

Superseded by a more recent version



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

I.371

(08/96)

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
Superseded by a more recent version

(Previously CCITT Recommendation)

Superseded by a more recent version

ITU-T I-SERIES RECOMMENDATIONS INTEGRATED SERVICES DIGITAL NETWORK

GENERAL STRUCTURE	I.100–I.199
Terminology	I.110–I.119
Description of ISDNs	I.120–I.129
General modelling methods	I.130–I.139
Telecommunication network and service attributes	I.140–I.149
General description of asynchronous transfer mode	I.150–I.199
SERVICE CAPABILITIES	I.200–I.299
Scope	I.200–I.209
General aspects of services in ISDN	I.210–I.219
Common aspects of services in the ISDN	I.220–I.229
Bearer services supported by an ISDN	I.230–I.239
Teleservices supported by an ISDN	I.240–I.249
Supplementary services in ISDN	I.250–I.299
OVERALL NETWORK ASPECTS AND FUNCTIONS	I.300–I.399
Network functional principles	I.310–I.319
Reference models	I.320–I.329
Numbering, addressing and routing	I.330–I.339
Connection types	I.340–I.349
Performance objectives	I.350–I.359
Protocol layer requirements	I.360–I.369
General network requirements and functions	I.370–I.399
ISDN USER-NETWORK INTERFACES	I.400–I.499
Application of I-series Recommendations to ISDN user-network interfaces	I.420–I.429
Layer 1 Recommendations	I.430–I.439
Layer 2 Recommendations	I.440–I.449
Layer 3 Recommendations	I.450–I.459
Multiplexing, rate adaption and support of existing interfaces	I.460–I.469
Aspects of ISDN affecting terminal requirements	I.470–I.499
INTERNETWORK INTERFACES	I.500–I.599
MAINTENANCE PRINCIPLES	I.600–I.699
B-ISDN EQUIPMENT ASPECTS	I.700–I.799
ATM equipment	I.730–I.749
Management of ATM equipment	I.750–I.799

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

Superseded by a more recent version

ITU-T RECOMMENDATION I.371

TRAFFIC CONTROL AND CONGESTION CONTROL IN B-ISDN

Summary

The I.371 set of Recommendations address 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. 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 revised by ITU-T Study Group 13 (1993-1996) and was approved under the WTSC Resolution No. 1 procedure on the 27th of August 1996.

Superseded by a more recent version

FOREWORD

ITU (International Telecommunication Union) is the United Nations Specialized Agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of the ITU. The 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 Conference (WTSC), which meets every four years, establishes the topics for study by the ITU-T Study Groups which, in their turn, produce Recommendations on these topics.

The approval of Recommendations by the Members of the ITU-T is covered by the procedure laid down in WTSC Resolution No. 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.

INTELLECTUAL PROPERTY RIGHTS

The ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. The 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, the ITU had/had not 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.

© ITU 1997

All rights reserved. No part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the ITU.

Superseded by a more recent version

CONTENTS

	Page
1	Scope..... 1
2	References..... 1
3	Abbreviations and terminology 2
4	Introduction..... 5
4.1	General objectives 5
4.2	Generic functions..... 6
4.3	A reference configuration for traffic control and congestion control..... 6
4.4	Events, actions, time-scales and response times..... 7
4.5	Quality of service and network performance..... 8
5	Traffic parameters and descriptors 9
5.1	Definitions 9
5.1.1	Traffic parameters..... 9
5.1.2	Traffic descriptors..... 9
5.2	Requirements 10
5.3	Traffic contract 10
5.3.1	Traffic contract definition..... 10
5.3.2	Traffic contract and quality of service..... 11
5.3.3	Traffic contract and cell loss priority 11
5.3.4	Traffic contract and tagging option 11
5.3.5	Impact of cell delay variation on UPC/NPC and resource allocation 12
5.4	Traffic parameter specifications 14
5.4.1	Peak cell rate..... 16
5.4.2	Sustainable cell rate 19
5.5	ATM transfer capabilities 20
5.5.1	General..... 20
5.5.2	Applicability of ATM transfer capabilities to applications..... 22
5.5.3	Deterministic bit rate transfer capability (DBR) 22
5.5.4	Statistical bit rate transfer capability (SBR) 26
5.5.5	ATM block transfer capabilities (ABT) 32
5.5.6	Available bit rate transfer capability (ABR) 40
6	Functions for traffic control and congestion control 46
6.1	Introduction..... 46
6.1.1	Traffic control and congestion control functions 46
6.2	Traffic control functions 47

Superseded by a more recent version

	Page
6.2.1 Use of virtual paths for network resource management	47
6.2.2 Connection admission control	49
6.2.3 Usage parameter control and network parameter control.....	50
6.2.4 Priority control.....	57
6.2.5 Traffic shaping.....	57
6.2.6 Fast Resource Management (FRM).....	58
6.3 Congestion control functions	58
6.3.1 Selective cell discard	58
6.3.2 Explicit forward congestion indication	58
6.3.3 Reaction to UPC/NPC failures	59
6.4 Traffic control interworking functions	59
6.4.1 Traffic control interworking with FMBS	59
7 Procedures for traffic control and congestion control	60
7.1 Resource management cell format.....	60
Annex A – Generic Cell Rate Algorithm GCRA(T,τ).....	61
Annex B – Application of the GCRA to SBR conformance definition	63
Annex C – ABT/DT control messages across a standardized interface	67
Annex D – ABT/IT control messages across a standardized interface.....	69
Annex E – ABR control messages across a standardized interface.....	70
Annex F – Traffic characteristics relevant to ATCs	72
Appendix I – Examples of application of the equivalent terminal for the peak cell rate definition.....	73
Appendix II – Transcoding rules from signalling information onto OAM traffic parameters at the ATM layer.....	76
Appendix III – Throughput behaviour of the generic cell rate algorithm.....	77
Appendix IV – UPC/NPC accuracy requirements.....	78

Superseded by a more recent version

Recommendation I.371

TRAFFIC CONTROL AND CONGESTION CONTROL IN B-ISDN

(revised in 1996)

1 Scope

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

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

- The main body describes the objectives and mechanisms of traffic control and congestion control.
- Examples of application of the concepts are in Annexes A, B, C, D, E and F.

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 the negotiated network performance objectives 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.

ATM layer traffic control refers to the set of actions taken by the network to avoid congested conditions.

ATM layer congestion control refers to the set of actions taken by the network to minimize the intensity, spread and duration of congestion. These actions are triggered by congestion in one or more network elements.

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; all 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.

- [1] ITU-T Recommendation I.326 (1995), *Functional architecture of transport networks based on ATM*.
- [2] ITU-T Recommendation I.113 (1993), *Vocabulary of terms for broadband aspects of ISDN*.
- [3] ITU-T Recommendation I.150 (1995), *B-ISDN asynchronous transfer mode functional characteristics*.
- [4] ITU-T Recommendation I.311 (1996), *B-ISDN general network aspects*.

Superseded by a more recent version

- [5] CCITT Recommendation I.321 (1991), *B-ISDN protocol reference model and its application*.
- [6] ITU-T Recommendation I.356 (1996), *B-ISDN ATM layer cell transfer performance*.
- [7] ITU-T Recommendation I.361 (1995), *B-ISDN ATM layer specification*.
- [8] ITU-T Recommendations I.363.1 (1996), I.363.3 (1996) and I.363.5 (1996), *B-ISDN ATM adaptation layer specification*.
- [9] ITU-T Recommendation I.413 (1993), *B-ISDN user-network interface*.
- [10] ITU-T Recommendations I.432.1 (1996), I.432.2 (1996), I.432.3 (1996) and I.432.4 (1996), *B-ISDN user-network interface – Physical layer specification*.
- [11] ITU-T Recommendation I.610 (1995), *B-ISDN operation and maintenance principles and functions*.
- [12] ITU-T Recommendation I.731 (1996), *Types and general characteristics of ATM equipment*.
- [13] ITU-T Recommendation I.732 (1996), *Functional characteristics of ATM equipment*.
- [14] ITU-T Recommendation Q.2650 (1995), *Interworking between Signalling System No. 7 Broadband ISDN User Part (B-ISUP) and Digital subscriber Signalling System No. 2 (DSS 2)*.
- [15] ITU-T Recommendation Q.2660 (1995), *Interworking between Signalling System No. 7 Broadband ISDN User Part (B-ISUP) and Narrow-band ISDN User Part (N-ISUP)*.
- [16] ITU-T Recommendation Q.2761 (1995), *Functional description of the B-ISDN User Part (B-ISUP) of Signalling System No. 7*.
- [17] ITU-T Recommendation Q.2762 (1995), *General functions of messages and signals of the B-ISDN User Part (B-ISUP) of Signalling System No. 7*.
- [18] ITU-T Recommendation Q.2763 (1995), *Signalling System No. 7 B-ISDN User Part (B-ISUP) – Formats and codes*.
- [19] ITU-T Recommendation Q.2764 (1995), *Signalling System No. 7 B-ISDN User Part (B-ISUP) – Basic call procedures*.
- [20] ITU-T Recommendation Q.2931 (1995), *Digital subscriber Signalling System No. 2 – User-Network Interface (UNI) layer 3 specification for basic call/connection control*.
- [21] ITU-T Recommendations Q.2961.1 (1995), and Q.2961.2 (1997), *Digital subscriber Signalling System No. 2 – Additional traffic parameters*.
- [22] ITU-T Recommendation Q.2962 (1996), *Digital subscriber Signalling System No. 2 – Connection characteristics negotiation during call/connection establishment phase*.
- [23] ITU-T Recommendation Q.2963.1 (1996), *Digital subscriber Signalling System No. 2 – Connection modification: Peak cell rate modification by the connection online*.

3 Abbreviations and terminology

This Recommendation uses the following abbreviations.

AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ABT	ATM Block Transfer

Superseded by a more recent version

ACR	Allowed Cell Rate
ATC	ATM Transfer Capability
ATM	Asynchronous Transfer Mode
ATM_PDU	ATM Protocol Data Unit
B-ISDN	Broadband ISDN
BCR	Block Cell Rate
B-NT1	Broadband Network Termination 1
B-NT2	Broadband Network Termination 2
B-TE	Broadband Terminal Equipment
BECN	Backward Explicit Congestion Notification
CAC	Connection Admission Control
CBR	Constant Bit Rate
CCR	Current Cell Rate
CDV	Cell Delay Variation
CEQ	Customer Equipment
CI	Congestion Indication
CRF(VC)	Virtual Channel connection-related Functions
CRF(VP)	Virtual Path connection-related Functions
CLP	Cell Loss Priority (bit)
CLR	Cell Loss Ratio
CRC	Cyclic Redundancy Check
CS	Convergence Sublayer
CT	Connection Termination
CTD	Cell Transfer Delay
DBR	Deterministic Bit Rate
DGCRA	Dynamic GCRA
DT	Delayed Transmission
ECR	Explicit Cell Rate
EDC	Error Detection Code
EFCI	Explicit Forward Congestion Indication
FIFO	First-In-First-Out
FMBS	Frame Mode Bearer Service
FRM	Fast Resource Management
GCRA	Generic Cell Rate Algorithm
GFC	Generic Flow Control
IBT	Intrinsic Burst Tolerance
INI	Inter-Network Interface
IT	Immediate Transmission
IWF	Interworking Function
LCT	Last Conformance Time

Superseded by a more recent version

MBS	Maximum Burst Size
MCR	Minimum Cell Rate
NE	Network Element
NI	No Increase
NPC	Network Parameter Control
NRM	Network Resource Management
OAM	Operation And Maintenance
PDU	Protocol Data Unit
PC	Priority Control
PCR	Peak Cell Rate
PEI	Peak Emission Interval
PHY	Physical Layer
PM	Performance Monitoring
PTI	Payload Type Indicator
QOS	Quality of Service
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
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
VPC	Virtual Path Connection
VPCT	Virtual Path Connection Termination
VPI	Virtual Path Identifier
VPLT	Virtual Path Link Termination
VSA	Virtual Scheduling Algorithm

The following terminology applies to this Recommendation:

– *On a virtual path connection*

User data cell

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.

Superseded by a more recent version

User OAM cell

Any F4 end-to-end OAM cell on the VPC generated by the user.

User RM cell

Any RM cell with VCI = 6 and PTI = 110 generated by the user.

User-generated cell

Any user data, user OAM or user RM cell.

– *On a virtual channel connection*

User data cell

Any CLP = 0 and CLP = 1 cell generated by the user, exclusive of F5 OAM cells and RM cells with PTI = 110.

User OAM cell

Any F5 end-to-end OAM cell on the VCC generated by the user.

User RM cell

Any RM cell with PTI = 110 on the virtual channel generated by the user.

