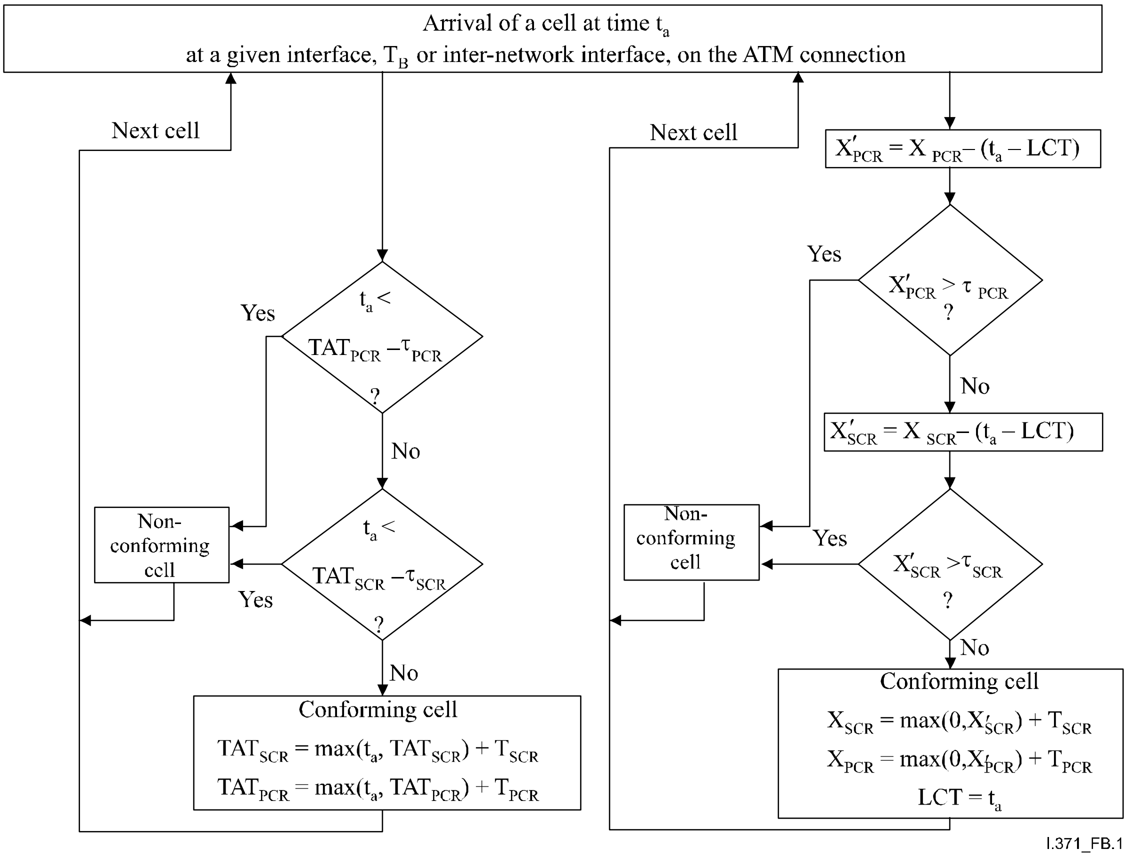

Virtual scheduling algorithm	
TAT_{SCR}, TAT_{PCR}	Theoretical Arrival Times

Continuous-state leaky bucket algorithm	
X_{SCR}, X_{PCR}	Values of the leaky bucket counters
X'_{SCR}, X'_{PCR}	Auxiliary variables
LCT	Last Conformance Time

t_a	Time of arrival of a cell to the given interface
T_{SCR}	Reciprocal of SCR for CLP = 0 + 1 cell flow
T_{PCR}	Reciprocal of PCR for CLP = 0 + 1 cell flow
τ_{SCR}	Tolerance associated with T_{SCR} ($= \tau_{IBT} + \tau'_{SCR}$)
τ_{PCR}	Tolerance associated with T_{PCR}

At the time of arrival t_a of the first cell of the connection to cross the given interface,
$TAT_{SCR} = TAT_{PCR} = t_a$

At the time of arrival t_a of the first cell of the connection to cross the given interface,
$X_{SCR} = X_{PCR} = 0$ and $LCT = t_a$

Figure B.1/I.371 – Reference algorithm for Sustainable Cell Rate (SCR) and Peak Cell Rate (PCR) traffic descriptors for CLP = 0 + 1 cell flow

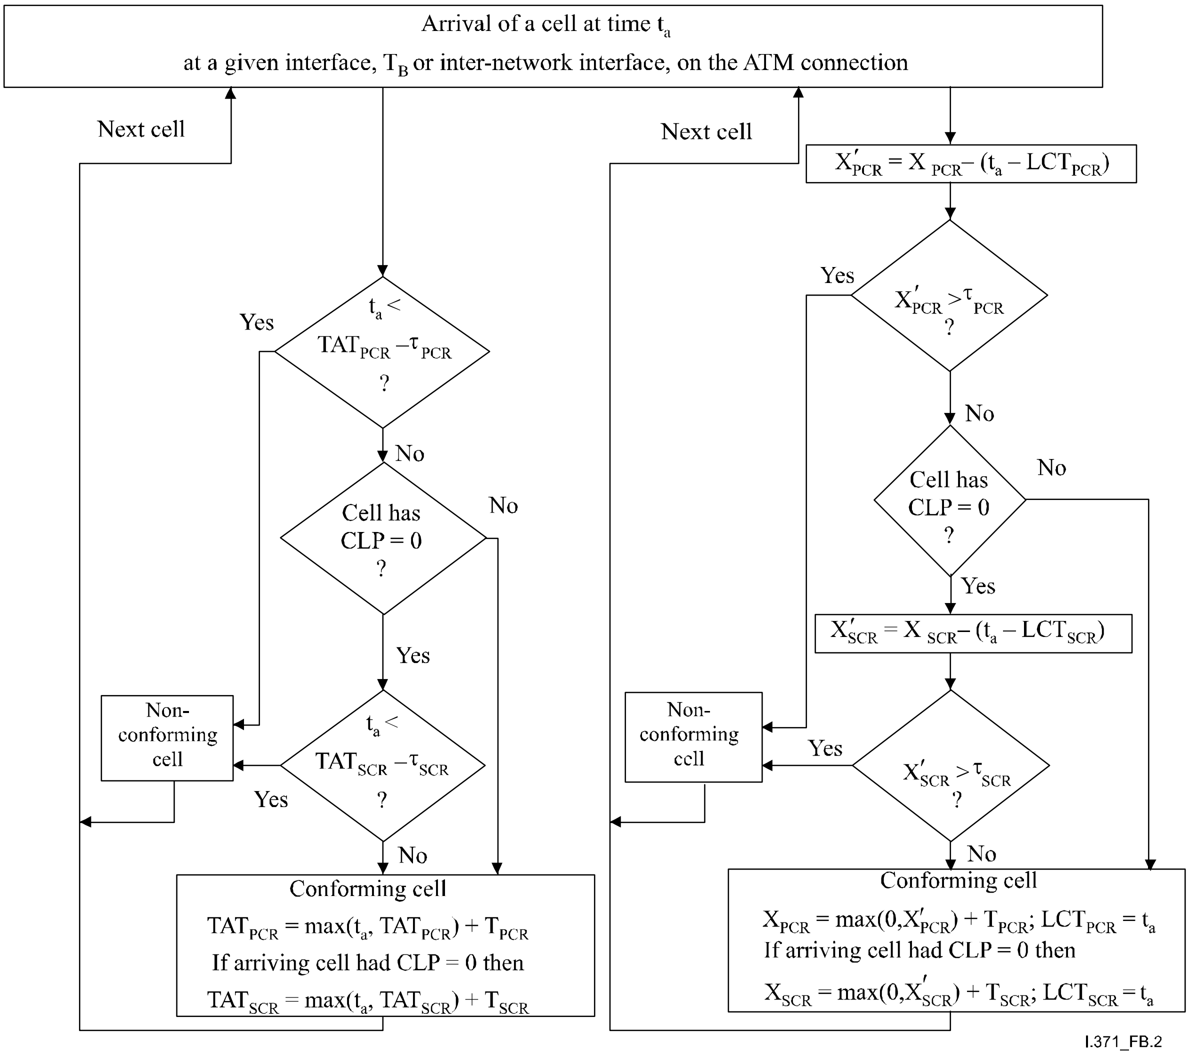

Virtual scheduling algorithm

TAT_{SCR}, TAT_{PCR} Theoretical Arrival Times

Continuous-state leaky bucket algorithm

X_{SCR}, X_{PCR} Values of the leaky bucket counters

X'_{SCR}, X'_{PCR} Auxiliary variables

LCT_{SCR}, LCT_{PCR} Last Conformance Times

t_a Time of arrival of a cell to the given interface
T_{SCR} Reciprocal of SCR for CLP = 0 cell flow
T_{PCR} Reciprocal of PCR for CLP = 0 + 1 cell flow
τ_{SCR} Tolerance associated with T_{SCR} (= $\tau_{IBT} + \tau'_{SCR}$)
τ_{PCR} Tolerance associated with T_{PCR}

At the time of arrival t_a of the first cell of the connection to cross the given interface,
$TAT_{SCR} = TAT_{PCR} = t_a$

At the time of arrival t_a of the first cell of the connection to cross the given interface,
$X_{SCR} = X_{PCR} = 0$ and $LCT_{SCR} = LCT_{PCR} = t_a$

Figure B.2/I.371 – Reference algorithm for Sustainable Cell Rate (SCR) traffic descriptor for CLP = 0 cell flow and Peak Cell Rate (PCR) traffic descriptor for CLP = 0 + 1 cell flow (Tagging does not apply)

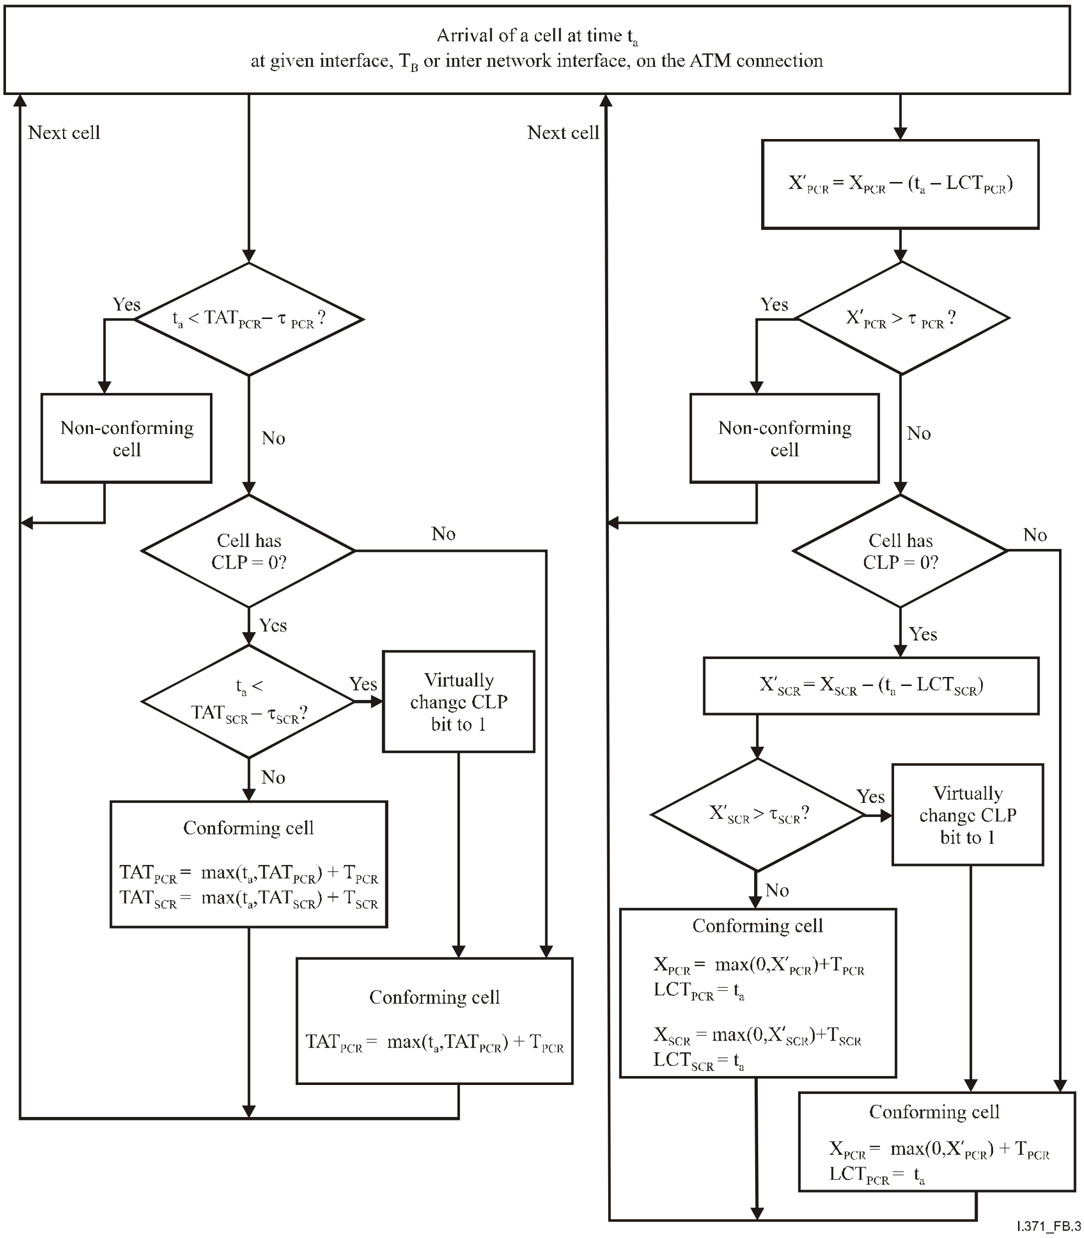

Virtual scheduling algorithm		**Continuous-state leaky bucket algorithm**	
TAT_{SCR}, TAT_{PCR}	Theoretical Arrival Times	X_{SCR}, X_{PCR}	Values of the leaky bucket counters
		X'_{SCR}, X'_{PCR}	Auxiliary variables
		LCT_{SCR}, LCT_{PCR}	Last Conformance Times

t_a Time of arrival of a cell to the given interface

T_{SCR} Reciprocal of SCR for CLP = 0 cell flow

T_{PCR} Reciprocal of PCR for CLP = 0 + 1 cell flow

τ_{SCR} Tolerance associated with T_{SCR} (= $\tau_{IBT} + \tau'_{SCR}$)