User-generated cell

Any user data, user OAM or user RM cell.

How segment OAM cells are handled is currently not specified in this Recommendation.

In this Recommendation and in consistency with Recommendations 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 from the user. QOS commitments are referred to where the network actually commits to meet QOS objectives, assuming the user generated cell flow conforms to a traffic contract. QOS indications pertain when there is not such a traffic contract between the user and the network, e.g. in cases where traffic engineering rules are used to operate the network and do not allow for commitments to the user.

4 Introduction

4.1 General objectives

The primary role of traffic control and congestion control parameters and procedures is to protect the network and the user in order to achieve network performance objectives. An additional role is to optimize the use of network resources.

The uncertainties of broadband traffic patterns and the complexity of traffic control and congestion control suggest a step-wise approach for defining traffic parameters and network traffic control and congestion control mechanisms. This 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.

Superseded by a more recent version

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 network performance Recommendations.
- ATM layer traffic control and congestion control should not rely on AAL protocols which are B-ISDN service specific, nor on higher layer protocols which are application specific. Protocol layers above the ATM layer may make use of information which may be provided by the ATM layer to improve the utility those protocols can derive from the network.
- The design of an optimum set of ATM layer traffic controls and congestion controls should minimize network and end-system complexity while maximizing network utilization.

4.2 Generic functions

To meet these objectives, the following functions form a framework for managing and controlling traffic and congestion in ATM networks and may be used in appropriate combinations. 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 re-negotiation 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). Routing is part of connection admission control actions.
- 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 controls are functions that differentiate how cells are handled relative to each other by the network in terms of time priority or loss priority.

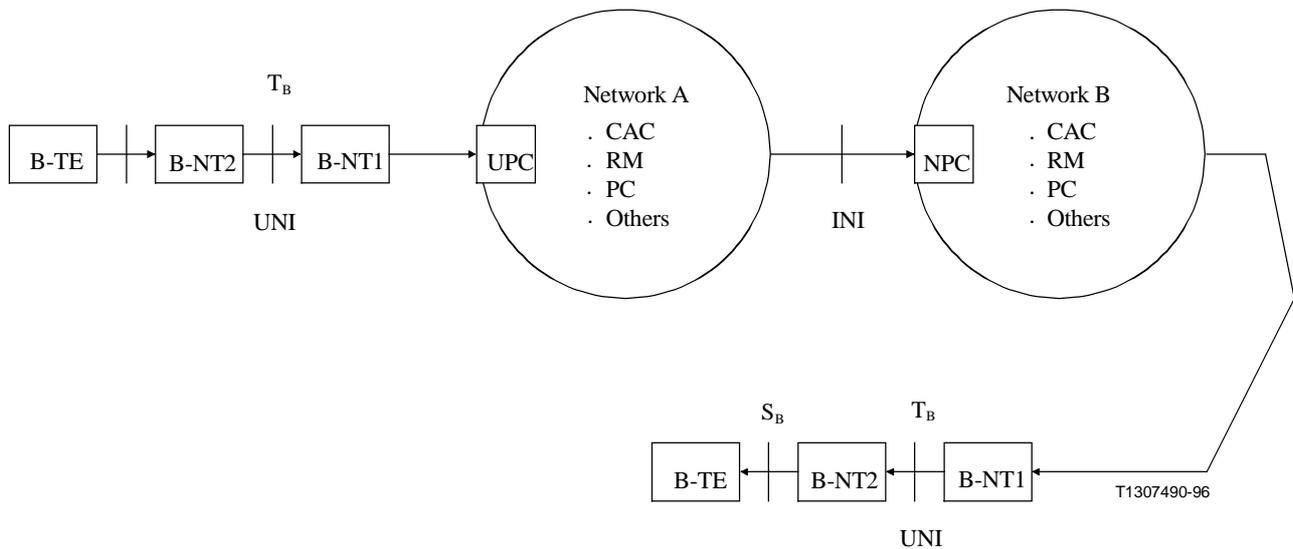
As a general requirement, it is desirable that a high level of consistency be achieved between the above traffic control capabilities.

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 5.5). This Recommendation includes a set of such capabilities which is intended to meet the requirements of different sets of broadband applications.

4.3 A reference configuration for traffic control and congestion control

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

Superseded by a more recent version



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

NOTE 1 – NPC may apply as well at some intra-network interfaces.

NOTE 2 – The arrows are indicating 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 Recommendation I.413). It is included in the figure for the sake of completeness and consistency with Recommendation I.413.

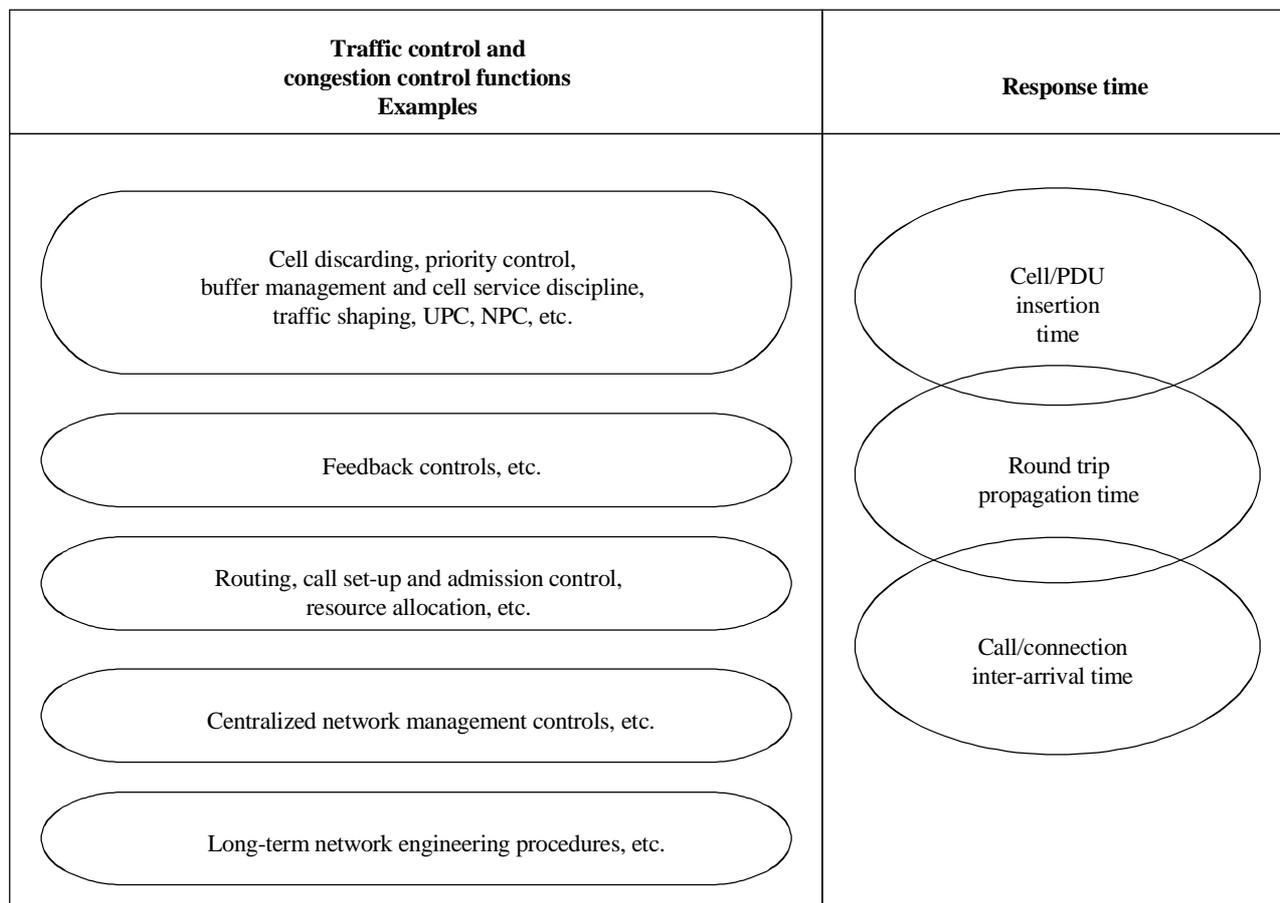
NOTE 5 – In this Recommendation, UNI refers to the interface at the T_B reference point.

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

4.4 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.

Superseded by a more recent version



T1307500-96

Figure 2/I.371 – Control response times

4.5 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 Recommendation 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 Recommendation 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 Recommendation I.356. When the conformance relative to a particular parameter is specified as being "unspecified", the ITU 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 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

Superseded by a more recent version

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).

Network performance objectives at the ATM SAP are intended to capture the network ability to meet the requested ATM layer quality of service. 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

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.

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 and may be derived from service type (e.g. telephone, videophone).

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

Service type can be used for implicit declaration by the user of a complete set of traffic parameters, e.g. by declaring the service requested (voice, etc.). Service type may also include implicit declaration of QOS requirements. Such a descriptor would be used for example as an address of a look-up table delivering the corresponding set of traffic characteristics. In case it is used by a traffic source, it would therefore not be necessary to convey any other traffic parameter belonging to the source traffic descriptor (see 5.1.2) via signalling. Additionally, service type may also be used to describe traffic characteristics of a source. This applies for example when typical source behaviours (e.g. variable bit rate video using standard coding schemes) are learned from operational experience or other means and used by network operators to apply specific traffic engineering rules, which may result in QOS indications rather than commitments. The use of service type is currently not specified in this Recommendation.

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).

Superseded by a more recent version

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 his 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 5.5), 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 5.5). 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 Recommendations I.311 and I.732) have to be declared.

Superseded by a more recent version

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

Recommendation 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 5.5.4) 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 Recommendation 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 5.5.3.3.

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 option

With regard to the tagging option (see 6.2.3.6):

- The user may request the tagging option to be applied to the CLP = 0 user generated traffic in excess. In that case, the user to network information will be "tagging requested". Tagging will be applied to the connection or not, depending on whether the network implements the tagging option. The network to user information should explicitly be "tagging applied" or "tagging not applied".
- The user may request the tagging option not to be applied to the CLP = 0 traffic in excess. In that case the user to network information will be "tagging not allowed". The only possible answer (implicit or explicit) from the network is "tagging not applied". In this case, tagging is disabled at the UPC/NPC and traffic in excess will be discarded. As a consequence, the CLP bit remains unchanged from source UNI to destination UNI.
- The default settings for tagging option are: from the user "tagging not allowed" and for the network "tagging not applied".

Superseded by a more recent version

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.

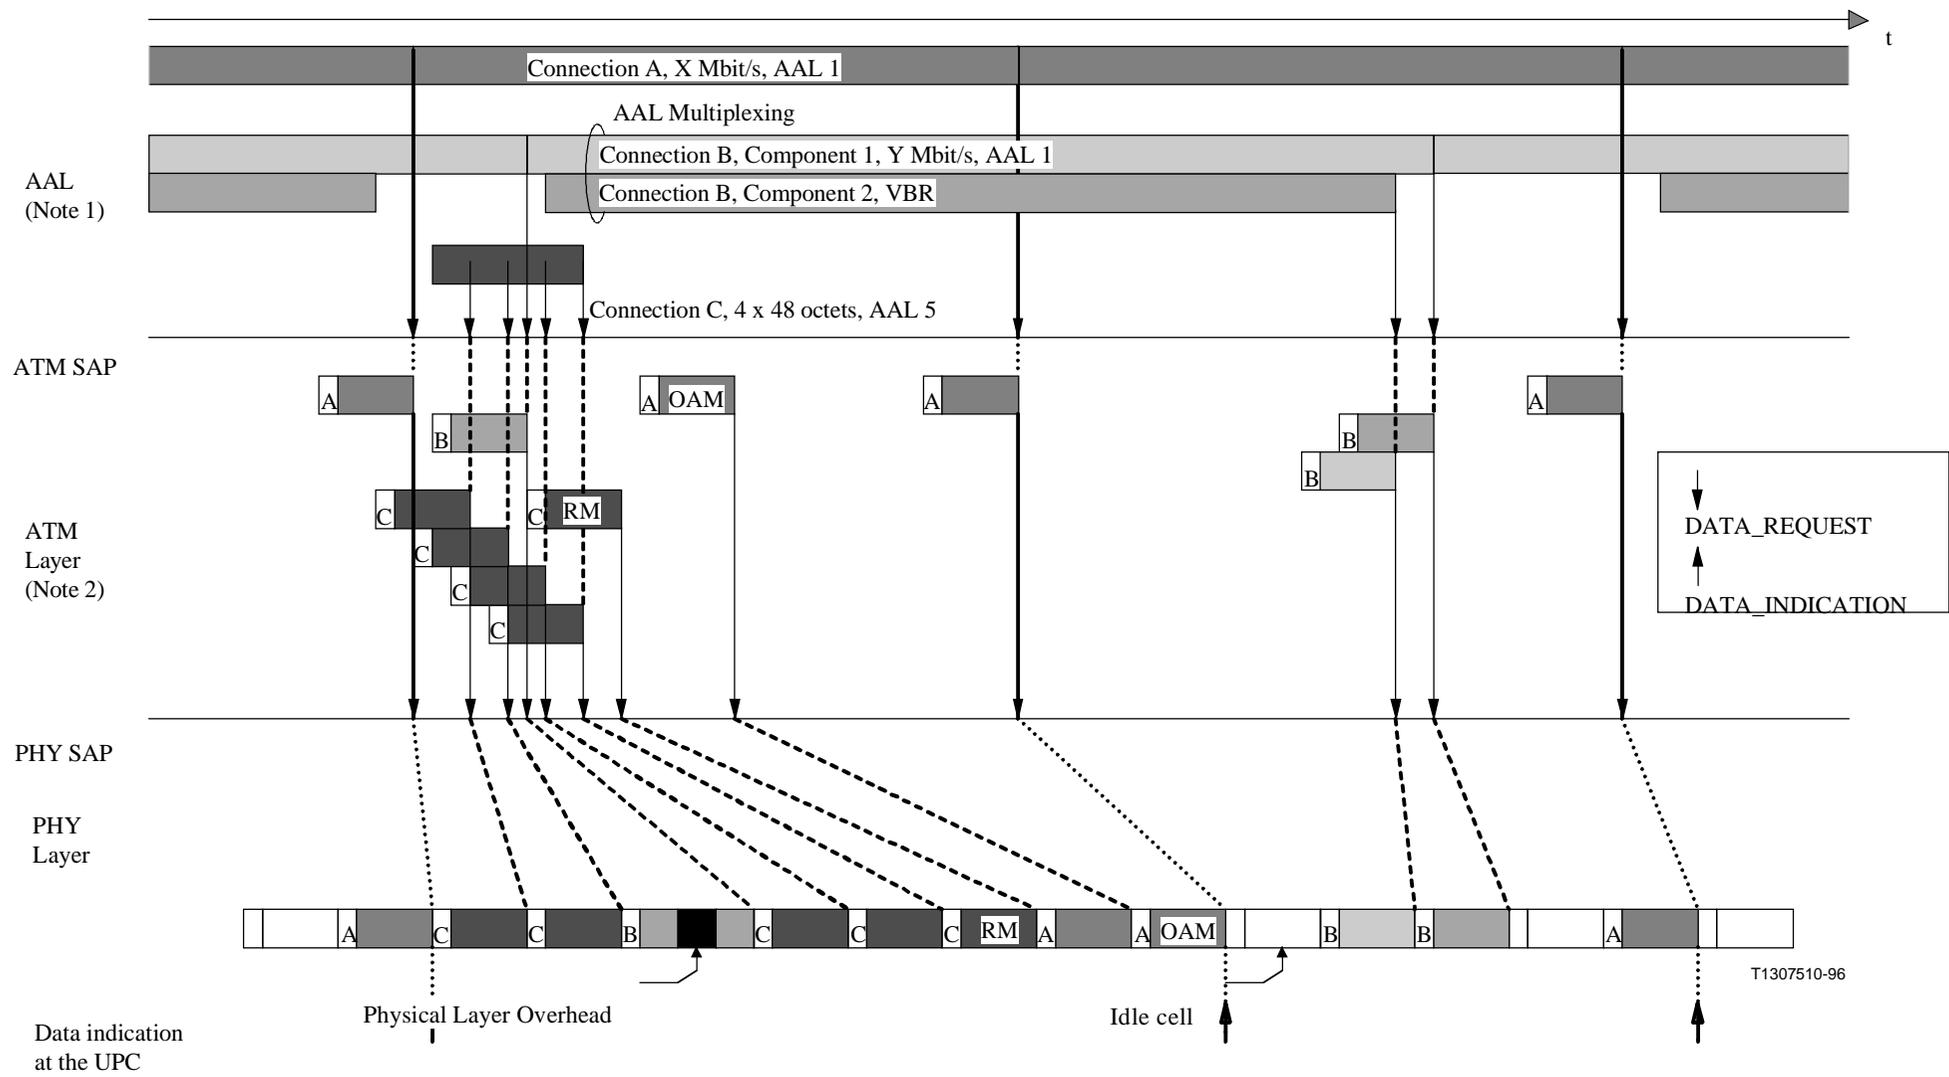
The conformance definition to the source traffic descriptor at a given interface (e.g. a UNI or an inter-network interface), as well as the execution of the UPC/NPC functions, require that the CDV attributed to the upstream portion of the connection and affecting each of the relevant parameters be specified.

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

Superseded by a more recent version

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 5.5).

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:

$$2^{e-32} \left(1 + \frac{w}{2^{10}} \right) \text{ seconds}$$
$$0 \leq e \leq 41$$
$$0 \leq w \leq 1023$$

Conformance to service type is currently not specified in this Recommendation.

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 5.5). 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.

Superseded by a more recent version

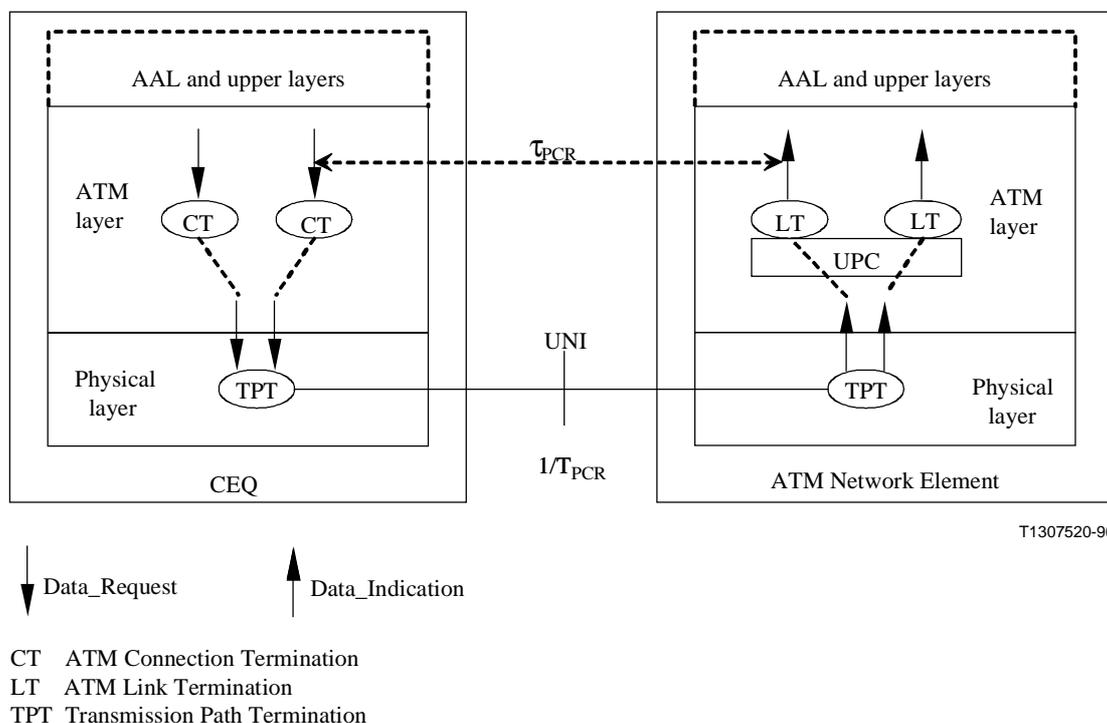


Figure 4/I.371 – Reference configuration for the specification of traffic parameters

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 Recommendations I.731 and I.732.

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).

Superseded by a more recent version

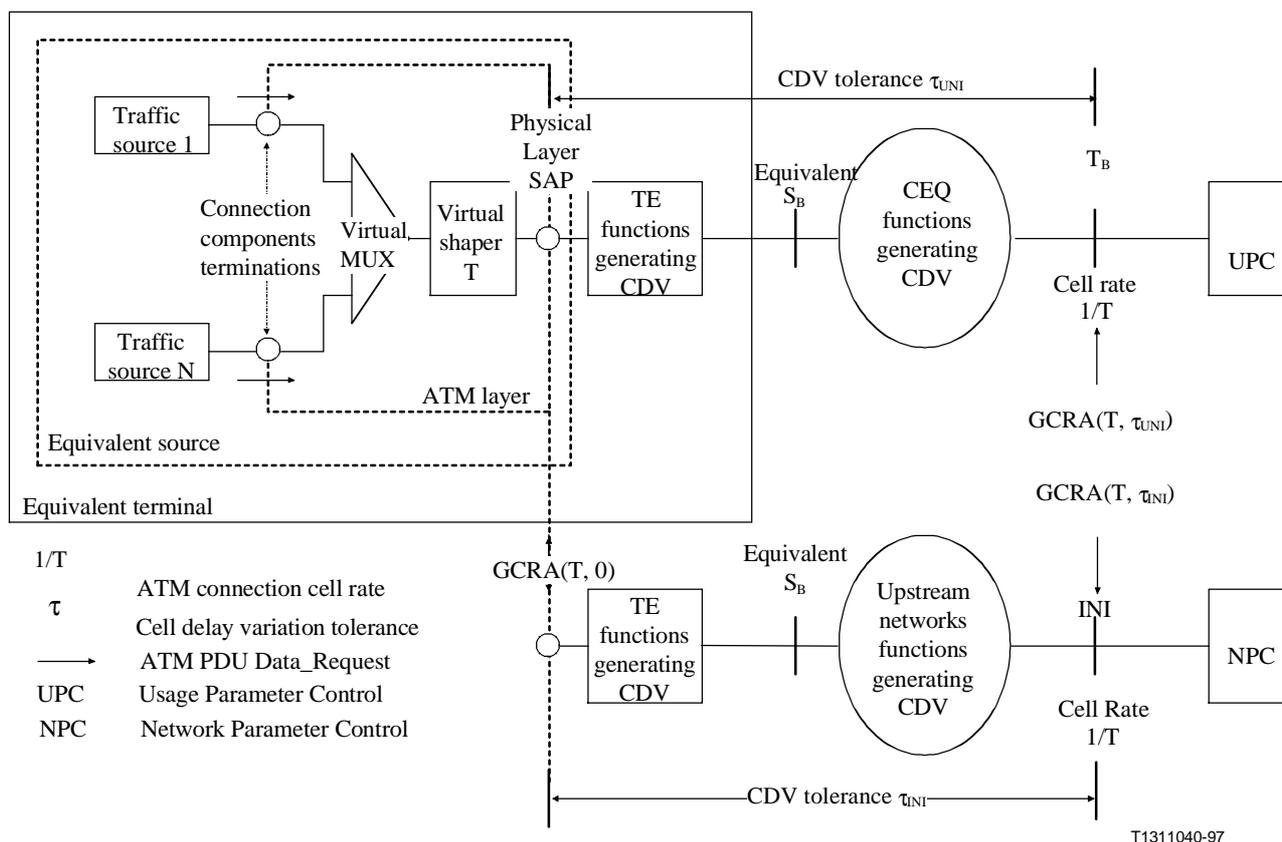


Figure 5/I.371 – Equivalent source and equivalent terminal for the definition of a cell rate of an ATM connection

(This Figure is illustrative only and does not imply any implementation.)

5.4.1 Peak cell rate

The following definition applies to ATM connections for any ATM transfer capability defined in 5.5.

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.

Superseded by a more recent version

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 5.5 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.

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%.

$$\Lambda_{PCR} = 2^{m_{PCR}} \cdot \left(1 + \frac{k_{PCR}}{512} \right) \text{ cells per second}$$
$$0 \leq m_{PCR} \leq 31$$
$$0 \leq k_{PCR} \leq 511$$

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 $2.33 \cdot 10^{-10}$ seconds. Relative difference between any pair (same values of m_{PCR} and k_{PCR}) of peak cell rate and $1/\text{peak emission interval}$ 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 ($x \rightarrow 1/x$).

Superseded by a more recent version

$$T_{PCR} = 2^{-(m_{PCR}+1)} \cdot \left(1 + \frac{1023 - k'_{PCR}}{1024}\right) \text{ seconds}$$
$$k'_{PCR} = \left\lfloor \frac{2047k_{PCR} - 512}{k_{PCR} + 512} \right\rfloor + 1$$
$$0 \leq m_{PCR} \leq 31$$
$$0 \leq k_{PCR} \leq 511$$

where $\lfloor x \rfloor$ stands for rounding down x 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 formulas:

$$m_{PCR} = \left\lceil \log_2 \left(\frac{PCR_{sig}}{1023} \right) + 9 \right\rceil$$
$$k_{PCR} = \left\lceil \frac{PCR_{sig}}{2^{m_{PCR}-9}} - 512 \right\rceil$$

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.