τ_{PCR} Tolerance associated with T_{PCR}

At the time of arrival t_a of the first cell of the connection to cross the given interface, $TAT_{SCR} = TAT_{PCR} = t_a$

At the time of arrival t_a of the first cell of the connection to cross the given interface, $X_{SCR} = X_{PCR} = 0$ and $LCT_{SCR} = LCT_{PCR} = t_a$

Figure B.3/I.371 – Reference algorithm for Sustainable Cell Rate (SCR) traffic descriptor for CLP = 0 cell flow and Peak Cell Rate (PCR) traffic descriptor for CLP = 0 + 1 cell flow (Tagging applies)

Annex C

ABT/DT control messages across a standardized interface

In ABT/DT, the following Block Cell Rate (BCR) modifications are possible and performed through the exchange of the following messages across a standardized interface.

1) BCR decrease (request RM cell) initiated by the source: the source does not wait for a response RM cell from the network and immediately decreases its transmission rate – Figure C.1.

2) BCR increase (request RM cell) initiated by the source: the source waits for a response RM cell from the network (acknowledgement RM cell), which is issued by the egress UNI; moreover, the egress UNI sends a BCR increase request to the destination, which is not acknowledged – Figure C.2.

3) BCR modification (request RM cell) initiated by the destination: if successful, a BCR modification request is sent by the ingress node to the source, which sends an acknowledgement to the destination – Figure C.3.

4) BCR modification initiated by the network in the forward direction (in case of non-conformance or if SCR is set equal to 0): the network sends a bandwidth modification acknowledgement to the user, which is acknowledged – Figure C.4.

The beginning of an ATM block as illustrated in Figures C.1 to C.4 is also the end of the previous ATM block, if not the first. These figures currently only apply to the rigid mode (elastic/rigid bit set to 1).

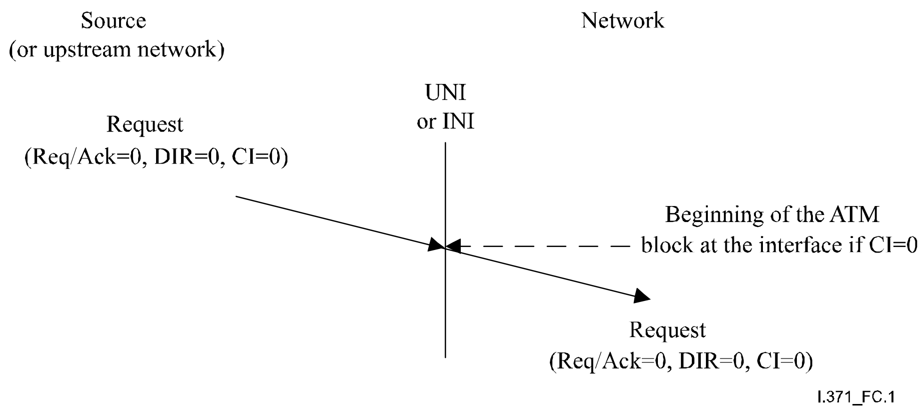

Figure C.1/I.371 – BCR decrease initiated by the source
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

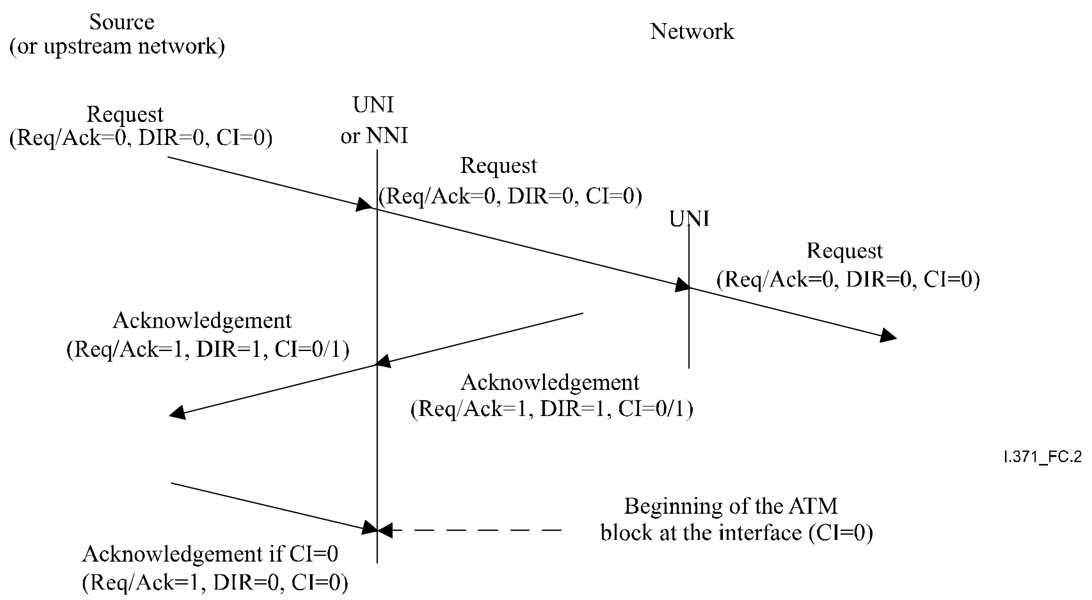

Figure C.2/I.371 – BCR increase initiated by the source
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

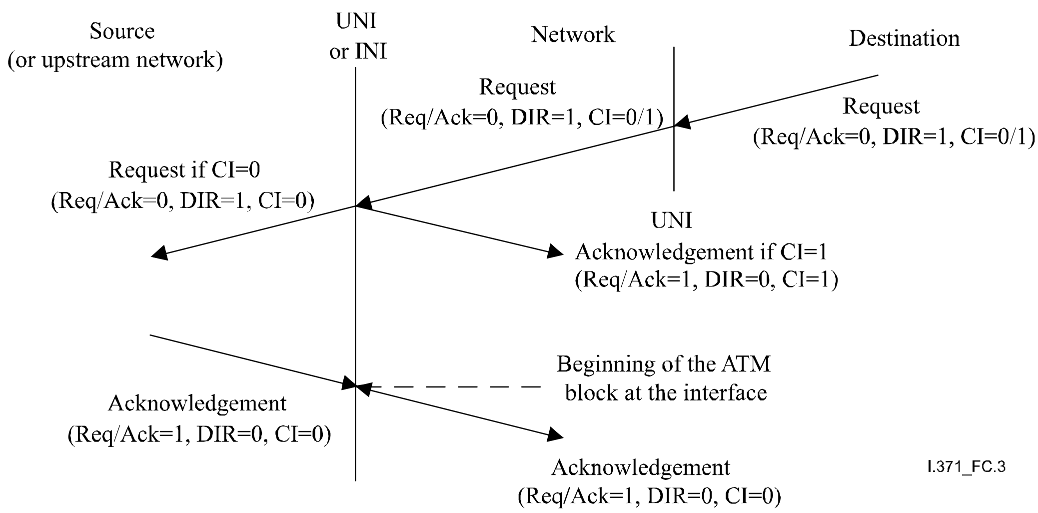

Figure C.3/I.371 – BCR modification initiated by the destination
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

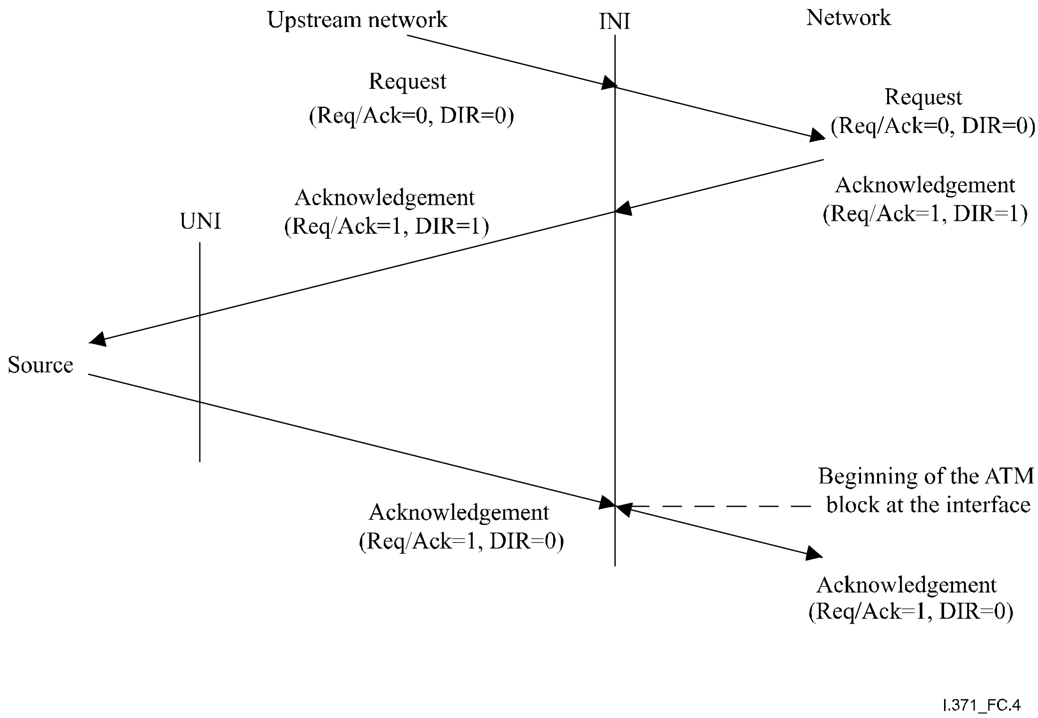

Figure C.4/I.371 – BCR modification initiated by the network in the forward direction (RM cells have maintenance = 0, traffic management = 1, CI = 0, elastic/rigid = 1)

Annex D

ABT/IT control messages across a standardized interface

In ABT/IT, the following Block Cell Rate (BCR) modifications are possible and performed through the exchange of the following messages across a standardized interface.

1) BCR modification (request RM cell) initiated by the source: the source does not wait for a response RM cell from the network and immediately modifies its transmission rate – Figure D.1.

2) BCR modification initiated by the network in the forward direction (in case of non-conformance or if SCR is set equal to 0): the network sends a bandwidth modification request to the destination, which then sends an acknowledgement to the source, which is in turn also acknowledged – Figure D.2.

The beginning of an ATM block as illustrated in Figures D.1 and D.2 is also the end of the previous ATM block, if not the first. These figures currently only apply to the rigid mode (elastic/rigid bit set to 1).

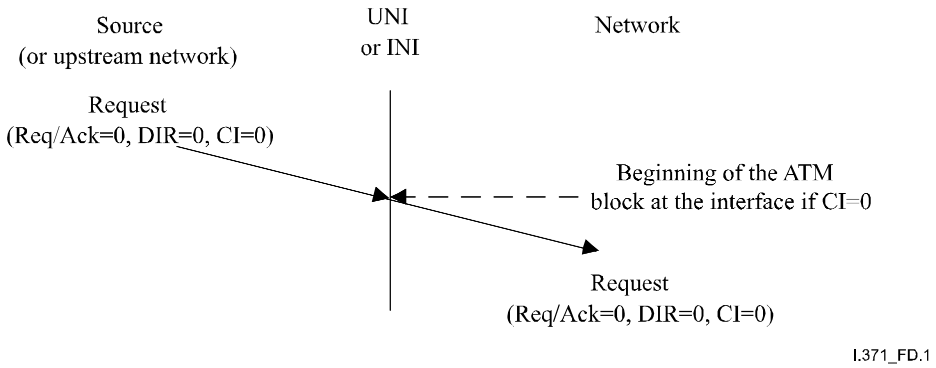

Figure D.1/I.371 – BCR modification initiated by the source (RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

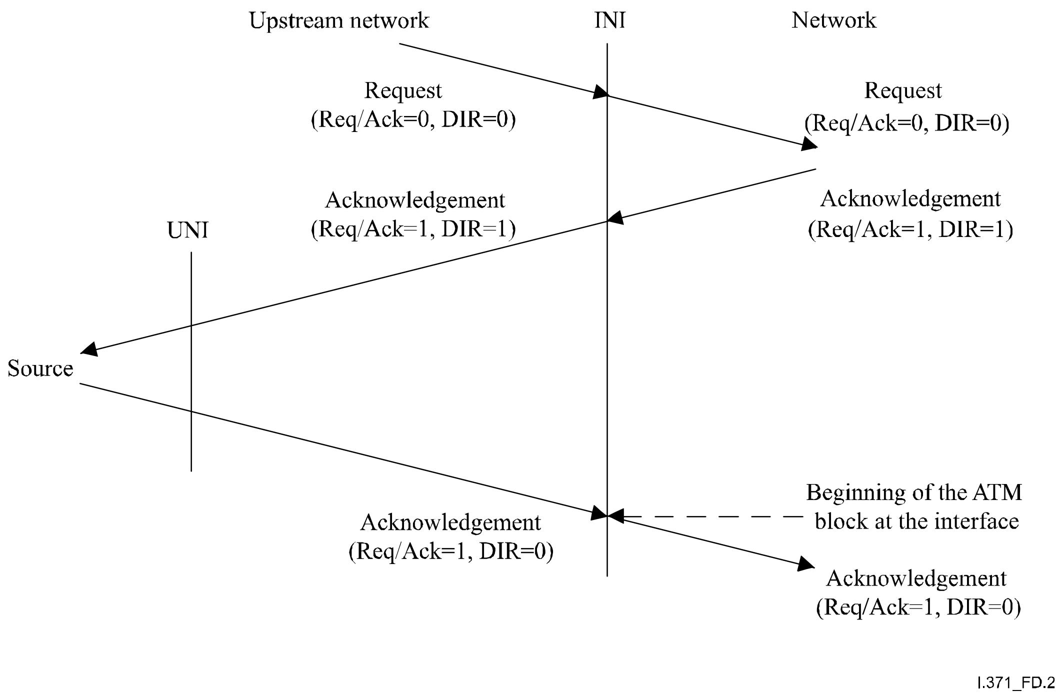

Figure D.2/I.371 – BCR modification initiated by the network in the forward direction (RM cells have maintenance = 0, traffic management = 1, CI = 0, elastic/rigid = 1)

Annex E

Avoidance of multiple outstanding BCR negotiations

To avoid multiple outstanding network-generated BCR negotiations in a network, the following priority principles between different network-generated BCR negotiations within a given network are introduced as follows:

1) A BCR negotiation request initiated by an upstream network has priority over any BCR negotiation initiated by the network considered or a downstream network. Given this priority principle, if a BCR negotiation with lower priority is pending in the network considered, then this network should interrupt the BCR negotiation with low priority and enable the higher priority level BCR negotiation to be processed.
2) If a BCR negotiation has been initiated by the network considered or an upstream network, then the network considered should deny any BCR negotiation request issued by a downstream network.

For the implementation of the above priority principles between BCR negotiations, it is desirable that two BCR negotiations processed by a given network are not identified by the same sequence number. The sequence number of the response given by a network following a BCR request should be compatible with the sequence number of the request and the priority principles between network-generated BCR negotiations. Different methods that achieve this requirement are described in Appendix V.

The interruption or denial of BCR negotiations is performed by discarding physically request or acknowledgement RM cells so that these cells do not cross a standardized interface. Examples where the above priority principles apply for ABT/DT are depicted in Figures E.1 and E.2.

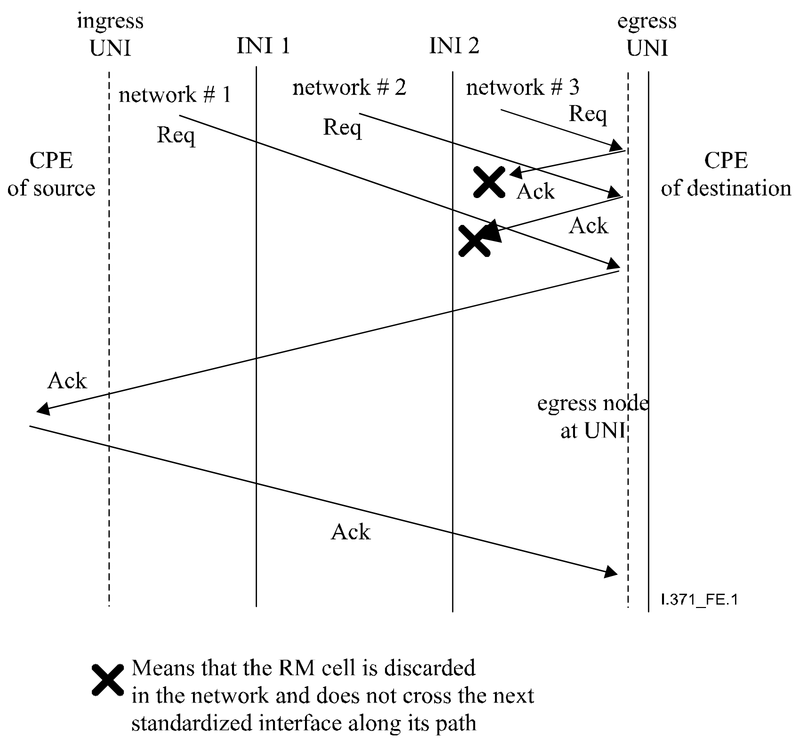

Figure E.1/I.371 – BCR negotiations denial in the originating network, or a downstream network (RM cells have different sequence numbers)

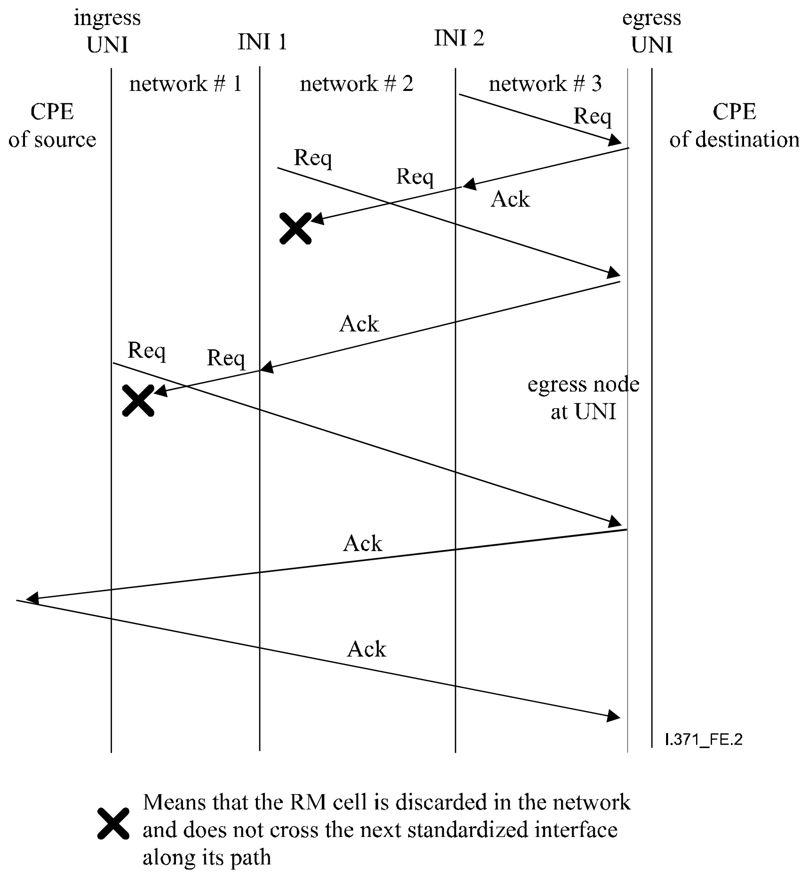

Figure E.2/I.371 – BCR negotiations denial in an upstream network (RM cells have different sequence numbers)

Annex F

ABR control messages across a standardized interface

With the ABR capability the source can obtain feedback information on transfer characteristics of the connection via:

1) Emitting an RM cell which is looped back by the destination and where network elements, as well as the destination, may alter the ECR, QueueLength, CI, or NI fields, see Figure F.1 for the case of one RM cell outstanding on the bidirectional connection and Figure F.2 for the case of multiple RM cells outstanding.

2) A network element, or the destination, originates an RM cell on the backward direction, see Figure F.3.

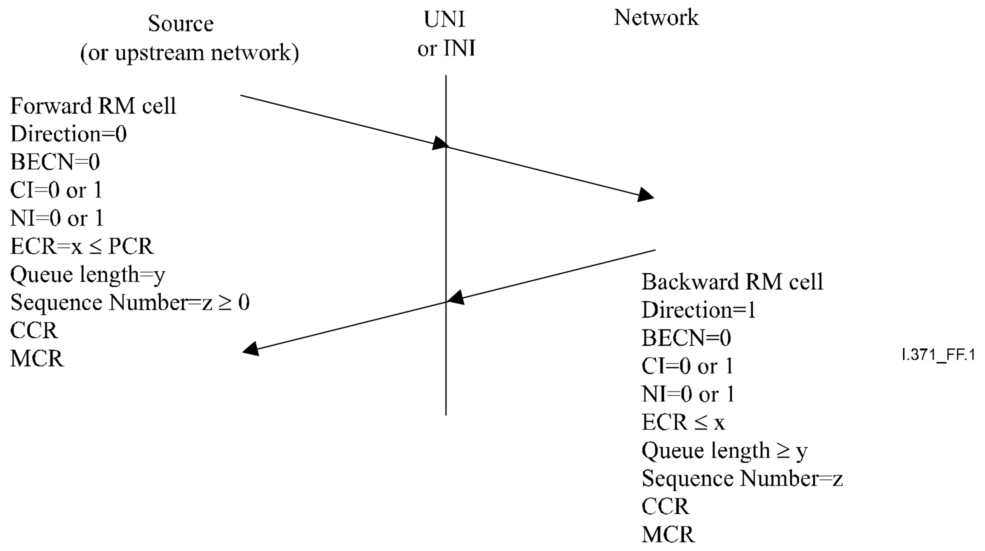

Figure F.1/I.371 – One RM cell outstanding in the bidirectional connection

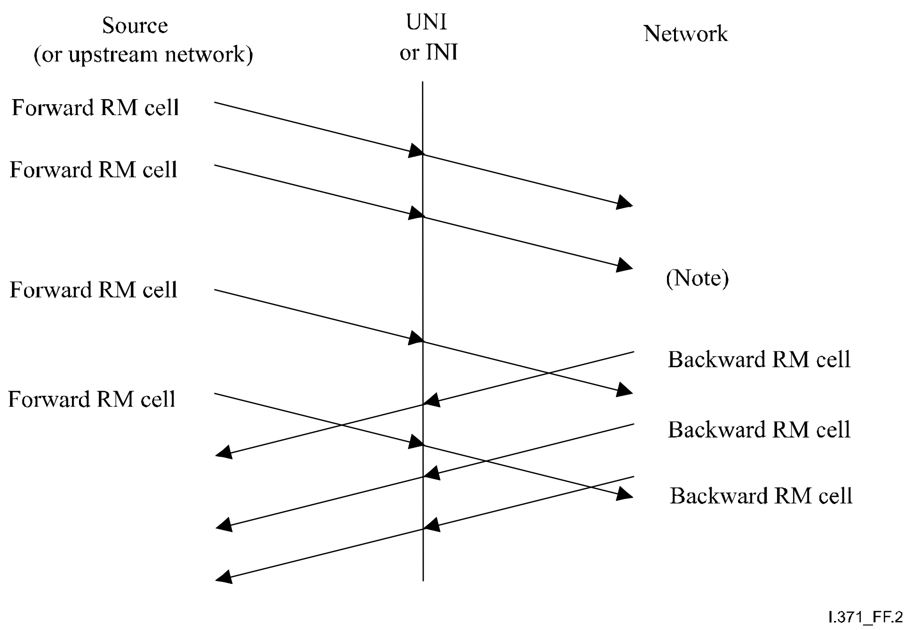

NOTE – For illustrative purposes, the destination is assumed to have not yet turned around the previous forward RM cell when the present forward RM cell arrives and the destination discards the older RM cell and turns around the present one.