Superseded by a more recent version

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:

$$\tau_{PCR} = 2^{e_{PCR}-32} \cdot 2^9 \cdot \left(1 + \frac{w_{PCR} \cdot 2^5}{2^{10}} \right) \text{seconds}$$
$$0 \leq e_{PCR} \leq 31$$
$$0 \leq w_{PCR} \leq 31$$

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.

Superseded by a more recent version

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 $d_{max} - d_{min}$, the difference of the cell transfer delays (between the source and the UNI) of the slowest and fastest cells in the connection (or 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.

5.5 ATM transfer capabilities

5.5.1 General

5.5.1.1 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

Superseded by a more recent version

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 Recommendation I.356). There may be more than one QOS class for a given ATC (see Recommendation 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 service request from a user may implicitly specify the requested ATM transfer capability. Specifying a service type allows for implicit declaration of some or all traffic parameters of a connection. Service type may also be used to implicitly specify the ATM transfer capability that is requested. This may be done in conjunction with explicit declaration of some parameters.

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.

5.5.1.2 Multiplexing and interactions of ATCs

VCCs with the same ATC may be multiplexed into a DBR VPC. Also, a DBR VCC or VPC 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 into a single VPC (e.g. ABR VCCs and ABT VCCs within an SBR VPC);
- the carriage of VCCs within 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).

Superseded by a more recent version

5.5.2 Applicability of ATM transfer capabilities to applications

This subclause 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 subclause 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 ATM transfer capability.

Two examples illustrate the above points:

- The DBR capability 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 which, 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.

5.5.3 Deterministic bit rate transfer capability (DBR)

5.5.3.1 Definition and service model

The deterministic bit rate transfer capability is used by connections that request a static amount of bandwidth that is continuously available during the connection lifetime. This amount of bandwidth is characterized by a peak cell rate value.

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. In the DBR capability, the source may emit cells at, or below the negotiated peak cell rate (and may also even be silent), for periods of time.

Superseded by a more recent version

The DBR capability may be used for both VPCs and VCCs. Neither the tagging option nor selective cell discard of CLP = 1 cells 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.

5.5.3.2 Source traffic descriptor and CDV tolerances

DBR uses peak cell rate and associated CDV tolerance 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 $T_{PCR}(\text{data})$ and PCR for end-to-end user OAM cells with the associated peak emission interval $T_{PCR}(\text{OAM})$.
- ii) One peak cell rate, namely, PCR for user generated cells with the associated peak emission interval $T_{PCR}(\text{agg})$.
- iii) Service type.

If the user intends to transmit user OAM cells and does not allow aggregate shaping of user data and user OAM, the first source traffic descriptor shall apply. The use of the second source descriptor 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 CLP = 0 + 1 descriptor.

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).

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 $\tau_{PCR}(\text{data})$;
- the CDV tolerance corresponding to the PCR of end-to-end user OAM cells $\tau_{PCR}(\text{OAM})$.

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

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, $\tau_{PCR}(\text{agg})$, is required.

In this case, $\tau_{PCR}(\text{agg})$ is computed as $\tau_{PCR}(\text{agg})=T_{PCR}(\text{agg})+\tau_{PCR}(\text{data})$, where $\tau_{PCR}(\text{data})$ is either conveyed via signalling or assigned on a per subscription basis and $T_{PCR}(\text{agg})$ is computed as:

$$T_{PCR}(\text{agg}) = \frac{1}{\frac{1}{T_{PCR}(\text{data})} + \frac{1}{T_{PCR}(\text{OAM})}}$$

When source traffic descriptor ii) is used, only the value of CDV tolerance corresponding to the PCR of user generated CLP = 0 + 1 cells, $\tau_{PCR}(\text{agg})$, is required. In this case, $\tau_{PCR}(\text{agg})$ 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 $T_{PCR}(\text{data})$, $\tau_{PCR}(\text{data})$, and $T_{PCR}(\text{OAM})$, then the user could derive the values for $T_{PCR}(\text{agg})$ and $\tau_{PCR}(\text{agg})$ by the same formulas as in the previous paragraph.

Superseded by a more recent version

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 Recommendation I.610); in this case, the following default rules apply for computing default T_{OAM} , T_{agg} :

- in the case when the operator uses separate conformance definition,
 $T_{PCR}(OAM) = n.T_{PCR}(data)$;
- in the case when the operator uses aggregate conformance definition:

$$T_{PCR}(agg) = \frac{n}{n+1} T_{PCR}(data)$$

These default rules are assuming that OAM user generated cell flows only consist in 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 and associated CDV tolerances from the relevant information conveyed by the present signalling protocol (see Recommendation Q.2931) are given in Appendix II.

The peak cell rate values shall not be re-negotiated at standardized interfaces by means of RM procedures during the lifetime of the connection, but could be re-negotiated via signalling or via network management procedures. The re-negotiation of the PCR may imply also the modification of the associated CDV tolerance.

5.5.3.3 Conformance definition and QOS commitments

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.

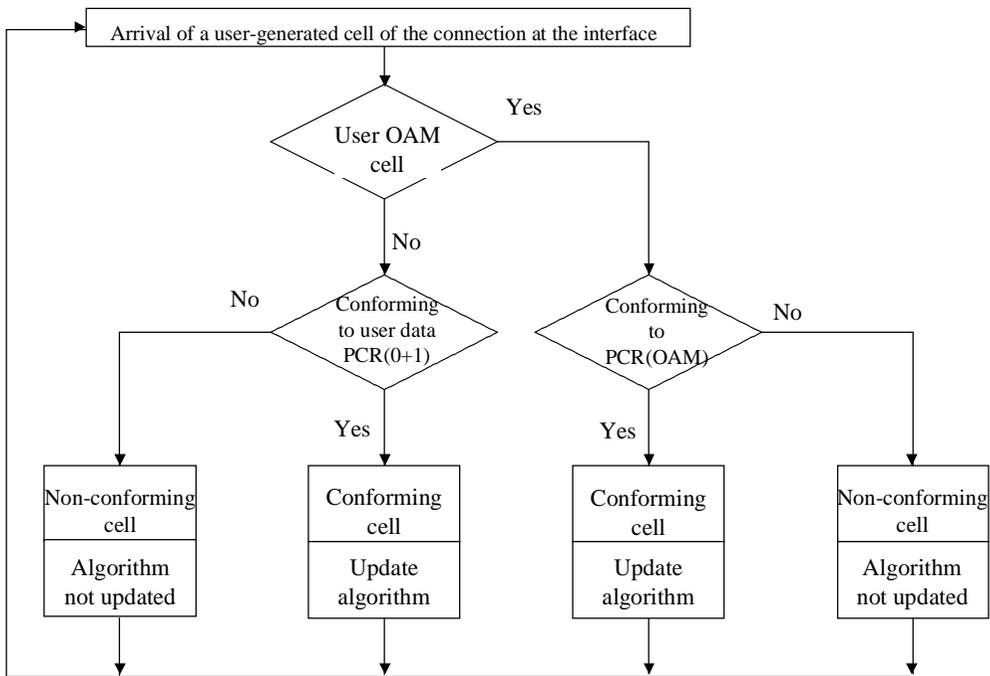
Conformance of a cell flow according to this definition can be estimated by the 1-point measurement process in Recommendation I.356.

Conformance to service type is currently not specified in this Recommendation.

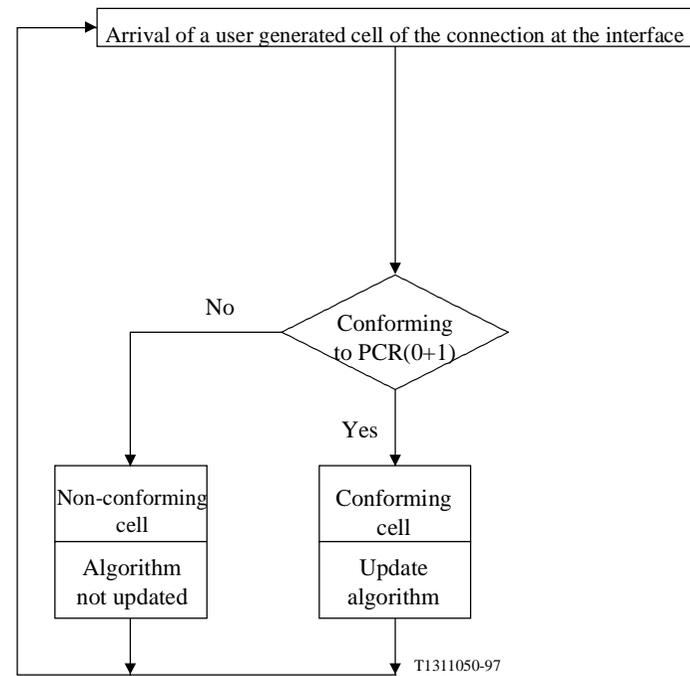
If the user intends to transmit user OAM cells, and if the user does not allow aggregate shaping of user data and user OAM and if the network performs PCR shaping actions (see 6.2.5), then the conformance definition is separate for the user OAM component and the user data component. Otherwise, the conformance definition is aggregate.

- 1) 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.
- 2) If conformance is aggregate, cell conformance of the aggregate user generated CLP = 0 + 1 cell flow is checked with respect to the PCR descriptor negotiated for user generated CLP = 0 + 1 traffic [in case of source traffic descriptor ii)] or 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 [in case of source traffic descriptor i)].

Separate and aggregate conformance definition for the DBR capability are illustrated in Figure 6.



a) Separate conformance definition



b) Aggregate conformance definition

Figure 6/I.371 – Conformance definition for DBR

Superseded by a more recent version

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 6.2.3). Where present, UPC/NPC mechanisms applying to user data cells are performed on the user data aggregate 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 6.2.3.2.

QOS aspects

The QOS class is negotiated at call establishment or on a subscription basis for the connection. The basic QOS class for the DBR capability shall include a specified CLR commitment for the aggregate 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. Conformance definition is specified for the aggregate cell flow, 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 Recommendation I.356):

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.
- When there are non-conforming cells, the method for determining how many cells are assured of the ATM layer QOS is network specific (see Recommendation I.356).
- RM cells, if any, are treated as user data cells.

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 Recommendation I.356).

5.5.4 Statistical bit rate transfer capability (SBR)

5.5.4.1 Definition and service model

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

Superseded by a more recent version

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 is not specified, and this capability may or may not support applications with stringent delay requirements.

The SBR capability may be used for both VPCs and VCCs. The SBR capability may support the tagging option. Selective cell discard may apply to SBR. The use of VC RM cells on a VCC and VP RM cells on a VPC is not currently envisaged to operate SBR; however, such cells that would still be present on the connection are considered as part of the relevant cell flows.

5.5.4.2 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 one of the three subsequent source traffic descriptors:

- i) PCR and SCR/IBT; or
- ii) PCR and service type; or
- iii) service type.

In the SBR capability, the user OAM cell flow is aggregated with the user data cell flow in the traffic parameters (PCR and SCR/IBT, and service type) and in the conformance definition.

In the case where source traffic descriptor i) is used, the following configurations of the PCR and SCR/IBT traffic parameters with the CLP bit and the tagging option 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.
- 2) PCR traffic parameter for the user generated CLP = 0 + 1 cell flow and SCR/IBT traffic parameter for the CLP = 0 cell flow; tagging option 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; tagging option is applied.

Neither the tagging option nor selective cell discard of CLP = 1 cells apply to configuration 1. Selective cell discard may apply to configurations 2 and 3.

When source traffic descriptor i) is used with configuration 1), 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 source traffic descriptor i) is used with configuration 2) and 3), 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.

Superseded by a more recent version

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 re-negotiated at standardized interfaces by means of RM procedures during the lifetime of the connection, but could be re-negotiated via signalling or via network management procedures. The re-negotiation of the source traffic descriptor may imply also the modification of the associated CDV tolerances.

5.5.4.3 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 tagging option).
- 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.

Conformance to service type is currently not specified in this Recommendation.