Figure F.2/I.371 – Multiple RM cells are outstanding in the bidirectional connection, and where, for illustration purposes, the destination overwrites an RM cell

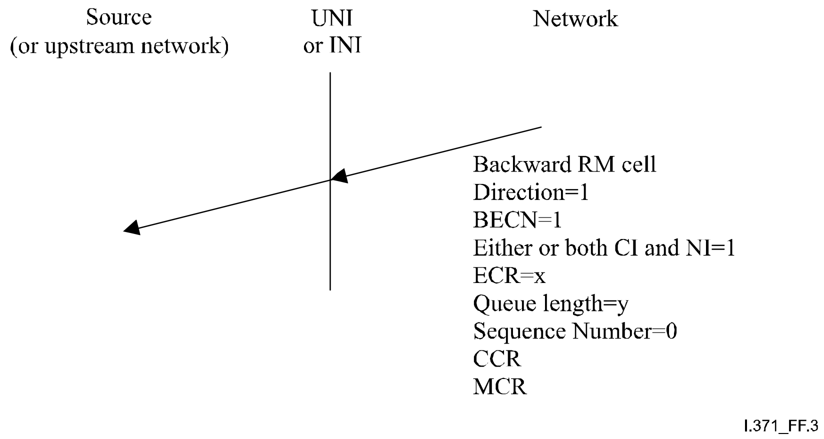

Figure F.3/I.371 – RM cell on the backward connection originated by a network element or by the destination

Appendix I

Examples of application of the equivalent terminal for the peak cell rate definition

An equivalent terminal has been used in 5.4.1 to define the peak cell rate of an ATM connection. The following two examples intend to clarify the concepts of peak emission interval T and the cell delay variation tolerance τ at T_B.

For the sake of simplicity, the transmission rate at T_B is approximated by 150 Mbit/s. Δ is the cell cycle time at the interface at T_B.

The terminology used here is the one from the virtual scheduling algorithm as shown on Figure A.1.

I.1 Configuration 1

This configuration (see Figure I.1) consists of a single terminal connected to T_B by a point-to-point single VCC.

ATM_PDU Data_Requests are generated every T = 1.25 Δ. This corresponds to a peak bit rate of 120 Mbit/s.

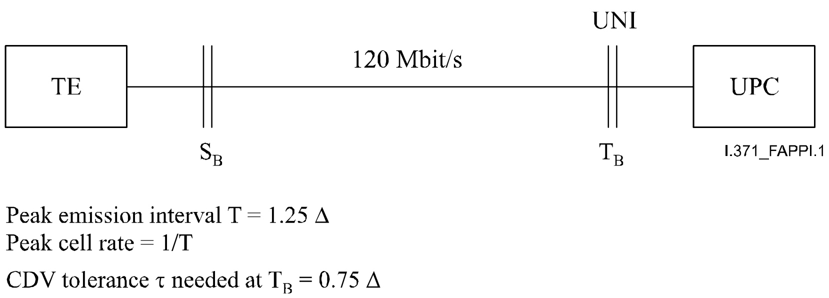

Figure I.1/I.371 – Traffic configuration 1

Figure I.2 shows the basic events on a time-scale and gives the needed CDV tolerance τ at T_B of configuration 1.

For the sake of simplicity, the propagation delay between the terminal and T_B is assumed to be zero.

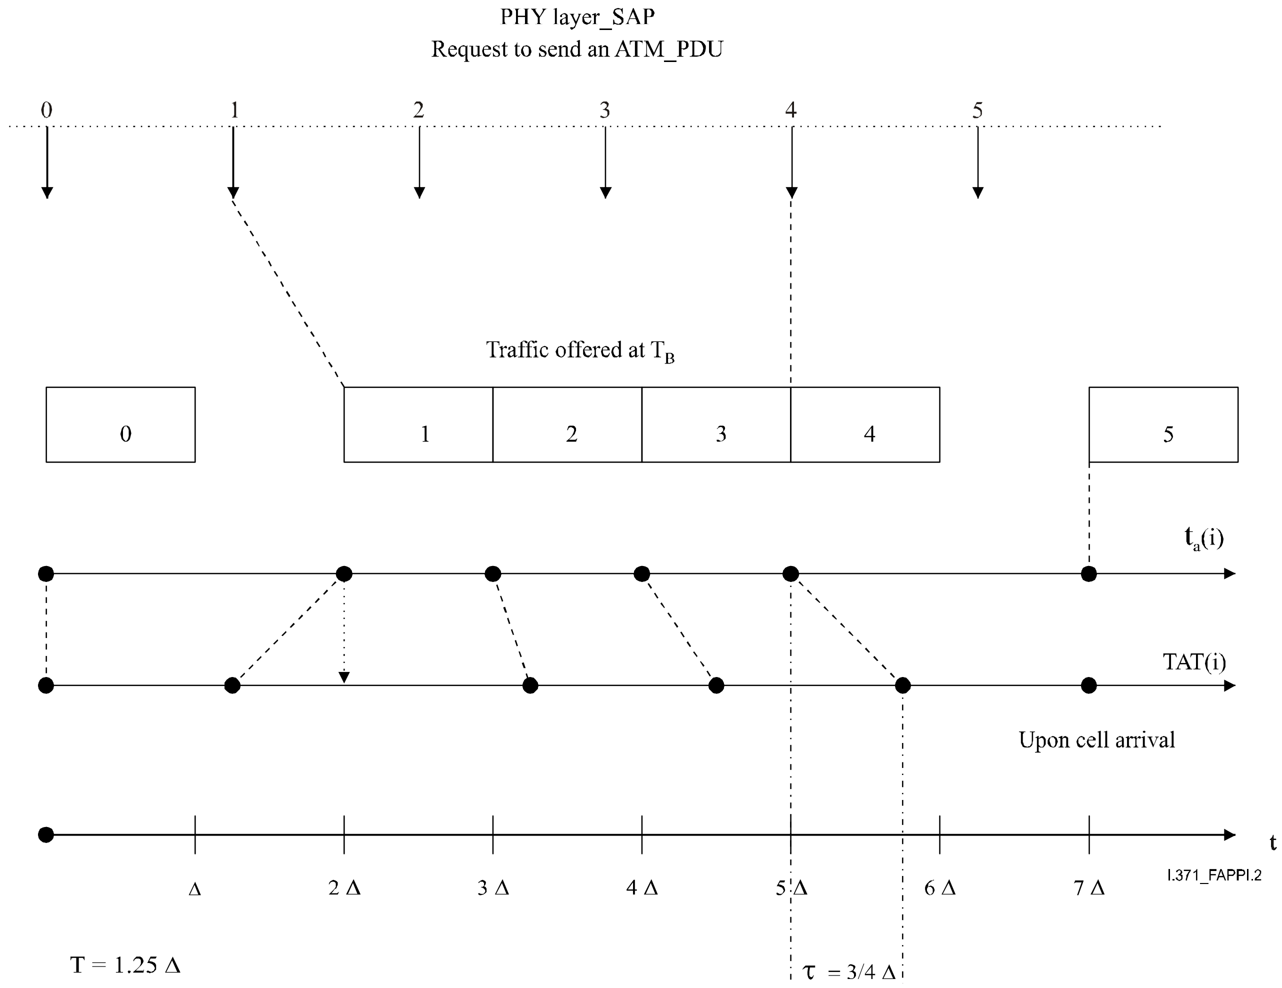

Figure I.2/I.371 – Illustration of CDV tolerance τ for traffic configuration 1

I.2 Configuration 2

This configuration (Figure I.3) consists of three terminals, each offering traffic on different VCCs. These three VCCs are multiplexed in the CEQ on to a VPC.

Terminals generate ATM_PDU Data_Requests every 10 Δ, 5 Δ and 10 Δ respectively, corresponding to 15 Mbit/s, 30 Mbit/s and 15 Mbit/s peak bit rates respectively.

The peak emission interval of the resulting VPC is T = 2.5 Δ, which corresponds to a peak bit rate of 60 Mbit/s.

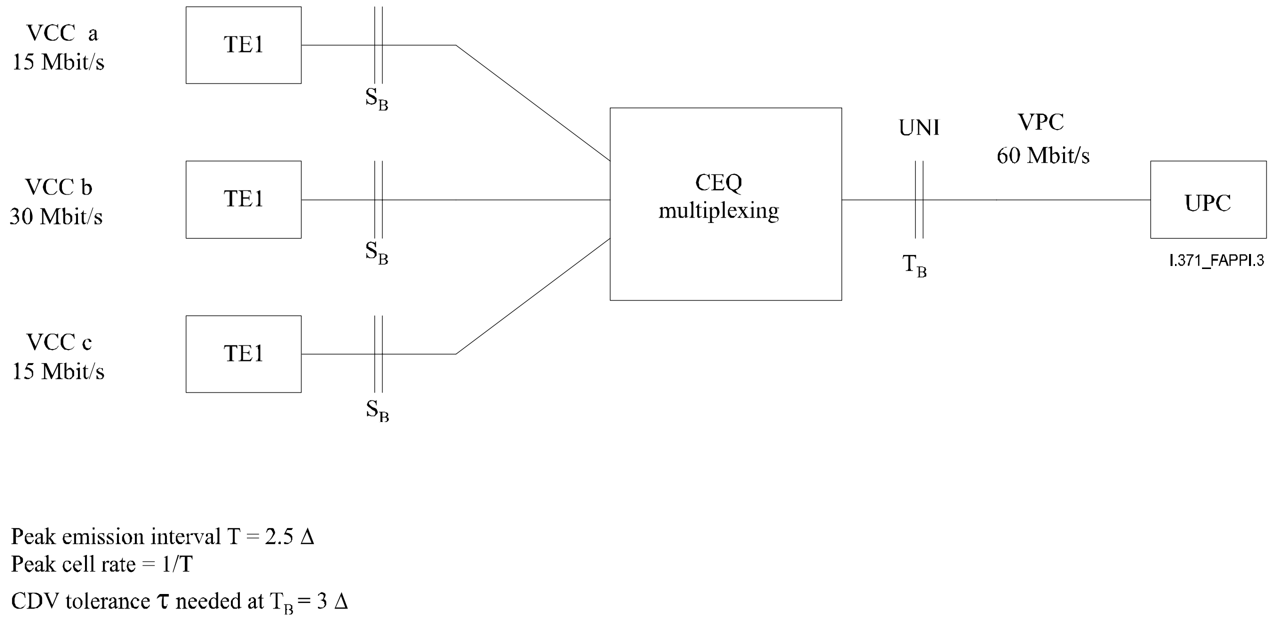

Figure I.3/I.371 – Traffic configuration 2

Figure I.4 shows the basic events and the needed CDV tolerance τ at T_B corresponding to configuration 2.

This figure and the used terminology are similar to Figure I.2.

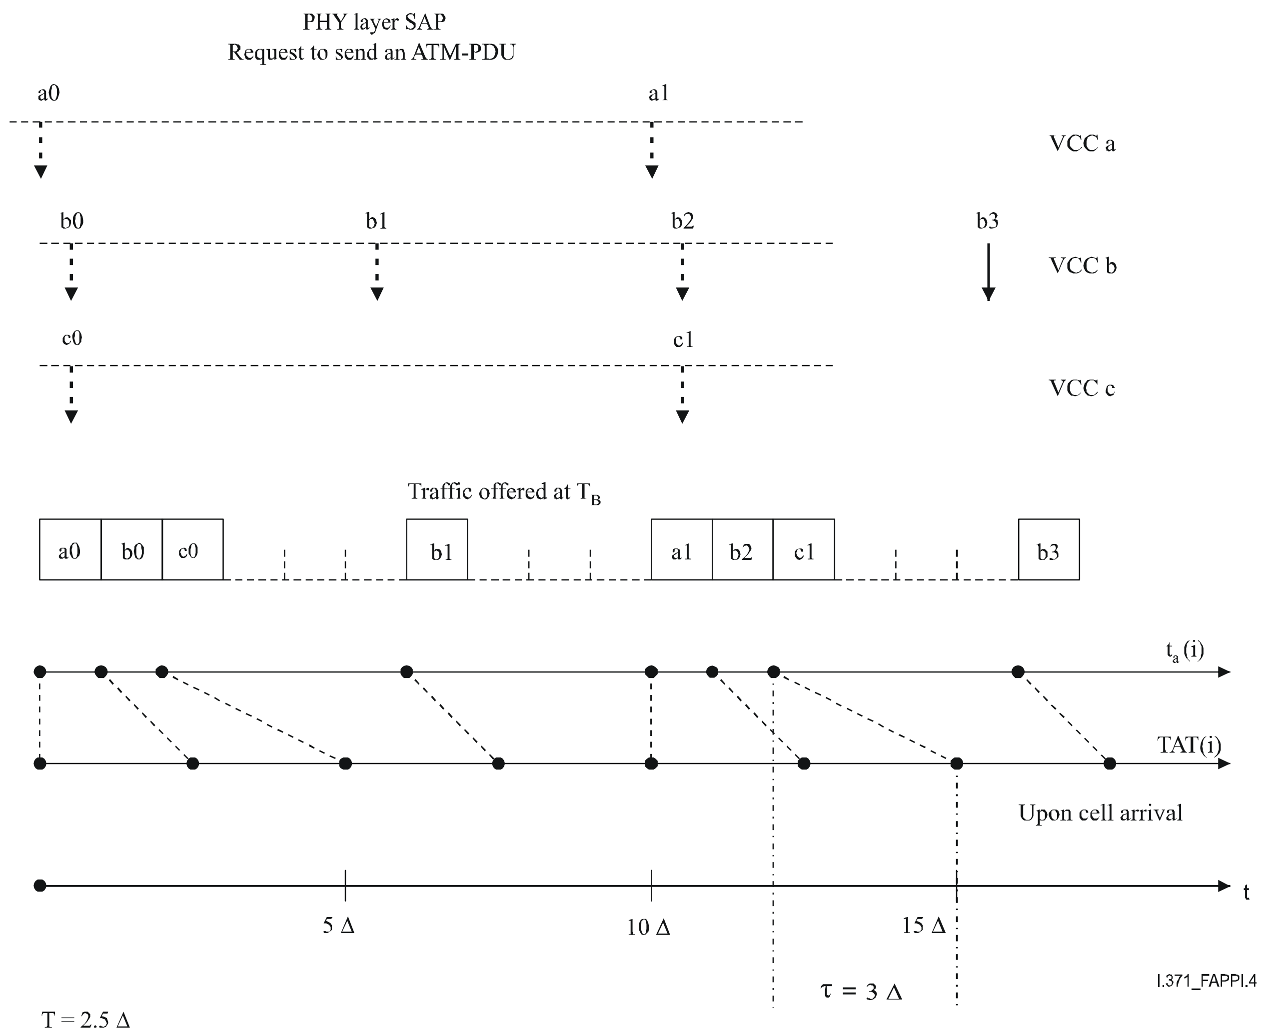

Figure I.4/I.371 – Illustration of CDV tolerance τ for traffic configuration 2

Appendix II

Transcoding rules from signalling information onto OAM traffic parameters at the ATM layer

The present release of ITU-T Rec. Q.2931 allows to signal only aggregate PCRs (user data plus user OAM). This Recommendation also allows to explicitly declare the presence of the user OAM cell flow, but without explicit OAM PCR value. The declaration is only implicit: the PCR for the user OAM component is either 1 cell per second, or 1% of the user data PCR or 0.1% of this PCR.

Furthermore, ITU-T Rec. Q.2931 does not provide the means of per call negotiating the CDV tolerance value(s). The user and the network have to rely on default values, that are negotiated for example on a subscription basis. It is therefore necessary to provide rules for computing the CDV tolerance associated either with a separate or an aggregate declaration of an OAM component. This rule can be used by a user for an implicit declaration of the CDV tolerance(s) associated with his connection.

Let T_{PCR}(agg) be the signalled aggregate PEI and τ_{PCR} (data) be the CDV tolerance value for the user data traffic. Let also P_{OAM} denote the OAM indicator in the signalling message. This indicator takes the values 0, 10^{-3} or 10^{-2}. Under the assumption that the total OAM flow is forward PM, the nominal number of cells in a cell block is either 999 or 99, depending on the value taken by P_{OAM} (if P_{OAM} is 0, there is no user OAM cell stream other than the fault management cell stream). Note that the values 99 and 999 are not standardized values for OAM block sizes given in ITU-T Rec. I.610.

The three quantities T_{PCR}(agg), τ_{PCR} (data) and P_{OAM} are considered as known in the rest of Appendix II.

In case of a separate conformance definition, the default values for traffic descriptors are given as:

$$T_{PCR}(OAM) = \tau_{PCR}(OAM) = \frac{T_{PCR}(agg)}{P_{OAM}}$$

$$T_{PCR}(data) = \frac{T_{PCR}(agg)}{1 - P_{OAM}}$$

In case of an aggregate conformance definition, the default values for traffic descriptors are given as:

$$\tau_{PCR}(agg) = T_{PCR}(agg) + \tau_{PCR}(data)$$

Appendix III

Throughput behaviour of the Generic Cell Rate Algorithm (GCRA)

This appendix describes an unexpected cell discard phenomenon when applying the reference conformance testing algorithm or GCRA to a CBR flow that slightly exceeds its agreed upon PCR. This discard situation may arise, for instance, in UPC implementation testing.

The following example illustrates this discard phenomenon:

For reasons of simplicity, a deterministic bit rate connection is considered and it is assumed that a single peak cell rate traffic parameter for the user generated CLP = 0 + 1 cell flow applies. The cell delay variation experienced by the cell flow is solely introduced by the access to the ATM slotted transfer medium. It is also assumed that the state variables used in the GCRA have an infinite precision and are not bounded. Finally, it is assumed that any cell identified as being non-conforming is subsequently discarded. The following symbols are defined:

– Λ_c: negotiated PCR with corresponding peak emission interval $T_c = \frac{1}{\Lambda_c}$

– Λ_{in}: GCRA offered input PCR with corresponding peak emission interval $T_{in} = \frac{1}{\Lambda_{in}}$

Assume that the GCRA is set with T_c and that the user is sending CBR traffic with peak emission interval $T_{in} = \frac{99}{100} T_c$. This corresponds to an excess traffic characterized by $\Delta = \frac{\Lambda_{in} - \Lambda_c}{\Lambda_c} = \frac{1}{99}$.

This means that the user is actually sending traffic about 1% above its contract.

Note that if the GCRA is set with some value T_c and a tolerance τ, then if a cell arrives and is not discarded, then for the GCRA the expression $\text{TAT} - t_a$ increases by $T_c - T_{in} = \frac{T_c}{100}$ when compared to the value of the same expression at the arrival of the previous user-generated cell.

Consequently:

– if $\tau = \frac{T_c}{100}$, every third cell is lost;

– if $\tau = \frac{T_c}{20}$, every 7th cell is lost.

This shows that the cell discard ratio (DR) highly depends on the CDV tolerance used in the GCRA and can be much larger than the intuitively expected value for DR which is Δ (≈1%). In general, one can show that the following relation approximately holds:

$$\text{for } \Delta << 1 \text{ and } \tau \le T_c, DR \approx \frac{T_c}{\tau}\Delta \text{ and for } \tau \ge T_c, DR = \Delta$$

This means that this unexpected cell discard phenomenon can only be observed if τ is chosen smaller than T_c.

Appendix IV

UPC/NPC accuracy requirements

The UPC/NPC requirements in 7.2.3.2.1 ensure that for a given connection, the number of cells discarded at the UPC is not larger than the number of cells identified as non-conforming by a conformance test at the UNI. However, the current UPC accuracy requirements allow that the cell rate enforced by the UPC be larger than the cell rate used for conformance testing at the UNI. The only accuracy requirement at the UNI is that the UPC should be capable of coding a cell rate at most 1% larger than the cell rate used in the conformance definition. This requirement pertains for cell rates as low as 160 cell/s; for cell rates between 100 and 160 cell/s, the requirement is that the coding should not be more than 1.6 cell/s larger than the cell rate used for conformance definition. These requirements apply both to PCR and SCR.

As a consequence, the number of cells of a connection which were discarded at a UPC meeting these accuracy requirements may be less than the number of non-conforming cells at the UNI.

When this connection is enforced at an INI, it may happen that the rate coded by the NPC is between the rate encoded by the UPC and the rate used for conformance definition. In this case, the NPC may discard some additional cells that would have been discarded at the UPC if the UPC would have used the conformance rate or an intermediate rate between the conformance rate and the rate encoded by the NPC.

Although this additional cell discard by the NPC is allowed from a conformance viewpoint, it could cause difficulties if performance monitoring is performed on a segment starting after the UPC and passing through one or several NPCs. In this case, the extra cells discarded by the NPCs would be counted as lost cells. This may then result in an increased number of errored cell blocks.

Moreover, the accuracy requirements on the UPC are requirements on the capabilities of the UPC. They are not requirements on how a network operator chooses to use the UPC. In particular, a network operator may choose to set the parameters in the UPC with a margin greater than 1%.

It is also noted that this problem arises only if the connection contains non-conforming cells at the UNI.

If two or more network operators agree to set up a performance monitoring segment across one or more NPCs, they may improve the accuracy of their performance monitoring in the following way: the NPCs involved should be set to a cell rate that is larger than or equal to the cell rate encoded at the UPC.

Appendix V

Examples of methods ensuring unicity of RM cell numbering in ABT

In order to implement the priority scheme (described in Annex E), that allows to discriminate between potentially conflicting BCR requests initiated by the network, it is necessary in some cases to rely on the Sequence Number (SN) value; this can be done only if cells corresponding to different BCR negotiations are allocated different SN values. However, Request RM cells that are generated by different networks may carry identical SN values unless a specific scheme is implemented. No scheme is currently recommended to ensure this property. Three possible methods are described in this appendix.

V.1 Segmentation of the SN field between different networks

It is possible to segment the coding of the 4-octet SN field between the networks along the connection. This would naturally prevent two different BCR negotiations being identified with the same SN, since a given network should not initiate a new BCR renegotiation while one initiated by itself is pending.

V.2 Proprietary handling of SN field

For instance, if a network is processing a BCR negotiation identified by a given sequence number and if this network receives a BCR request with higher priority but with the same SN value, the network may change the sequence number of this latter BCR transaction for processing in this network and the downstream networks but the network considered should also restore the initial sequence number value in the response to upstream networks. The actions taken when different RM cells have the same sequence number values are depicted in Figure V.1.

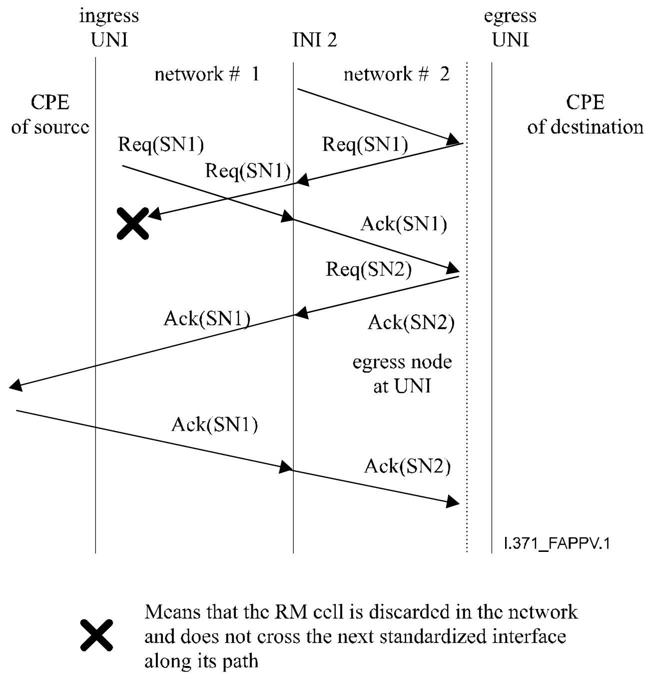

Figure V.1/I.371 – Avoidance of conflicts between sequence numbers

V.3 Segmentation of the SN field for indicating relative location of RM cell

The following scheme may be envisaged: out of the four available octets in the SN field of the ABT RM cell, three are used to assign a number (NA) to each RM cell that is generated, and one (RL) is used to identify the location of the network where it is observed at a given time, relative to the network that generated it.

- A cell is generated by the network, in the forward direction, as a Req cell, with a given NA and RL = 0;
- NA is not modified when the RM cell crosses an interface;
- when the request crosses a standardized interface, RL is increased by 1;
- when the cell is turned back as an Ack at the destination UNI, RL is unchanged;
- when the Ack crosses a standardized interface, RL is decreased by 1 until RL = 0;
- the Ack cell is then passed across the interface, as a Req, with RL = 1;
- when the request crosses a standardized interface, RL is increased by 1.

The above scheme ensures that, in the network that has originated the BCR renegotiation, an Ack cell corresponding to a given Req cell carries exactly the same sequence number (NA, RL). Moreover, two RM cells generated in different networks have necessarily different sequence numbers (different RL values).

In order to identify the priority level of a given cell, one stores the number (NA, RL) carried by the RM/Req cell at the given interface. Upon reception of either an Ack, or another Req, it is possible to identify the cell with the highest priority.

Appendix VI

Derivation of conformance definition parameters for ABT

Consider an ABT connection, which is conforming to the peak cell rate $1/T$, the sustainable cell rate $\Lambda^0{}_{SCR}$, and the maximum block size MBS^0 at the PHY-SAP of the equivalent terminal. These parameters are specified in the traffic contract. Moreover, the forward and backward user request RM cell flows at the interface considered are conforming to the GCRA(T_{RM}, τ_{RM}) and $GCRA(T'_{RM}, \tau'_{RM})$, respectively. The number $S(0,t)$ of cells, which can be transmitted at the PHY-SAP of the equivalent terminal over the time interval $(0,t)$ satisfies:

$$S(0,t) = \sum_{\substack{\text{number of ATM} \\ \text{blocks in } (0,t)}} \rho_i \left(t_i^d - t_i^f\right) \leq \Lambda_{SCR}^0 \times t + MBS^0$$

where ρ_i is the BCR in cell/s of ATM block i and t_i^d and t_i^f are the starting and finishing times of ATM block i, respectively. t_i^d and t_i^f are in fact the transmission times of the leading and the trailing RM cells of ATM block i, respectively.

Define t''_{SCR} as in 6.6.1.4.4. To determine the worst case for the amount of resources consumed by the ABT connection, assume that the leading RM cells experience the minimum virtual cell transfer delay and that the trailing RM cells of ATM blocks experience the maximum virtual cell transfer delay. The size of ATM block i is then at most increased by $t''_{SCR} \times \rho_i$ cells.