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

Superseded by a more recent version

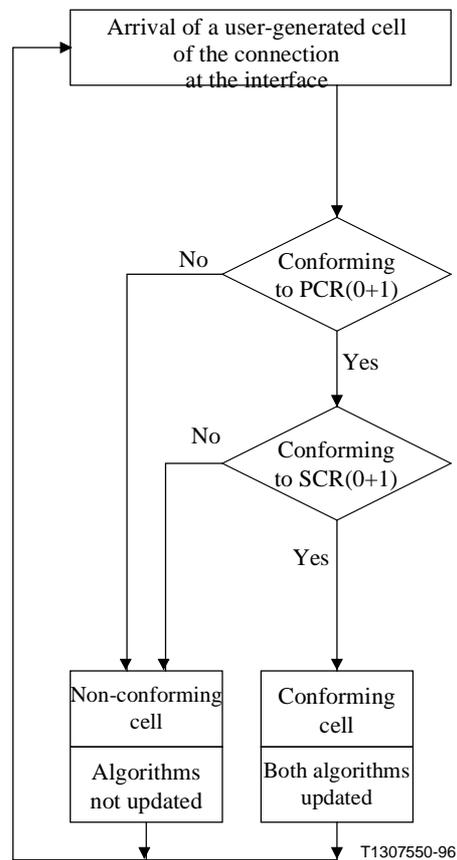
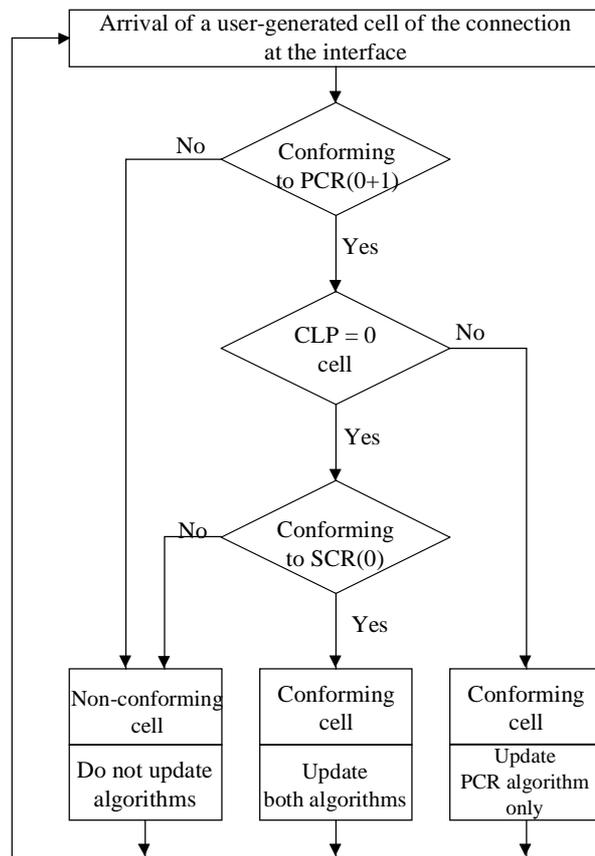


Figure 7/I.371 – Conformance definition for SBR configuration 1

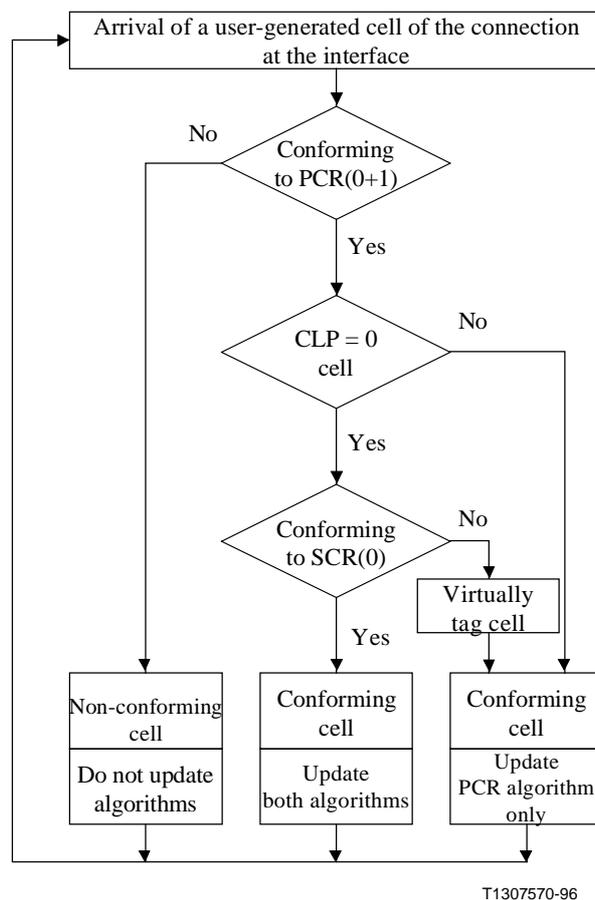
Superseded by a more recent version



T1307560-96

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

Superseded by a more recent version



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, the quality of service (QOS) commitments is for the aggregate $CLP = 0 + 1$ cell flow. Thus, the QOS commitments are independent of the CLP bit. There is a specified QOS commitment on the cell loss ratio of the aggregate cell flow, and the value will be specified (see Recommendation I.356). There may be a QOS commitment on delay.

For traffic-parameter configurations 2 and 3 above there is a specified QOS commitment on the cell loss ratio of the $CLP = 0$ cell flow, and the value will be specified (see Recommendation I.356). The QOS commitment on the cell loss ratio of the aggregate $CLP = 0 + 1$ cell flow is unspecified. Hence, the QOS commitment on the cell loss ratio of the $CLP = 1$ cell flow is also unspecified. There may be a QOS commitment concerning delay, and if there is, it applies to the $CLP = 0 + 1$ cell flow.

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.

Superseded by a more recent version

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.

5.5.5 ATM block transfer capabilities (ABT)

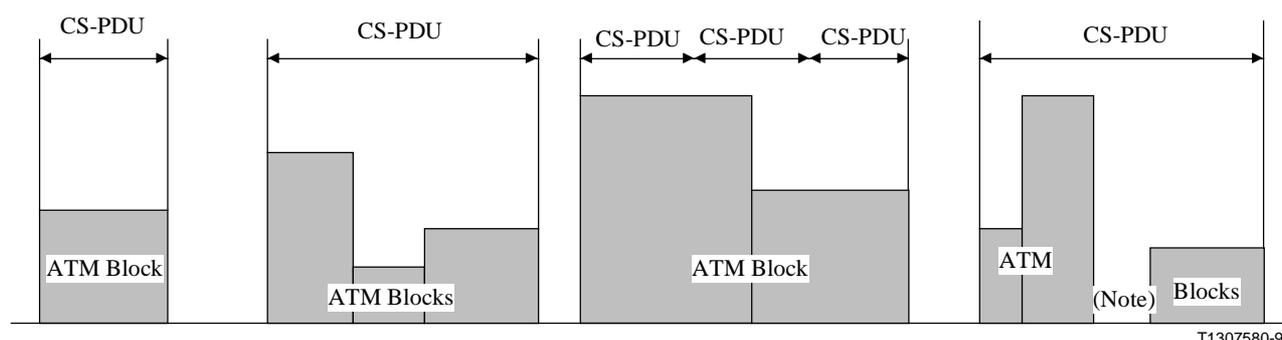
An ATM Block Transfer (ABT) capability is an ATM layer mechanism for providing a service where the ATM layer transfer characteristics are negotiated on an ATM block basis. Within an ATM block accepted by the network, the network allocates sufficient resources such that the QOS received by the ATM block 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).

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.

Superseded by a more recent version

ABT does not support the tagging option. Selective cell discard functions based on the CLP 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 re-negotiation 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 re-negotiated 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).

5.5.5.1 ABT with delayed transmission (ABT/DT)

5.5.5.1.1 Definition and service model

In ABT/DT, during the connection lifetime, the BCR of successive ATM blocks is dynamically negotiated with the network. BCR re-negotiations may be initiated by either end user on both the forward and the backward direction. For a given direction, only one re-negotiation 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 re-negotiations in either direction, two BCR re-negotiations 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 re-negotiations, the BCR re-negotiation initiated by the destination in the backward direction has priority over the BCR re-negotiation initiated by the source in the forward direction. Furthermore, BCR re-negotiations initiated by the network have priority over BCR re-negotiations initiated by end users (see 5.5.5.1.4).

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. 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 5.5.5.1.2), then a

Superseded by a more recent version

new BCR reservation should be accepted by the network within a finite time interval, consistent with QoS commitments (see 5.5.5.1.4).

In the case that a sustainable cell rate equal to 0 has been specified, the network may accept or may deny BCR re-negotiations 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 re-negotiated the allocated BCR.

5.5.5.1.2 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 re-negotiation rate, namely the peak cell rate PCR(RM) of the ABT/DT RM cell flows.

The above parameters may be declared by means of a service type (see 5.1.1). They are static and it is not specified in this Recommendation that they will be re-negotiated 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 \neq 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.

5.5.5.1.3 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 re-negotiated 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 1 contains the RM cell format for ABT transfer capabilities.

Superseded by a more recent version

Table 1/I.371 – RM cell format for ABT

Field	Octet(s)	Bit(s)	Coding
ATM Header (Note 1)	1-5	All	As in Rec. I.361
Protocol ID	6	All	2 (ABT/DT) 3 (ABT/IT)
Message type: Direction	7	8	(Note 2)
Message type: Traffic management cell (Note 3)	7	7	
Message type: Congestion indication	7	6	
Message type: Maintenance (Note 4)	7	5	
Message type: Req/Ack (Note 5)	7	4	
Message type: Elastic/rigid (Note 5)	7	3	
Message type: Reserved	7	1-2	7.1
CLP = 0 + 1 BCR (user data + user OAM cells)	8-9	All	
User OAM BCR	10-11	All	
Reserved (Note 6)	12-13	All	7.1
Block size	14-17	All	(Note 7)
Sequence number	18-21	All	(Note 8)
Reserved	22-51	All	7.1
Reserved	52	3-8	7.1
CRC-10	52	1-2	
	53	All	
<p>NOTE 1 – Only ABT RM cells with CLP bit set to 0 are currently specified in this Recommendation.</p> <p>NOTE 2 – The direction bit equals 0 for forward RM cells and equals 1 for backward RM cells.</p> <p>NOTE 3 – This bit is used as BECN in ABR.</p> <p>NOTE 4 – This bit is used as no-increase in ABR.</p> <p>NOTE 5 – These bits are currently not used in ABR.</p> <p>NOTE 6 – This field is used to convey the MCR in the ABR capability.</p> <p>NOTE 7 – Integer valued. Least significant bit is bit 1 of octet 17.</p> <p>NOTE 8 – Integer valued. Least significant bit is bit 1 of octet 21.</p>			

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 re-negotiation 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.

Superseded by a more recent version

Congestion Indication (CI): when the ABT/DT RM cell is involved in a BCR modification, this bit indicates whether the BCR re-negotiation 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 $\text{maintenance} = 0$, and ABT/DT RM cells used for maintenance of ABT/DT procedures are identified by $\text{maintenance} = 1$. 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 $\text{Req/Ack} = 0$, the RM cell is a BCR modification request;
- 2) when sent by the user and $\text{Req/Ack} = 1$, 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 $\text{Req/Ack} = 1$, the ABT/DT RM cell is an acknowledgement RM cell of a BCR modification;
- 4) when sent by the network and $\text{Req/Ack} = 0$, 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 $\text{CLP} = 0 + 1$ 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:

$$\Lambda = \left[2^m \cdot \left(1 + \frac{k}{512} \right) \right] \cdot nz$$
$$0 \leq m \leq 31 \text{ and } 0 \leq k \leq 511$$
$$nz \in \{0,1\}$$

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 $\text{CLP} = 0 + 1$ cell flow but applies to the user OAM cell flow.

Block size and sequence number

These fields are used in ABT/IT (see 5.5.5.2.3).

Superseded by a more recent version

5.5.5.1.4 Conformance definition and QOS commitments

Conformance definition

Conformance definitions at the cell level and at the block level are required. This is under study.

QOS aspects

At call establishment or at connection subscription time, the user may negotiate a QOS class for the connection. QOS commitments 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. 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 needs 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, 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 re-negotiation. The cell level QOS is in any case assured as long as the network has not re-negotiated 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.

5.5.5.2 ABT with immediate transmission (ABT/IT)

5.5.5.2.1 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.

Superseded by a more recent version

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 given direction. 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 5.5.5.2.4), 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. In that case, 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.

5.5.5.2.2 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 5.5.5.1.2).

5.5.5.2.3 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 is given in Table 1 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 5.5.5.1.3).

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.

Superseded by a more recent version

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.

5.5.5.2.4 Conformance definition and QOS commitments

Conformance definition

Conformance definitions at the cell level and at the block level are required. This is under study.

QOS aspects

At call establishment or at connection subscription time, the user may negotiate a QOS class for the connection. QOS commitments 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. 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 DBR connection 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 needs 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, 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

Superseded by a more recent version

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 re-negotiation. The cell level QOS is in any case assured as long as the network has not re-negotiated 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).

5.5.6 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).

5.5.6.1 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 this case, the VPC is allocated a static bandwidth. Dynamically changing the bandwidth of the VPC via ABR is currently not specified in this Recommendation.

Superseded by a more recent version

In ABR, user data cells have the CLP bit set to zero. Tagging is currently not supported by ABR. It is for further study whether ABR may also use CLP = 1 user data cells and whether 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.

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. The procedures and restrictions regarding the generation of out-of-rate RM cells is currently not specified in this Recommendation.

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 5.5.6.4).

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 bi-directional 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 11).

Superseded by a more recent version

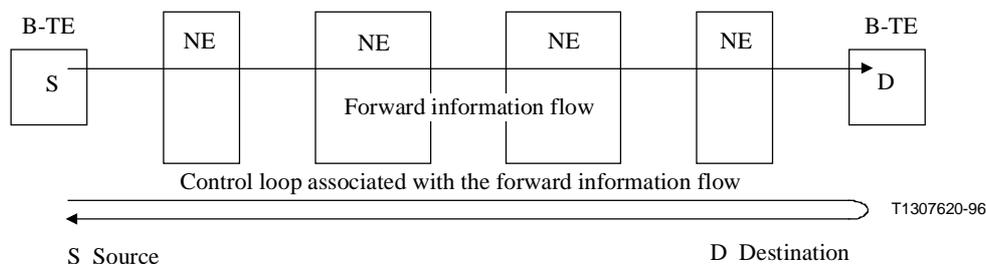


Figure 11/I.371 – Example of a source to destination ABR control loop

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 12 illustrates an ABR virtual connection which incorporates segmentation.

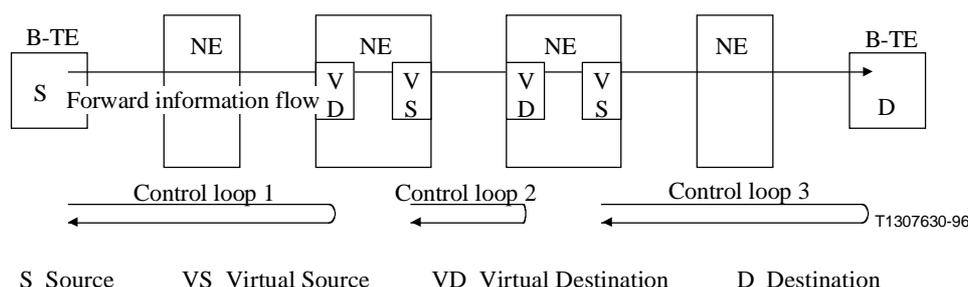


Figure 12/I.371 – Example of a segmented virtual connection

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.

5.5.6.2 Source traffic descriptor and CDV tolerances

There are two types of source traffic parameters associated with ABR: traffic parameters which are negotiated at call establishment and cannot be changed by RM procedures and dynamically changing traffic parameters which may be changed by RM procedures (see 5.5.6.3). With respect to conformance at interfaces, the cell rates in ABR are defined with respect to tolerance parameters, see 5.5.6.4.

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

- A PCR which is the maximum cell rate requested by the application (see 5.4.1).
- An MCR 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).

Superseded by a more recent version

Tolerances associated to the source traffic descriptor are under study.

The above parameters may be declared by means of a service type (see 5.1.1).

5.5.6.3 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 the MCR and the PCR.

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

Table 2 contains the RM cell format for the ABR transfer capability.

Table 2/I.371 – Format for the RM cell that supports ABR

Field	Octet(s)	Bit(s)	Coding
Header	1-5	All	As in Rec. I.361
Protocol ID	6	All	1
Message type: Direction	7	8	(Note 1)
Message type: BECN cell	7	7	(Note 2)
Message type: Congestion indication	7	6	(Note 3)
Message type: No-increase	7	5	(Note 4)
Message type: (Note 5)	7	4	
Message type: (Note 6)	7	3	
Message type: Reserved	7	1-2	7.1
Explicit Cell Rate (ECR)	8-9	All	(Note 7)
Current Cell Rate (CCR)	10-11	All	(Note 7)
Minimum Cell Rate (MCR)	12-13	All	(Note 7)
Queue length	14-17	All	(Note 8)
Sequence number	18-21	All	(Note 9)
Reserved	22-51	All	7.1
Reserved	52	3-8	7.1
CRC-10	52	1-2	
	53	All	
NOTE 1 – The direction bit equals 0 for forward RM cells and equals 1 for backward RM cells.			
NOTE 2 – The BECN-cell bit equals 1 for BECN cells and equals 0 otherwise.			
NOTE 3 – The congestion indication bit equals 1 to indicate congestion and equals 0 otherwise.			
NOTE 4 – The no-increase bit equals 1 to indicate that no rate increase by the source is allowed and equals 0 otherwise.			
NOTE 5 – This bit is req/ack in the ABT capability.			
NOTE 6 – This bit is elastic/rigid in the ABT capability.			
NOTE 7 – 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 5.5.6.3.1 below.			
NOTE 8 – Integer valued. Least significant bit is bit 1 of octet 17.			
NOTE 9 – Integer valued. Least significant bit is bit 1 of octet 21.			

Superseded by a more recent version

5.5.6.3.1 Details on the fields

Protocol ID

ABR RM cells are identified by protocol ID 1.

Message type

Message type is an octet containing five single bit fields plus three 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: 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. (BECN is an acronym for Backward Explicit Congestion Notification.)

Congestion Indication (CI): this bit indicates imminent 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:

$$ECR = \left[2^m \cdot \left(1 + \frac{k}{512} \right) \right] \cdot nz \text{ cells per second}$$
$$0 \leq m \leq 31 \text{ and } 0 \leq k \leq 511$$
$$nz = \{0,1\}$$

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.

Superseded by a more recent version

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.

5.5.6.4 Conformance definition and QOS commitments

Conformance definition

A conformance definition is required. This is under study.

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.

ABR capability without using a conformance definition

The ABR capability may also be used in a mode where the conformance definition in this subclause 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 subclause 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.

Superseded by a more recent version

6 Functions for traffic control and congestion control

6.1 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 malfunctioning 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.

6.1.1 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 (6.2.1).
- ii) Connection admission control (6.2.2).
- iii) Usage/network parameter control (6.2.3).
- iv) Priority control (6.2.4).
- v) Traffic shaping (6.2.5).
- vi) Fast resource management (6.2.6).

Congestion control functions

- vii) Selective cell discard (6.3.1).
- viii) Explicit forward congestion indication (6.3.2).

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

- ix) Connection admission control that reacts to and takes account of the measured load on the network.
- x) Variation of usage monitored parameters by the network. For example, reduction of the peak rate available to the user.
- xi) Other traffic control 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 network performance may be provided on ATM connections by proper routing, traffic shaping, priority control and resource allocation to meet the required ATM layer QOS for these connections.

Superseded by a more recent version

6.2 Traffic control functions

6.2.1 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.

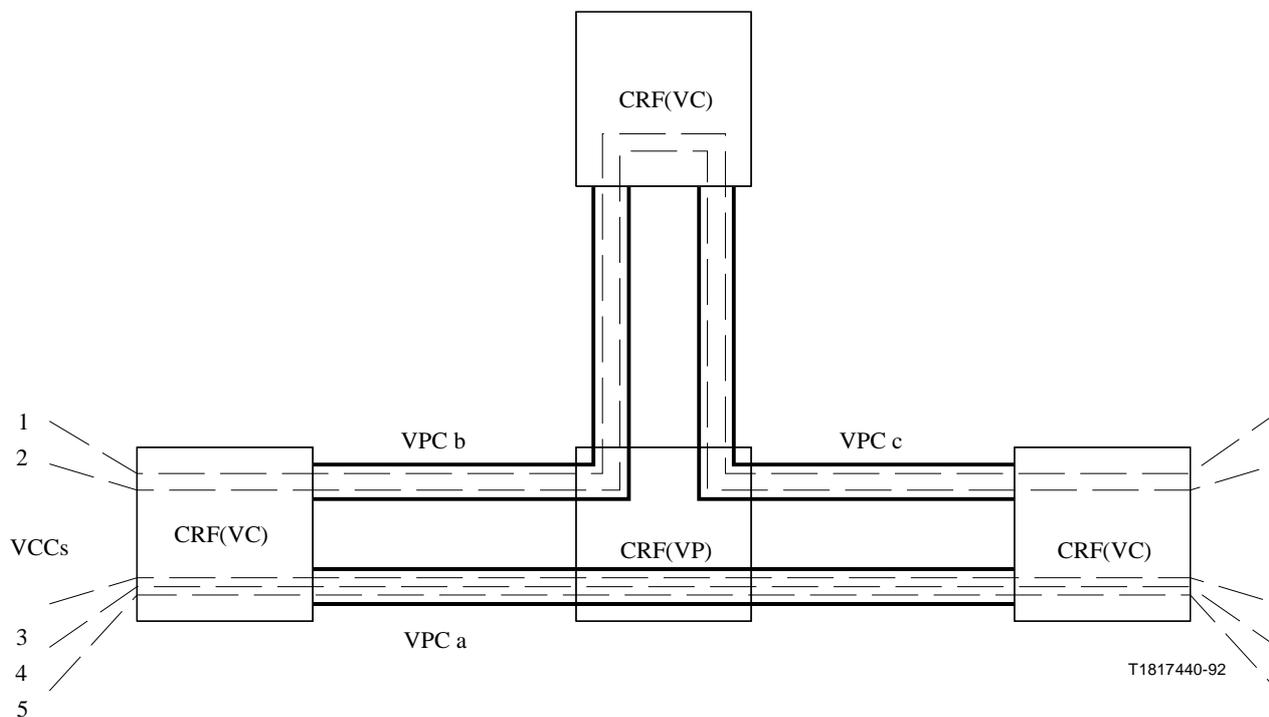
The peer-to-peer network performance on a given VCC depends on the performances of the consecutive VPCs used by this VCC and on how it is handled in CRF(VC)s (connection-related functions at the VC level, e.g. a VC multiplexer, a VC switch, see Figure 13).

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).

Superseded by a more recent version



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.

NTOE 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 13/I.371 – Mapping cell loss ratios for virtual channel connections and virtual path connections

On the basis of the applications of VPCs contained in Recommendation 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.

Superseded by a more recent version

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.

6.2.2 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 re-negotiation 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 re-negotiated 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 re-negotiations. 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 re-negotiated with an appropriate procedure, either off-line (e.g. at subscription) or on-line via management.

Depending on the selected ATC (see 5.5), dynamic modifications of traffic characteristics may also use ATM layer resource management procedures (see 6.2.6).

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 re-negotiation 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);
- routing and allocation of network resources.

Superseded by a more recent version

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 6.2.4). 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 Recommendation I.356.

6.2.3 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).

6.2.3.1 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.
- 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.

Control tasks of the UPC/NPC are cell discarding and optionally cell tagging.

Superseded by a more recent version

6.2.3.2 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 standardise 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 Recommendation 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 tagging action must operate as a single server using First-In-First-Out (FIFO) service discipline for each ATM connection.

6.2.3.2.1 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 Recommendation 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 Recommendation I.356).

Superseded by a more recent version

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.

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 14). 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 .