It then follows that under the assumption that the ABT connection is conforming to the peak cell rate $1/T$, the sustainable cell rate Λ_{SCR}^0 and the maximum block size MBS^0 at the PHY-SAP of the equivalent terminal, the number S' of cells which can be transmitted at the interface satisfies, by noting that $\rho_i \leq \frac{1}{T}$,

$$\begin{aligned} S'(0,t) = & \sum_{i=1}^{n(t)} \rho_i \left(t_i^d - t_i^f\right) + \sum_{i=1}^{n(t)} \rho_i \times \tau''_{SCR} \\ & \leq t \times \Lambda_{SCR}^0 + MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times n(t) \end{aligned} \qquad \text{(VI-1)}$$

where $n(t)$ is the number of ATM blocks over $(0,t)$.

The number $n(t)$ of ATM blocks actually depends on the transmission mode and on the traffic contract for the user request RM cell flows.

In the immediate transmission mode (ABT/IT), taking into the traffic contract on the user request RM cell flow issued by the source on the forward direction, this number satisfies:

$$n(t) \leq \frac{t}{T_{RM}} + \sigma_{RM} \qquad \text{(VI-2)}$$

where $\sigma_{RM} = \left\lfloor 1 + \frac{\tau_{RM}}{T_{RM} - \Delta} \right\rfloor$ with Δ denoting the cell transmission time. It follows that:

$$S'(0,t) \leq t\left(\Lambda_{SCR}^0 + \frac{1}{T} \times \tau''_{SCR} \times \frac{1}{T_{RM}}\right) + MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times \sigma_{RM} \qquad \text{(VI-3)}$$

As a consequence, the cell stream at the interface is characterized by the sustainable Λ_{SCR} and the (fractional) maximum burst size MBS defined by:

$$\Lambda_{SCR} = \min\left(\Lambda^0_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \frac{1}{T_{RM}}, \frac{1}{T} \right)$$
$$MBS = MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times \left\lfloor 1 + \frac{\tau_{RM}}{T_{RM} - \Delta} \right\rfloor \qquad \text{(VI-4)}$$

The tolerance τ_{SCR} is determined by using the relation:

$$\tau_{SCR} = (MSB - 1)(T_{SCR} - T) \qquad \text{(VI-5)}$$

In the delayed transmission mode (ABT/DT), since BCR negotiations can be initiated by both the source and the destination, not only the number of ATM blocks due to the source but also those due the destination should be taken into account. The user request RM cell flow of the destination should be conforming at the interface considered to the $GCRA(T'_{RM}, \tau'_{RM})$ (the parameters T'_{RM} and τ'_{RM} are known at connection establishment). The aggregation of the user request RM cell flows generated by both the source and the destination may give rise at most to $n(t)$ ATM blocks over the time interval $(0,t)$ with:

$$n(t) \le t\left(\frac{1}{T_{RM}} + \frac{1}{T'_{RM}} \right) + \sigma''_{RM} \qquad \text{(VI-6)}$$

where:
$$\sigma''_{RM} = \left\lfloor 2 + \frac{\tau_{RM}}{T_{RM} - \Delta} + \frac{\tau'_{RM}}{T'_{RM} - \Delta} \right\rfloor$$

It follows that the connection is characterized at the interface considered by the sustainable Λ_{SCR} and the (fractional) maximum burst size MBS defined as:

$$\Lambda_{SCR} = \min\left(\Lambda^0_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \left(\frac{1}{T_{RM}} + \frac{1}{T'_{RM}} \right), \frac{1}{T} \right)$$
$$MBS = MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times \left\lfloor 2 + \frac{\tau_{RM}}{T_{RM} - \Delta} + \frac{\tau'_{RM}}{T'_{RM} - \Delta} \right\rfloor \qquad \text{(VI-7)}$$

Tolerance τ_{SCR} relevant to the ATM block conformance definition for ABT/DT is deduced by using equation (VI-5).

Remark – The corrective terms in the above formulae giving the parameters to take into account in the block level conformance definition depend on the traffic characteristics of the RM cell flows. In general, T_{RM} is taken large enough (a fraction of the round-trip time through the network for ABT/IT and several times this round-trip time for ABT/DT). Moreover, τ_{RM} should be chosen sufficiently small so as to avoid clumps of RM cells. It thus turns out that the corrective terms are, in general, small when compared to the intrinsic parameters.

Appendix VII

Source, destination, and network element reference behaviours for ABR

VII.1 Source reference behaviour

In order to obtain full use of the dynamic bandwidth of an ABR connection, a source needs to send RM cells in the forward direction (i.e., a Forward RM cell) of the information flow. A source receives RM cells in the backward direction (i.e., a Backward RM cell) unless these cells have been lost in the network. For efficient operation of the close loop control, a source needs to adapt regularly to the changing network conditions. An ABR source interprets an errored backward RM cell as not received (for EDC field, see 8.1).

User data cells are emitted with the CLP bit set to 0.

The first cell sent by the ABR source should be an in-rate forward RM cell. The source should insert an in-rate forward RM cell at least every (N_{RM}–1) other in-rate cells. It should also insert at least one backward RM cell between two forward in-rate RM cells if backward RM cells are waiting for transmission. The parameter N_{RM} should be network-specific or set to a default value.

In forward RM cells, the source should set the MCR field to MCR, and set the CCR field to be equal to the current ACR.

At the PHY-SAP of the equivalent terminal (see 5.4), an active source should emit in-rate cells at a rate no more than the current Allowed Cell Rate (ACR). The value of ACR shall never exceed PCR, nor shall it ever be less than MCR.

At connection establishment, the source uses the values TBE and FRTT to possibly further reduce the negotiated IACR as follows: If the value $\max(MCR, \lceil TBE/FRTT \rceil)$ is smaller than the negotiated IACR then IACR is reduced to that value, where $\lceil x \rceil$ stands for rounding up x to the nearest integer value.

A source should update its ACR according to the information received in backward RM cells:

1) If the ECR value is less than ACR, then ACR should be reduced to ECR, but not less than MCR.

2) If the ECR value is greater than the ACR, then ACR may be increased (unless the backward RM cell is a BECN, in which case the ACR shall not be increased). The increase in ACR should be limited by a fixed increment RIF × PCR that provides for stepwise convergence to ECR. If the incremented ACR is greater than ECR, it is set to ECR. Setting RIF to 1 would allow an immediate jump to ECR. The Rate Increase Factor (RIF) would be set by default or assigned at connection establishment.

3) A source may make use of the CI and NI bits:

   a) If the source receives a RM cell with CI = 1, then the value of ACR (in effect prior to the arrival of the backward RM cell) should be reduced by a multiplicative factor, but not further than MCR. Specifically, ACR should be reduced by at least ACR × RDF where the parameter RDF, the Rate Decrease Factor, can be set by default or assigned at connection establishment by management procedures or by signalling.

   b) If the backward RM cell has CI = 0 and NI = 0, then the ACR may be increased by at most the additive increment RIF × PCR to a rate not greater than the PCR.

   c) If the backward RM cell has NI = 1, then the source should not increase the ACR.

d) If the value of ACR resulting from steps 3) a) to 3) c) is greater than the value of ECR in the backward RM cell, then the ACR should be reduced to a value that is less than or equal to ECR, but not less than MCR. Otherwise, the source should use the ACR value calculated from the CI and NI bits alone.

4) In addition, if the source makes use of the QueueLength field and if the QueueLength is non-zero, then the sending rate should be further decreased, or no cells should be sent for a period of time, so as to allow the queue length to decrease. Procedures for calculating rate reductions and intervals based on non-zero QueueLengths are under study.

In addition to ACR updates due to reception of backward RM cells, a source should update its ACR according to the following rules:

5) When a source initializes, it should set the allowed cell rate, ACR, to at most the Initial Allowed Cell Rate (IACR) and the first in-rate cell sent should be a forward RM cell. The value of IACR is greater than or equal to MCR. At the beginning of the connection, the user is allowed to send at most the Transient Buffer Exposure (TBE) number of cells at the IACR without receiving a backward RM cell, which explicitly allocates an ACR. When TBE cells have been transmitted without any received RM cell, the source should reduce its rate, in steps or in a single step, to MCR. IACR can be negotiated between the network and the user at connection establishment. The value of TBE is assigned to the connection via management procedures or signalling.

6) A source which has not emitted any in-rate cells for a sufficiently long period should reduce its ACR to IACR if its ACR is above IACR to reflect the reallocation of network resources that may have taken place during the period of inactivity:

   Before sending an in-rate forward RM-cell, if ACR > IACR and the time that has elapsed since the last in-rate forward RM cell was sent is greater than ADT (the ACR Decrease Time), then ACR should be reduced to IACR; otherwise the value of ACR is not changed. ADT may be negotiated or set to a network specific default value. Its value could be in the order of a few hundred ms.

   When the source becomes active again, it should behave as in the previous item 5), using the (possibly reduced) allowed cell rate.

7) A source which has failed to receive a backward RM cell in a sufficiently long period of time should reduce its sending rate but need not reduce it below the MCR:

   Before sending an in-rate forward RM-cell, if at least MCR in-rate forward RM cells have been sent since the last valid backward RM cell with BECN = 0 was received, then ACR should be reduced by at least $ACR \times CDF$, unless that reduction would result in a rate below MCR, in which case ACR is set to MCR.

   MCR could be set to $\lceil TBE/N_{RM} \rceil$. Since in-rate forward RM cells are sent at least every N_{RM} cells, the first ACR decrease would happen after about TBE forward user data cells. CDF can be negotiated per connection or set to a network specific default value. It could have a value in the order of 1/16.

VII.2 Destination reference behaviour

A destination enables its corresponding source to estimate the bandwidth available from the network by returning RM cells to the source.

1) The destination should turn all RM cells received around to return them to the source. The direction bit, DIR, should be changed from "Forward" to "Backward".

2) If a destination cannot turn around a forward RM cell before it receives a subsequent forward RM cell to be turned around on the same VC, it may return only the most recent forward RM cell and discard the older forward RM cells. Alternatively, it may emit the older RM cell with the CLP bit set equal 1, and with the contents of the older cell possibly

overwritten by the contents of the newer cell. However, loss of a backward CLP = 1 RM cell between a standardized interface and the source can cause misalignment between the ACRs at the source and the conformance definition at the interface, which may impact the QoS of the connection. If a destination finds that it does not have an adequate ACR on the backward connection to support the emission of backward RM cells, it should consider itself to be in internal congestion and act as in item 4) below.

3) If an EFCI = 1 has been received on the data cell prior to the RM cell, then the destination should mark the backward RM cell. An implementation may either:
   a) reduce the ECR; or
   b) set the CI bit in the RM cell.
4) To declare itself congested, the destination may do one or more of the following:
   a) further reduce ECR to whatever rate it can support;
   b) set the CI bit and/or NI bit;
   c) further increase the value of the QueueLength field in the RM cell.

A destination can also generate a backward RM cell without having received a forward RM cell. These cells are BECN cells. These cells have the following characteristics:

- BECN cells have the CLP bit set to 0.
- The BECN bit in the message field must be set.
- The direction should be "backward".
- Either the CI or the NI bit is set to 1.

Other interactions between forward bandwidth, backward bandwidth and the frequency with which RM cells are sent are for further study.

VII.3 Network element reference behaviour

The network element may modify transiting RM cells based on the network element state. The need for network elements to insert forward RM cells is for further study.

A network element is not allowed to update the ABR RM cell fields protected by the EDC field if the CRC-10 code in the EDC field is wrong.

A network element shall implement at least one of the following methods to control congestion at queueing points:

1) The network element may reduce the ER field of forward and/or backward RM cell (Explicit Rate Marking).
2) The network element may set the EFCI flag in the data cell headers (EFCI marking).
3) The network element may set CI = 1 or NI = 1 in forward and/or backward RM cells (Relative Rate Marking).
4) The network element queueing point may set the queue length field of the RM cell to the maximum of the present value and the number of cells queued for this VC at this queueing point for this connection.

The explicit rate feedback provided by a network element is derived from the defined allocation policy.

In addition, the network element may segment the ABR control loop using a virtual source and destination (VS/VD control).

A network element may generate backward RM cells, called Backward Explicit Congestion Notification (BECN) cells. These cells have the following characteristics:

- BECN cells have the CLP bit set to 0.
- The BECN bit in the message field must be set.
- The direction should be "backward".
- The CI or the NI bit is set to 1.

The ECR in the BECN cell that the network element inserts should be not larger than the ECR in the last backward RM cell that left the network element in the same direction (if there are any such RM cells). See also VII.1 on source reference behaviour item 2).

The maximum rate of network-element-generated BECN cells is currently not specified but should be consistent with the mutual agreement limiting the aggregate rate of BECN cells that applies at standardized interfaces.

Backward RM cells may be serviced out of sequence with respect to data cells. Priority of forward RM cells is for further study. Bounds on the relative delay of RM cells with respect to data cells is under study.

In the special case when 1/ECR becomes large relative to the round-trip time, it is no longer reasonable to send at least one RM cell per round-trip time. This has the effect of increasing the feedback time beyond a round-trip time. As a consequence, one may require additional buffer allocations for those VCs. This is for further study.

VII.4 Impact of source reference behaviour on traffic characteristics

When a user follows the source reference behaviour described in VII.1, approximately one out of N_{RM} cells is a forward RM cell with CLP = 0. Under the assumption that each such cell is turned around by the destination, the source should provision enough resources to carry this RM traffic. The corresponding rate on the backward direction is proportional to the ACR in the forward direction. Therefore, the MCR on the backward direction should also accommodate $1/N_{RM} \times MCR_f$, where MCR_f is the MCR negotiated on the forward direction.

A smaller estimate of the requested MCR on the backward direction may be obtained under the assumption that fewer RM cells are looped back (i.e., RM cell consolidation). This may increase the control loop delay and decrease the control loop responsiveness in the forward direction.

Appendix VIII

2-store algorithm for determination of T(k) in ABR explicit mode

Clause 6.7.5.3 describes the ABR conformance definition, where 6.7.5.3.2 provides the reference algorithm for determination of $T(k)$. This appendix provides a simplified version of the reference algorithm where the number of stored combinations of $PACR(j)$ and $t_a(k)$ is limited to two. This 2-store algorithm is recognized to be less than optimal with respect to its tightness in order to reduce complexity.

VIII.1 2-store algorithm for determination of *T*(*k*) in ABR explicit mode

The algorithm is written in a format that determines the ACR at the interface as a continuous-time variable whose reciprocal at time $t_a(k)$ is $T(k)$; if ACR happens to be computed to a value less than 1 cell/s, $T(k)$ is set to 1 s. Thus, even if ACR is computed to be < 1 cell/s, the algorithm may identify all cells conforming if the rate is not higher than 1 cell/s. The increment for the DGCRA at each cell arrival in the forward direction is thus determined.

NOTE – At a given time instant, the ACR that is valid at the interface may differ from the ACR considered by the source to be valid. This may be for example because of a time lag, or because some backward RM cells considered at the interface have not reached the source.

The following algorithm computes two sets of counters (t_first,PACR_first) and (t_last,PACR_last). PACR_max is an auxiliary variable defined as Max(PACR_first,PACR_last).

- t_first is the time at which $T(k)$ is scheduled to be set to 1/PACR_first.
- If different from t_first, t_last is the planned update for t_first, at expiration of t_first; at that time PACR_first is updated to PACR_last.

PACR_first and PACR_last are determined on the basis of the value $PACR(j)$ of the ECR field carried in pertinent RM cells. Pertinent RM cells are backward RM cells with correct CRC-10 in the EDC field (see 8.1) that are either non-BECN cells, or BECN cells with ECR<PACR_last.

The algorithm described below has the following characteristics:

- At most, two rate modifications can be scheduled, which can be either increases or decreases from the current ACR.
- Since t_first, t_last, PACR_first and PACR_last are potentially updated every time a backward RM cell is observed in the backward direction, a given value of PACR_first or PACR_last may never be used in the DGCRA since, prior to its scheduled time to be applied, it can be revised by another backward RM cell.
- If less than two rate updates are scheduled, t_first=t_last and PACR_first=PACR_last.
- If no rate update is scheduled, PACR_first=PACR_last=ACR and t_first=t_last < tb(j).
- If at least one rate update is scheduled (PACR_first ≠ ACR), t_first cannot be delayed by a later rate update, and PACR_first can only be increased.
- At any time, PACR_last carries the ECR value of the last pertinent cell that has crossed the interface.
- If the ECR of a new pertinent cell is equal to PACR_last, no update takes place.
- $MCR \leq PACR\_first \leq PCR$ and $MCR \leq PACR\_last \leq PCR$.
- $tb(j) \leq t\_first \leq t\_last \leq tb(j) + \tau_2$ if least one rate update is scheduled.
- If $ACR<PACR\_first$, $t\_first \leq tb(j) + \tau_3$.
- If $PACR\_first<PACR\_last$, $t\_last \leq tb(j) + \tau_3$.

VIII.2 Adjustment of ACR(t) based on the ECR field in backward RM cells

- **Initialization:**
  t_first=t_last=0
  PACR_max=PACR_first=PACR_last=IACR

- **At each tb(j) which is the arrival time of a pertinent RM cell:**

```
compute PACR(j) = min( PCR, max(MCR, ECR in backward RM cell) )
if PACR(j) ≠ PACR_last:                 # else no update takes place
     if (t_first>tb(j))                 # is the scheduling list non-empty?
     # start update of a non-empty scheduling list
          if (PACR(j)≥ PACR_max)        # PACR (j) is an increase
                                        # over the current PACR_max
          # start processing an increase
               PACR_max=PACR(j)         # update PACR_max

               if (tb(j)+τ3>t_first)    # t_first and PACR_first unchanged
                    if ((t_first=t_last) or (t_last>tb(j)+τ3))
                         t_last=tb(j)+τ3     # else t_last in unchanged
                    endif
               endif                    # endif (tb(j) +τ3>t_first)

               else                     # tb (j) +τ3≤t_first
                    PACR_first=PACR(j)  # update PACR_first
                    if (PACR(j)≥ACR)    # PACR(j) is an increase over ACR
                         t_first=tb(j)+τ3    # else t_first is unchanged.
                    endif               # endif (PACR(j)≥ACR)
                    t_last=t_first      # a single rate update is scheduled
               endelse                  # endelse (tb (j) +τ3≤t_first)
          endif
          # end processing an increase
          else                          # PACR(j) is a decrease
                                        # over PACR_max
          # start processing a decrease
               PACR_first=PACR_max      # schedule highest rate at t_first
               if (PACR(j)<PACR_last)   # PACR(j) is a decrease
                                        # over PACR_last
                    t_last=tb(j)+τ2     # t_last is delayed
               endif                    # else t_last is unchanged
          endelse
          # end processing a decrease
          PACR_last=PACR(j)             # store new rate in PACR_last
     endif
     # end update of a non-empty scheduling list
     else                               # the scheduling list is empty
     # start update of an empty scheduling list
          if (PACR(j)>ACR)
               t_first=tb(j)+τ3         # an increase is scheduled (τ3 lag)
          else
               t_first=tb(j)+τ2         # a decrease is scheduled (τ2 lag)
          t_last=t_first                # a single rate update is scheduled
          PACR_max=PACR_first=PACR_last=PACR(j)
     endelse
     # end update of an empty scheduling list
endif                                   # endif for PACR (j) ≠PACR_last
```

- **At expiration of t_first:**

```
          ACR=PACR_first                # update ACR
          t_first=t_last                # update t_first
          PACR_first=PACR_last          # update PACR_first
```

```
          PACR_max=PACR_last
End of Adjustment of ACR(t) based on the ECR field in backward RM cells.

Begin determination of T(k)
     Initialize:
          T(1) = 1/IACR;
     At each arrival time ta(k) of a cell for k ≥ 2:
          If (ACR(k) < 1 cell/s)  then  T(k) = 1  else  T(k) = 1/ACR(k)
End determination of T(k).
```

Appendix IX

Applicability of ATM transfer capabilities to applications

This appendix illustrates the applicability of ATCs and QoS classes through examples of applications. A number of applications are listed with possible choices of ATC and QoS class. The selection of applications and ATC-QoS class combinations is intended as an illustration by example. It does not intend to be complete, nor does it exclude other choices. The specific properties of applications determine the combinations of ATC and QoS class that are suitable to support their demands.

Where rate parameters are listed, their associated tolerances are also relevant. See Table IX.1.

Table IX.1/I.371 – Examples of applications, ATCs, parameters and QoS classes

Example application	I.371 ATM transfer capability	QoS class	Transfer capability parameters	Remarks
Circuit emulation	DBR	QoS class 1	PCR	Committed low cell loss ratio for the duration of the connection when all cells are conforming to the relevant conformance tests.
Real-time Audio/Video communication	DBR	QoS class 1	PCR	For example, videophone or videoconference (assuming no graceful degradation capability).
	SBR1	QoS class 1	PCR, SCR/IBT	Encoder adapts the encoding so that the rate conforms to the SBR traffic descriptor.
Delay sensitive applications that produce variable bit rate traffic	DBR	QoS class 1	PCR	
Non real-time audio/video communications (requiring low cell loss)	SBR1	QoS class 2	PCR, SCR/IBT	
	DBR	QoS class 2	PCR	
Support of SMDS, FMBS, IP	SBR1 SBR2 SBR3	QoS class 2 QoS class 3 QoS class 3	PCR, SCR/IBT	

Table IX.1/I.371 – Examples of applications, ATCs, parameters and QoS classes

Example application	I.371 ATM transfer capability	QoS class	Transfer capability parameters	Remarks
Video-on-Demand (limited receive buffer)	DBR	QoS class 1	PCR	PCR = maximum required bandwidth of the application.
	ABT/DT	QoS class 1	PCR, SCR/IBT	Committed low cell loss ratio for the duration of each ATM block if all cells are conforming to the relevant conformance tests (equivalent to piecewise DBR). As an example, the transfer of each scene will be preceded by the request for PCR of the scene, determined through prior off-line analysis.
	SBR1	QoS class 1	PCR, SCR/IBT	
Elastic Audio/Video communication	ABR	QoS class 3	PCR, MCR	Minimum guaranteed bandwidth may be for example bandwidth needed for voice. Assumes the application accepts service fallback and degradation of video.
	ABT/DT (elastic mode)	QoS class 2	PCR, /IBT	Committed low cell loss ratio during each block (delimited partly by the network). SCR can represent minimum bandwidth as long as network does not set BCR<SCR (network specific policy).
File transfer Image transfer	ABR	QoS class 3	PCR, MCR	Minimum Cell Rate needed to maintain end point peer-to-peer protocol or to comply with a maximum transfer delay for the whole file.
Database enquiry	SBR2/SBR3	QoS class 3	PCR, SCR/IBT	For example, access to existing internet applications.

Table IX.1/I.371 – Examples of applications, ATCs, parameters and QoS classes

Example application	I.371 ATM transfer capability	QoS class	Transfer capability parameters	Remarks
Virtual Private Network (transport of any traffic through a VP network)	DBR	QoS class 1	PCR	The private network (re-)negotiates the VPCs. The way the traffic is organized within the VP is left to the private network. PCR = maximum required bandwidth.
	ABT/DT	QoS class 1	PCR, SCR/IBT	The private network negotiates the VPCs. The BCR of the VP is dynamically renegotiated using RM cells. The way the traffic is organized within the VP is left to the private network.
	SBR1	QoS class 1	PCR, SCR/IBT	The private network (re-)negotiates the VPCs. The way the traffic is organized within the VP is left to the private network.
IP support (edge-to-edge router connection)	DBR	QoS class U	PCR	IP traffic between two routers is put into a DBR VCC with QoS class U. Low cost best-effort service without minimum throughput and QoS support.

Appendix X

Additional material related to the F-GCRA

The following is a list of additional comments on the F-GCRA which are added to aid the reader's understanding of the behaviour of F-GCRA.

X.1 The support of QoS commitments through F-GCRA

For deriving the QoS commitments, the GFR ATC uses the Frame based Generic Cell Rate Algorithm F-GCRA(T,τ) defined in 6.8.3.2. The GFR ATC provides a QoS commitment in terms of a low cell loss ratio for at least the number of cells in conforming frames.

One expects that if the tolerance of the F-GCRA increases, the total number of cells in conforming frames will not decrease. However, this is not always true when the frames have different lengths. This will be shown by way of an example. Results will be stated under which conditions this unexpected phenomenon disappears.

Example

The following example shows that an increase of the tolerance of the F-GCRA may reduce the total number of cells in conforming frames if frames have different lengths. In the following, we will assume that all cells in CLP = 0 frames are conforming.

In the first part of the example, assume that the tolerance of the F-GCRA is $\tau = \tau_{IBT} + \tau_{MCR}$ and in the second part of the example assume that the tolerance is $\tau' = \tau_{IBT} + \tau'_{MCR}$ where τ'_{MCR} is larger than τ_{MCR}. For both parts of the example $T = 1/$ MCR.

Assume that before the arrival of the first cell of a frame, X' was always less than or equal to τ for the previous cells. This means that the values of X' in both parts of the example were the same until now. Assume that for a GFR connection, a short frame (frame length = 1) arrives that is followed by a long frame with length MFS >> 1. Assume that the F-GCRA parameter X' at the arrival of the first cell of the short frame is just slightly above τ but still below τ'.

Therefore, in the first part of the example the short frame would not pass the F-GCRA frame test. In that case it may happen that the subsequent long frame would pass.

On the other hand, in the second part of the example the short frame passes the F-GCRA frame test, but it may happen that the subsequent long frame would not pass.

As a result, out of the MFS + 1 arriving cells there are MFS cells in frames that passed in the first part of the example and there is one cell in frames that passed in the second part of the example. Therefore, the case of increased tolerance produced less cells in frames that passed. This is not what one expects.

Result

The following result shows that this unexpected phenomenon disappears if the increase in tolerance is "sufficiently" large. The proof of the result can be found in X.4.