Superseded by a more recent version

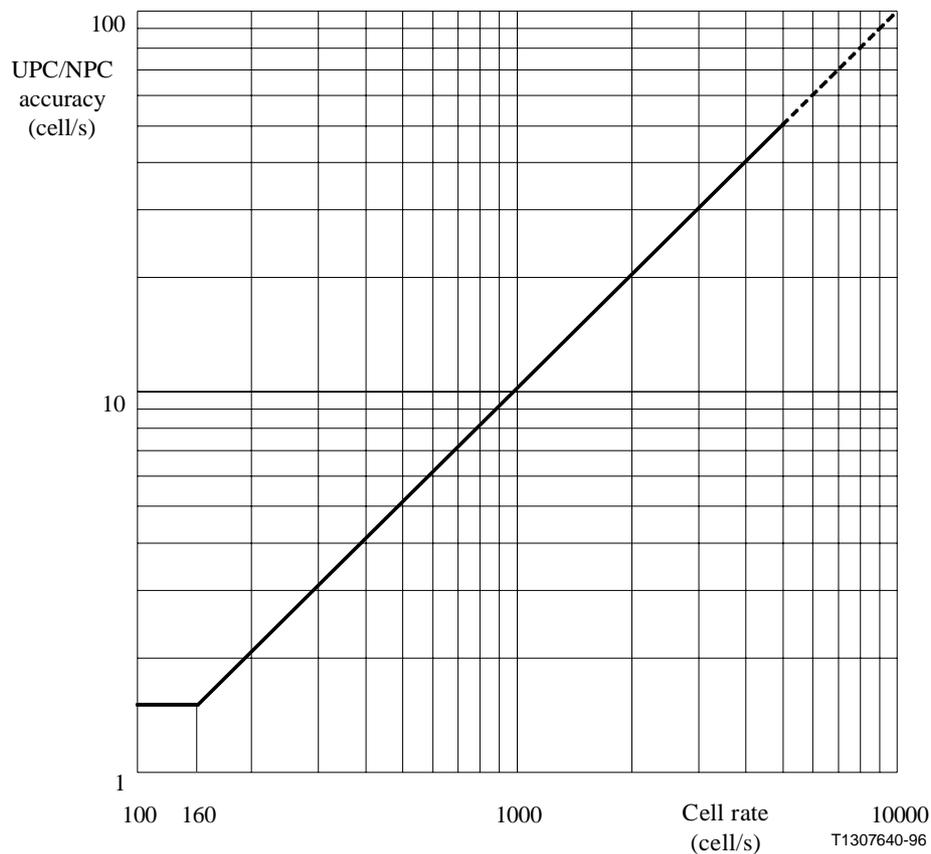


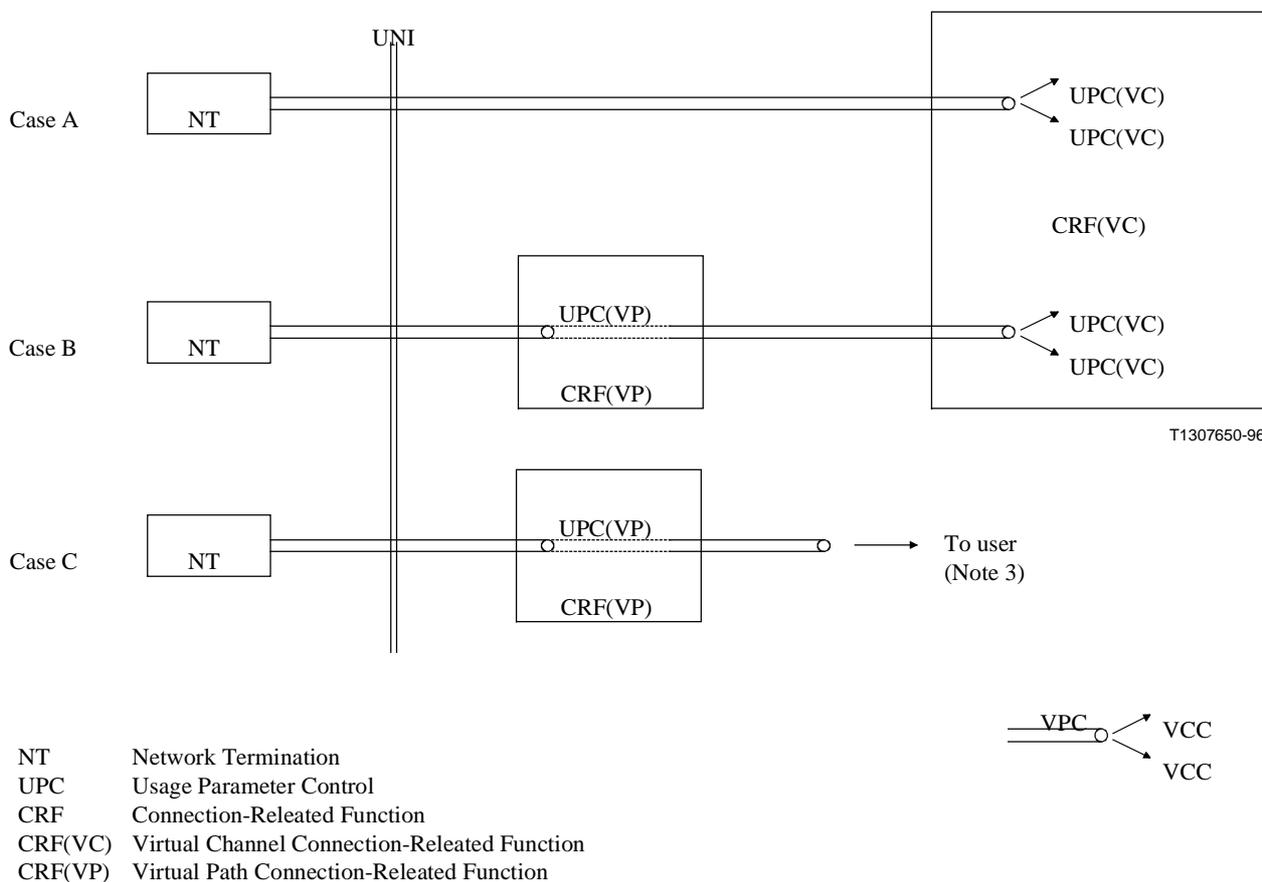
Figure 14/I.371 – UPC/NPC accuracy requirement

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

6.2.3.3 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 15:

Superseded by a more recent version



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 15/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 15): User connected directly to CRF(VC)

Usage parameter control is performed within the CRF(VC) on VCCs (action 1, 6.2.3.1).

Case B (Figure 15): User connected to CRF(VC) via CRF(VP)

Usage parameter control is performed within the CRF(VP) on VPCs only (action 2, 6.2.3.1) and within the CRF(VC) on VCCs only (action 1, 6.2.3.1).

Case C (Figure 15): 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, 6.2.3.1).

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

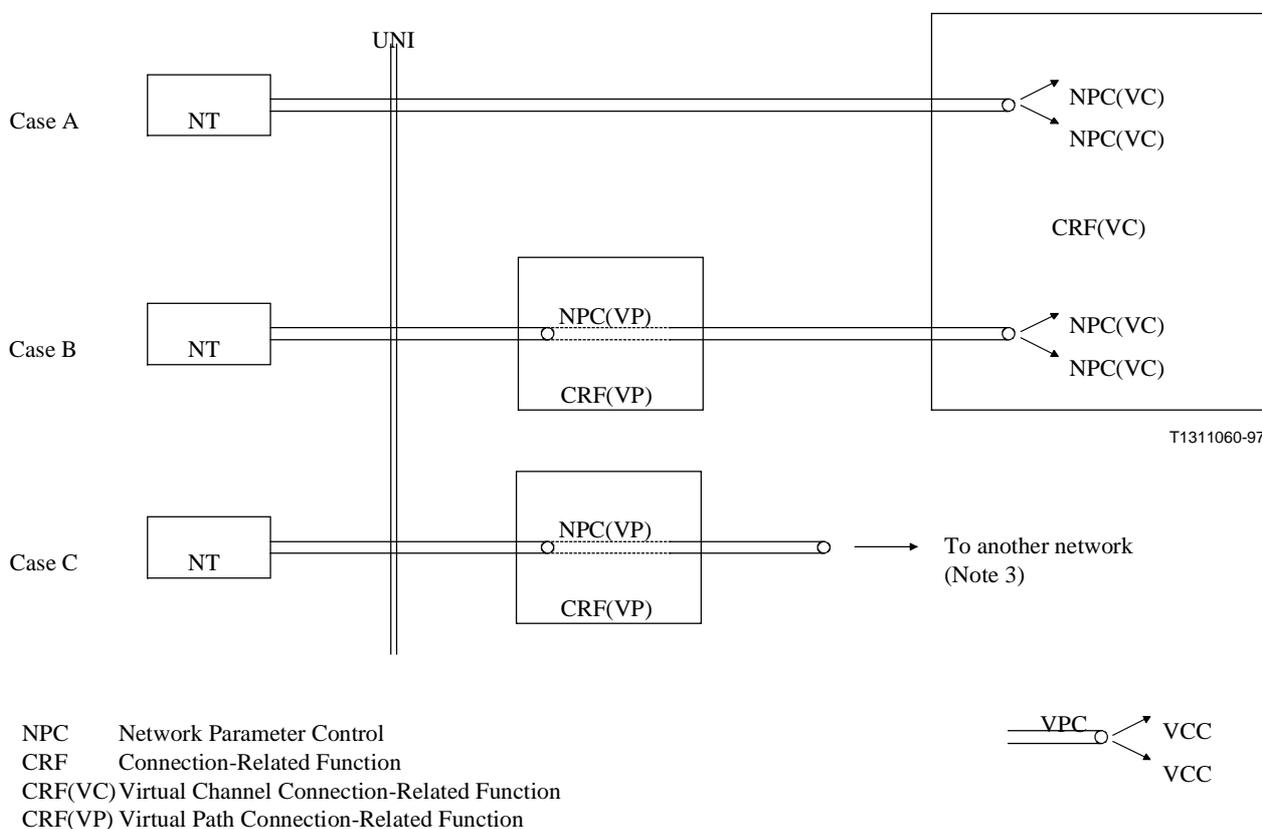
- at the VP level with the CRF(VP) (when the user requests the operator of the CRF(VP) for a VP connection);
- at the VC level with the CRF(VC) (when the user requests the operator of the CRF(VC) for a VC connection).

Superseded by a more recent version

Implication of this is under study.

6.2.3.4 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 16.



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 identify a negotiated VPC.

NOTE 3 – Another network may be a user.

Figure 16/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 16): Originating network connected directly to CRF(VC)

NPC is performed within the CRF(VC) (action 1, 6.2.3.1).

Case B (Figure 16): Originating network connected to the CRF(VC) via the CRF(VP)

NPC is performed within the CRF(VP) on VPCs only (action 2, 6.2.3.1) and within the CRF(VC) on VCCs only (action 1, 6.2.3.1).

Case C (Figure 16): Originating network connected to another network via CRF(VP)

NPC is performed within the CRF(VP) on VPCs only (action 2, 6.2.3.1).

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

Superseded by a more recent version

- 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.

6.2.3.5 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 clause 5). 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.

Enforcement of connections defined by a service type is currently not specified in this Recommendation.

6.2.3.6 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.

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

- a) cell passing;
- b) cell tagging (is user and network operator option – see 5.3.4); cell tagging operates on CLP = 0 cells only, by overwriting the CLP bit to 1;
- 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.

UPC/NPC and traffic shaping functions may be combined in some implementations, in which case cell re-scheduling actions will result from this association.

At the ATM block level (see 5.5.5), an action of the UPC/NPC may be:

- Initiating a modification of the resources allocated to the connection.
- Discarding all the remaining cells in the ATM block.

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 re-negotiation 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.

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.

6.2.3.7 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 an ATM connection utilizes the CLP capability on user request and if the tagging option is allowed, CLP = 0 cells identified by the UPC/NPC function performed on CLP = 0 flow as

Superseded by a more recent version

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 the tagging option 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, 5.5.4.3) 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.

6.2.3.8 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 re-negotiation of the amount of resources allocated to the connection (see 6.2.3.6);
- 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 re-negotiation attempt, the functions addressed above are currently not specified in this Recommendation.

6.2.4 Priority control

Selective cell discard is the only priority control function that is currently specified. It refers to the action of selectively discarding CLP = 1 cells (which are either submitted as lower priority cells by the user or tagged by the network UPC/NPC) when buffer overflow occurs within the network.

In the context of traffic control functions, selective cell discard may be performed by network elements while still meeting network performance objectives.

The applicability of selective cell discard depends on the ATM transfer capability (see 5.5).

Network elements shall not change the value of the CLP bit, except possibly in the case when the tagging option is used (see 6.2.3.6).

6.2.5 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.

Superseded by a more recent version

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 5.5).

6.2.6 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 5.5.5) and the ABR (see 5.5.6) ATM transfer capabilities make use of fast resource management functions to dynamically allocate resources to connections using these capabilities. Potential other fast resource management functions are for further study.

Fast resource management functions use resource management cells which are described in 7.1.

6.3 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.

6.3.1 Selective cell discard

In the context of congestion control, selective cell discard of $CLP = 1$ cells as defined in 6.2.4 may be applied within a congested network element, primarily in order to protect, as long as possible, high priority $CLP = 0$ flows.

6.3.2 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 in a congested state may set an explicit forward congestion indication in the cell header 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 is not in a congested state will not modify the value of this indication.

Superseded by a more recent version

The mechanism by which a network element determines whether it is congested is an implementation issue and is not subject to standardization. The mechanism by which the congestion indication is used by the higher layer protocols in the 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.

6.3.3 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).

6.4 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.

6.4.1 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 17) 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 17, there may be multiple network operators involved. The inter-operation between network operators within a B-ISDN is not addressed in this subclause.