The number of cells in conforming frames is determined with two F-GCRAs on the same cell flow: a reference F-GCRA(T,τ) and a second F-GCRA(T',τ'). Define the capacity of the reference F-GCRA as $C = 1 + \tau/T$ and the capacity of the second F-GCRA as $C' = 1 + \tau'/T'$.

If $T' \leq T$ and $C' \geq C +$ MFS, then the total number of cells in conforming frames as determined by the second F-GCRA is at least as large as the total number of cells in conforming frames as determined by the reference F-GCRA.

Consequences of the result

If the F-GCRA is not implemented with the exact parameters (T,τ) but with the parameters (T',τ') then the undesirable phenomenon described above will not arise if the parameters T' and τ' are chosen so that $T' \leq T$ and $\tau'/T' \geq \tau/T +$ MFS. Using such parameters (T',τ') may result in an increase by MFS cells of the buffer space to be reserved in a network element for the connection.

X.2 Example implementation illustrating how F-GCRA may be used to support QoS commitments

In this example a possible GFR implementation is described to show the relationship between F-GCRA and cell forwarding decisions taken in the implementation.

– A QoS reference counter is used per GFR connection. It is set to zero at the arrival of the first cell of the connection.
– At the arrival of the last cell of a conforming frame, the QoS reference counter is increased by the number of cells of the frame.
– When the last cell of a CLP = 0 frame of which all cells are conforming leaves the implementation, the QoS reference counter is decreased by the number of cells of the frame, but it is never decreased below zero.
– It is anticipated that the QoS reference counter returns to zero very often. This would mean that the implementation would provide at least the committed QoS for the GFR connection.

– It is anticipated that even when the QoS reference counter is zero, frames may leave the implementation if excess resources are available.

X.3 Implementation limits for the case of many non-conforming frames

The variables X' and X may grow beyond any limit in the case that many non-conforming frames are sent. In an implementation of the F-GCRA, X' and X have to be limited so they will not exceed a network specific value. For any connection, to have committed QoS in the network, its $\tau + T \times \text{MFS}$ should not exceed that network specific value.

X.4 Proof of a result related to F-GCRA

The following result and its proof were mentioned in X.1.

QoS result

Assume that QoS is determined by two frame-based GCRAs on the same cell flow. A reference F-GCRA(T,τ) and a second F-GCRA(T',τ'). Define the capacity of the reference F-GCRA as $C = 1 + \tau/T$ and the capacity of the second F-GCRA as $C' = 1 + \tau'/T'$. If $T' \leq T$ and $C' \geq C + \text{MFS}$ then $\text{QoS\_count}_n \leq \text{QoS\_count}_n'$ if the cell n is the last cell of a frame. Here QoS_count_n is the number of cells in conforming frames that passed the reference F-GCRA out of the first n cells. $\text{QoS\_count}_n'$ is defined similarly for the second F-GCRA.

Proof

The proof is by induction on m where $n = n_m$ is the number of cells at the end of the frame m. For $n = n_1$, $\text{QoS\_count}_n \leq \text{QoS\_count}_n'$ is trivial since the first frame receives QoS by both the F-GCRAs or by none.

Now let $\text{QoS\_count}_n \leq \text{QoS\_count}_n'$ for $n = n_m$. Since CLP = 1 frames bypass the F-GCRA, one may assume that all frames arriving at the Frame based GCRA start with CLP = 0 cells. Then, once the last cell of frame $m + 1$ has been processed by the F-GCRAs, there is only something to show if the frame $m + 1$ is conforming to the reference F-GCRA and not conforming to the second F-GCRA. Let the first cell of this frame be the cell $j = n_m + 1$. At the arrival of cell j, for the variable X'_j for the reference F-GCRA the following relationship holds: $X'_j \leq \tau$. For the corresponding variable X''_j for the second F-GCRA the following relationship holds: $X''_j > \tau'$.

For $k = 1, 2, \dots n$, let t_k be the arrival time of the cell k. Define X'_k and for all cells up to cell n, even for cells in frames that did not pass the frame test. For all the cells up to cell n, the last cell of the frame has arrived so that the cell conformance information about the cells in the frame is available. For the reference F-GCRA, set $X'_k = X - (t_k - LIT\_1)$ for frames where the incrementing is undone and for frames where no incrementing takes place. Similarly, define X''_k for the second F-GCRA. In addition, define $Y'_k = \max(X'_k, 0)$ and $Y''_k = \max(X''_k, 0)$.

For a cell $k \leq n$, the QoS_count_k can be defined naturally as follows: if the corresponding frame contains one or more non-conforming cells or is not conforming, the value QoS_count_k is the same as at the end of the previous frame. Otherwise it increases by one for every cell of the frame. Similarly for $\text{QoS\_count}_k'$. With this definition one also gets that $\text{QoS\_count}_k \leq \text{QoS\_count}_k'$ for $k = 1, \dots, n$.

Note that $Y''_j/T' = X''_j/T' > \tau'/T' = C'-1 \geq C-1 + \text{MFS} = \tau/T + \text{MFS} \geq Y'_j/T + \text{MFS}$ and thus:

$$Y''_j/T' - Y'_j/T > \text{MFS} \quad \text{(X-1)}$$

Let cell i be the last cell arriving before cell j such that $Y''_i = 0$. Then $1 \leq i < j$. Then:

$$Y''_i/T' - Y'_i/T \leq 0 \quad \text{(X-2)}$$

One then gets the following:

– For every cell of a conforming frame according to the second F-GCRA and not to the reference F-GCRA, the incrementing will result in an increase of the difference $Y''T' - Y'T$ by one.
– For every cell of a conforming frame according to the reference F-GCRA and not to the second F-GCRA, the incrementing will result in a decrease of the difference $Y''T' - Y'T$ by one.
– For all other cells of frames of which all cells are conforming, the incrementing will not change the difference $Y''T' - Y'T$.
– For every cell of a frame of which not all cells are conforming, the incrementing will result in an increase of $Y'T$ by one and in an increase of $Y''T$ by one. This will not result in an increase of $Y''T' - Y'T$.

Also since $Y'' > 0$ from cell $i + 1$ to j, and since $T \geq T'$, $Y''T'$ is decremented from cell to cell at most as much as $Y'T$. This means that the decrementing does not increase the difference $Y''T' - Y'T$.

Therefore inequalities (X-1) and (X-2) show that of the cells i up to $j - 1$, the second F-GCRA has found at least MFS cells more in conforming unmarked frames that passed the frame test than the reference F-GCRA. Therefore $QoS\_count_{j-1} + MFS \leq QoS\_count_{j-1}'$ or $QoS\_countn_m + MFS \leq QoS\_countn_m'$. Since the frame $m + 1$ is conforming, it has a size of at most MFS cells and one gets $QoS\_count_n \leq QoS\_count_n$ for $n = n_{m+1}$ which finishes the proof.

Appendix XI

Providing GFR QoS with the CF-GCRA

This appendix contains an algorithm called the cell conforming F-GCRA (CF-GCRA). The CF-GCRA is equivalent to the F-GCRA algorithm for connections containing *only* frames of which all cells are conforming. It is simpler than the F-GCRA and can also be used to provide GFR QoS because of the following: under the condition that all cells are conforming, it can be shown (the proof is similar to the proof in X.4), that the number of cells in frames that pass the CF-GCRA is at least as large as the number of cells in frames that pass the F-GCRA. This assumes that the CF-GCRA is not implemented with the exact parameters (T,τ) but with the parameters (T',τ') where $T' \leq T$ and $\tau'/T' \geq \tau/T + \text{MFS}$. With these settings of T' and τ', the minimum QoS commitments based on F-GCRA are met.

In the CF-GCRA below:

– t_a denotes the arrival time of the latest cell at a standardized interface.
– X denotes the value of the Leaky Bucket counter, as in the continuous-state leaky bucket algorithm.
– *LIT* denotes the Last Incrementing Time.
– *Frame_test_passed* denotes a connection specific variable that stores the frame test result.
– *Frame_tagging* denotes a connection specific variable that is only used in GFR2. It stores the frame tagging status. If frame tagging for GFR is implemented, then this status information could be used to change the CLP bit from 0 to 1.
– X' is an auxiliary variable.

Initialization:

– At the time of arrival t_a of the first cell of the connection to cross the given interface, $X = 0$ and $LIT = t_a$.

– The initial values of *frame_test_passed* and *frame_tagging* are irrelevant.

At the arrival of the *first* cell of the frame at a given interface T_B or inter-network interface, on the ATM connection:

GFR1

```
if (CLP = 1)
    then frame_test_passed = false
    else
        X' = X – (ta – LIT)
        if (X' > τ)
            then frame_test_passed = false
            else frame_test_passed = true
            X = max(0,X') + T
                LIT = ta
```

GFR2

```
if (CLP = 1)
    then frame_test_passed = false;
        frame_tagging = false
    else
        X' = X – (ta – LIT)
        if (X' > τ)
            then frame_test_passed = false;
                frame_tagging = true
            else frame_test_passed = true;
                frame_tagging = false
            X = max(0,X') + T
                LIT = ta
```

At the arrival of *subsequent* cells of a frame at a given interface T_B or inter-network interface, on the ATM connection:

GFR1 and GFR2

```
if (frame_test_passed = true)
    then X' = X – (ta – LIT)
    X = max(0, X') + T
    LIT = ta
```

Appendix XII

Expectations of a GFR network element behaviour

From clause 6.8.1 on the GFR service model, some minimal GFR implementation requirements may be derived. The following contains expectations on how a network element could support GFR in order to improve the GFR service.

– If the connection transmits a mixture of CLP = 0 and CLP = 1 frames at a constant overall cell rate which is below MCR and assuming that all cells are conforming, then the network element should deliver all of the connection's frames.

– If the connection transmits CLP = 0 frames at a cell rate lower than MCR and in addition CLP = 1 frames such that the overall cell rate is higher than the MCR and assuming that all cells are conforming, then the network element should deliver all the conforming frames (the CLP = 0 frames, as part of the commitments) and, in addition, deliver an overall rate of at least MCR for that connection.

Appendix XIII

Applicability of GFR ATM Transfer Capability to applications

This appendix extends Table IX.1, which illustrates the applicability of ATCs and QoS classes through examples of applications, with an example application for the GFR ATC defined in this Recommendation. Where rate parameters are listed, their associated tolerances are also relevant. See Table XIII.1.

Table XIII.1/I.371 – Examples of applications, ATCs, parameters and QoS classes

Example application	ATM transfer capability	QoS class	Transfer capability parameters	Remarks
IP support (edge-to-edge router connection)	GFR	QoS class 3 for conforming frames	PCR, MCR/IBT, MFS	IP traffic between two routers is put into a GFR VCC. Minimum throughput and QoS support and frame discard.

Appendix XIV

OAM support for GFR connections

This appendix shows how OAM support for a GFR connection is possible based on the GFR text in the main body of this Recommendation.

Note that the GFR definition assumes that the user generated data cells on a GFR connection are organized in the form of frames that are delineated at the ATM layer. Note also that from clauses 4.1 and 7.2.5, it follows that user generated end-to-end OAM cells inserted into a GFR connection would be considered as part of the GFR frames by the GFR conformance definition:

Extract from 4.1:

…QoS commitments are referred to where the network actually commits to meet upper bounds for some QoS parameters, assuming the user generated cell flow conforms to a traffic contract. … Segment OAM flows are not part of the traffic contract negotiated by the user. How they are handled is currently not specified in this Recommendation.

Extract from 7.2.5:

…For a VCC, the AUU indication may be used (e.g., as specified for AAL 5) to define a frame cell sequence as follows:

– *A frame cell sequence starts with the first user generated cell on the connection or with a user generated cell following a cell with the AUU indication set.*

– *A frame cell sequence terminates with a user cell that has the AUU indication set.*

If insertion of OAM cells is desired on a GFR connection, the following principles help to reduce or avoid cell conformance or frame conformance problems:

- If an OAM cell is to be inserted into the middle of a GFR frame, the OAM cell should receive the same CLP bit as the first cell of the frame.

- If an OAM cell is to be inserted after a cell with the AUU delineation, the CLP bit of the next frame may not be known. In this case setting the CLP bit of the OAM cell to 0 minimizes the OAM cell discard even though it may result in non-conformance of the next frame provided this frame turns out to be a CLP = 1 frame.
- If the user inserts OAM cells, the user should increase the value of MFS by one above the value of MFS needed for support of GFR user frames consisting entirely of user data cells.
- If the network inserts OAM cells into a GFR user connection, the value of MFS used for testing conformance should be increased by at least one above the value of MFS requested by the user. If the network inserts OAM cells into a network internal GFR connection the network provider should increase the value of MFS by one above the value of MFS needed for support of GFR user frames consisting entirely of user data cells.

NOTE – Since it can be expected that most of the frames in a GFR connection are frames with all CLP = 0 cells, inserting OAM cells with CLP = 0 can be usually achieved without frame conformance violation by inserting them into frames with all CLP = 0 cells.

ERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	**Integrated services digital network**
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communications
Series Y	Global information infrastructure, Internet protocol aspects and Next Generation Networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2006