Superseded by a more recent version

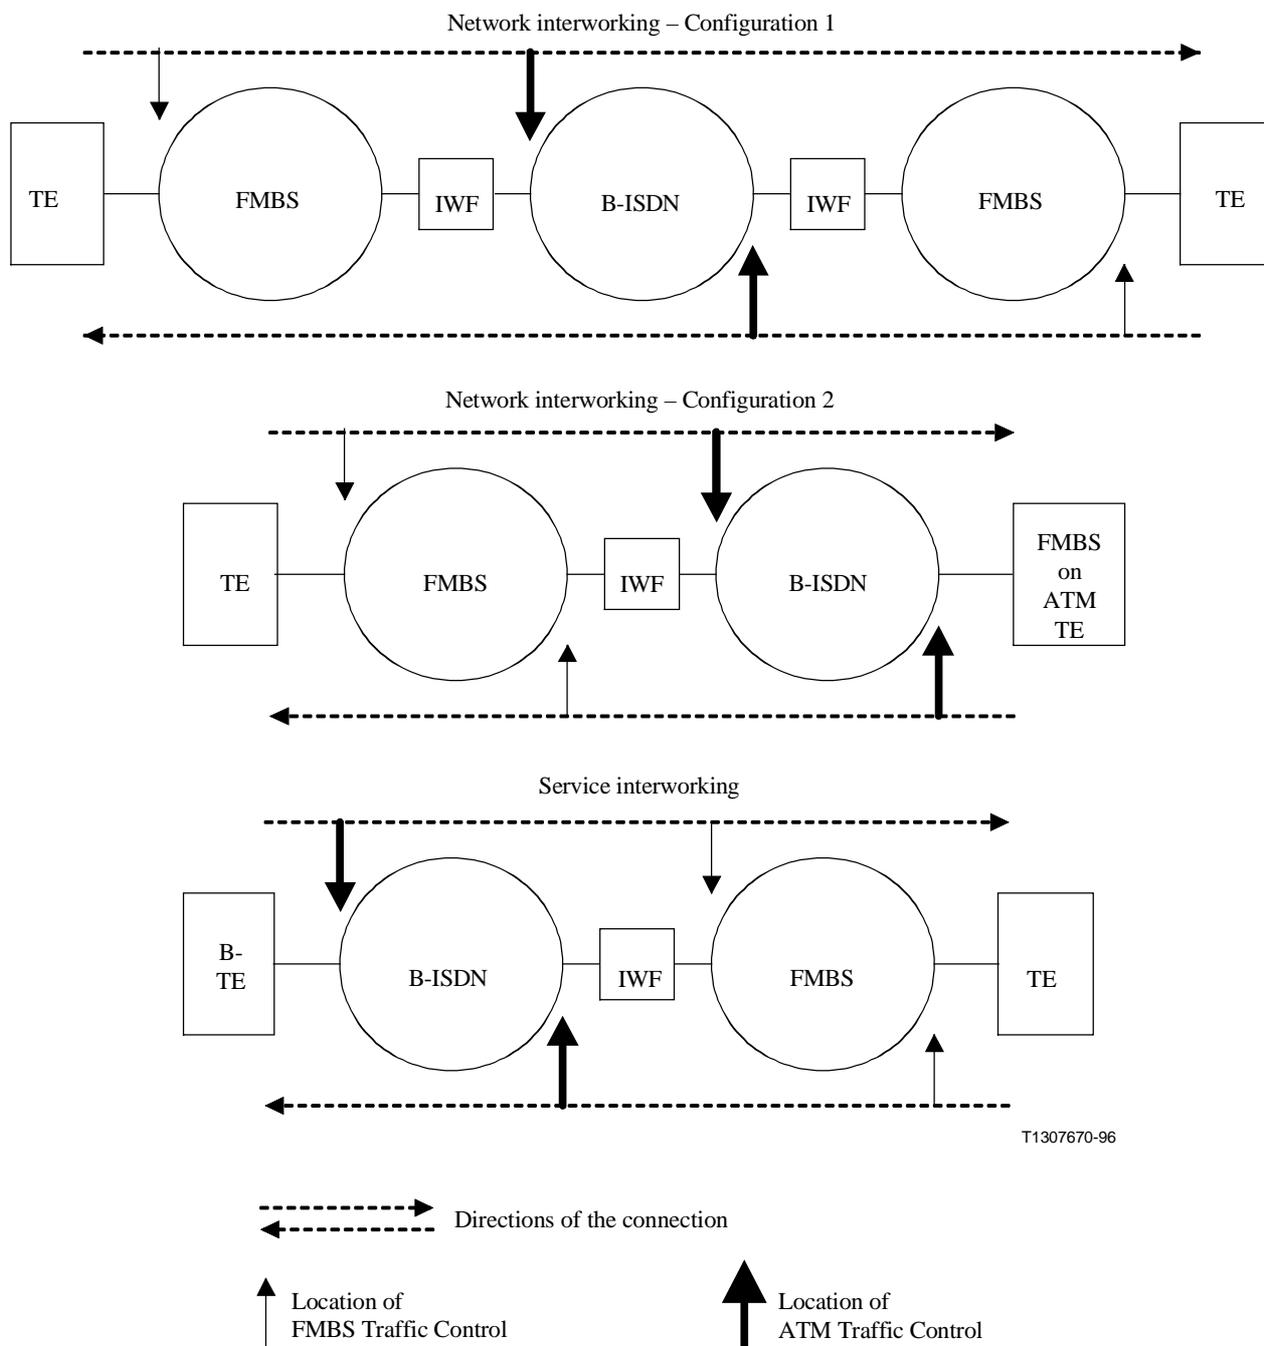


Figure 17/I.371 – Reference configurations for traffic control interworking between FMBS and B-ISDN

For interworking case 2 on Figure 17, there are two traffic contracts applying to the FMBS on ATM terminal. The implications on conformance definition are under study.

7 Procedures for traffic control and congestion control

7.1 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.

Superseded by a more recent version

The ATM layer RM cells contain fields common to all types of RM cells (see Figure 18) as well as specific fields for each type of RM cell.

ATM Header	RM Protocol Identifier	Function specific fields	Reserved	EDC (CRC-10)
5 octets	8 bits	45 octets	6 bits	10 bits

EDC Error Detection Code T1307680-96

Figure 18/I.371 – Resource 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 Recommendation I.361 for the coding of the ATM header for VC and VP resource management cells

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.

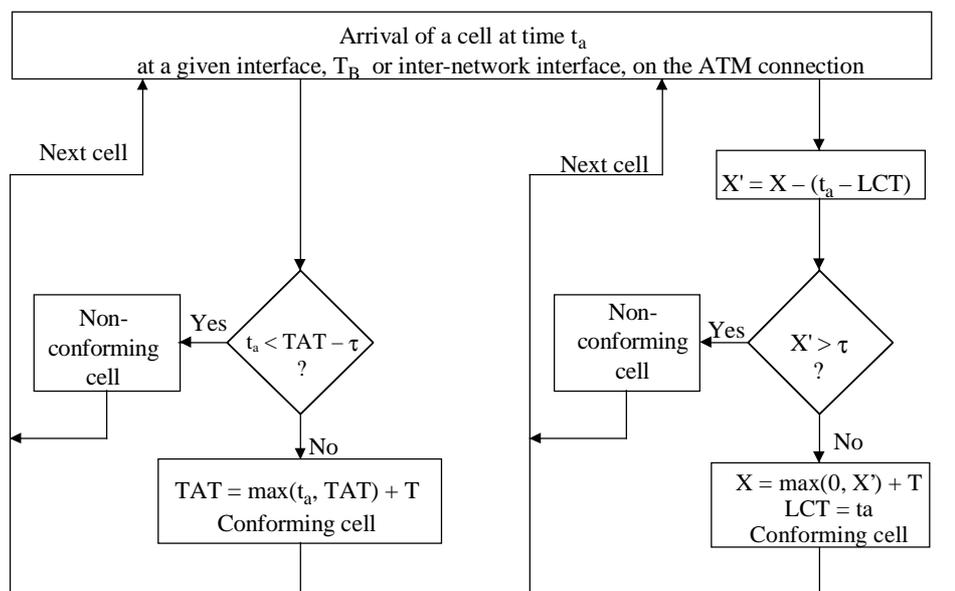
ANNEX A

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.

Superseded by a more recent version



T1307690-96

Virtual scheduling algorithm

TAT Theoretical Arrival Time
 t_a Time of arrival of a cell to the given interface

At the time of arrival t_a of the first cell of the connection to cross the given interface, $TAT = t_a$

Continuous-state leaky bucket algorithm

X Value of the Leaky Bucket counter
 X' Auxiliary variable
LCT Last Conformance Time

At the time of arrival t_a of the first cell of the connection to cross the given interface, $X = 0$ and $LCT = t_a$

Figure A.1/I.371 – Equivalent versions of the generic cell rate algorithm

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 k^{th} 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 $[t_a(k) + T]$.

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

Superseded by a more recent version

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 k^{th} cell $t_a(k)$, first the content of the bucket is provisionally updated to the value X' , which equals the content of the bucket 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.

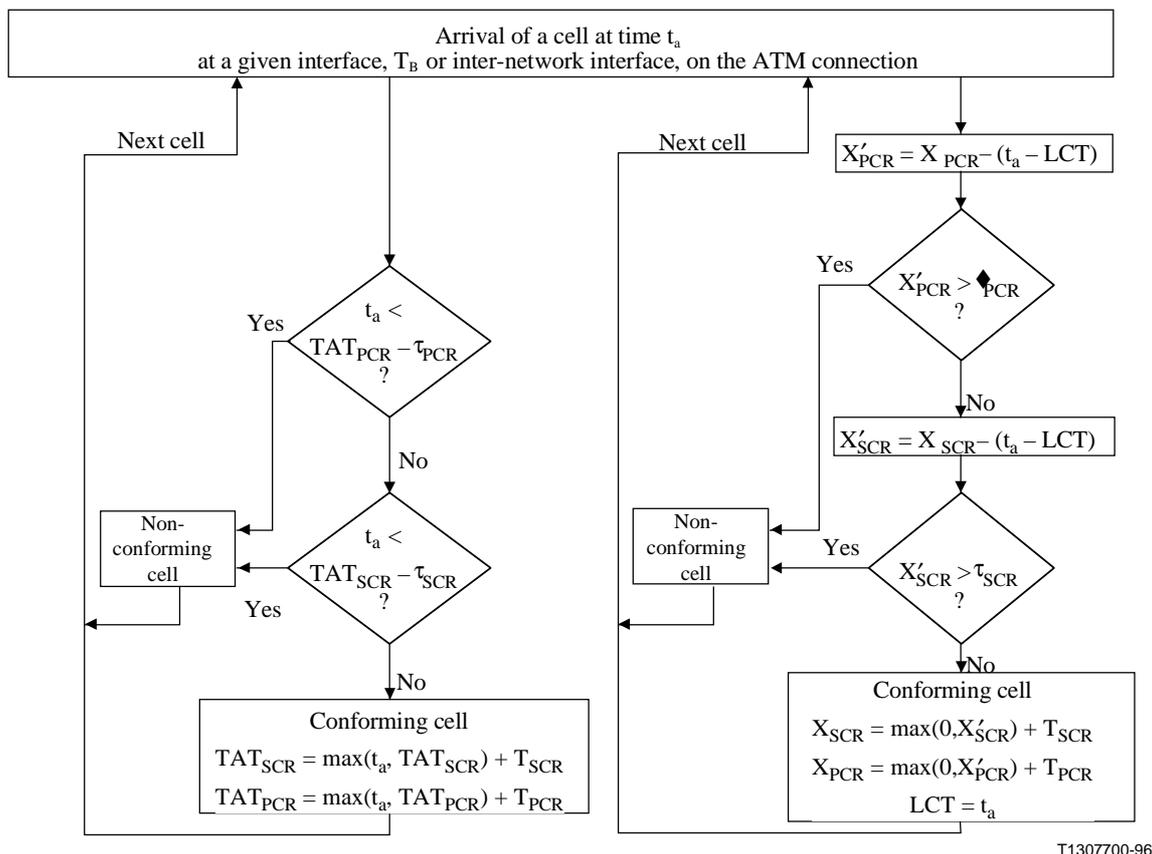
NOTE – The two algorithms presented Figure A.1 are equivalent in that sense that, for any sequence of cell arrival times $\{t_a(k), k \geq 1\}$, the two algorithms determine the same cells to be conforming and thus the same cells to be non-conforming. The two algorithms are easily compared if one notices 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.

Superseded by a more recent version



T1307700-96

Virtual scheduling algorithm

TAT_{SCR}, TAT_{PCR} Theoretical Arrival Times

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,

$$TAT_{SCR} = TAT_{PCR} = t_a$$

Continuous-state leaky bucket algorithm

X_{SCR}, X_{PCR} Values of the Leaky Bucket counters

X'_{SCR}, X'_{PCR} Auxiliary variables

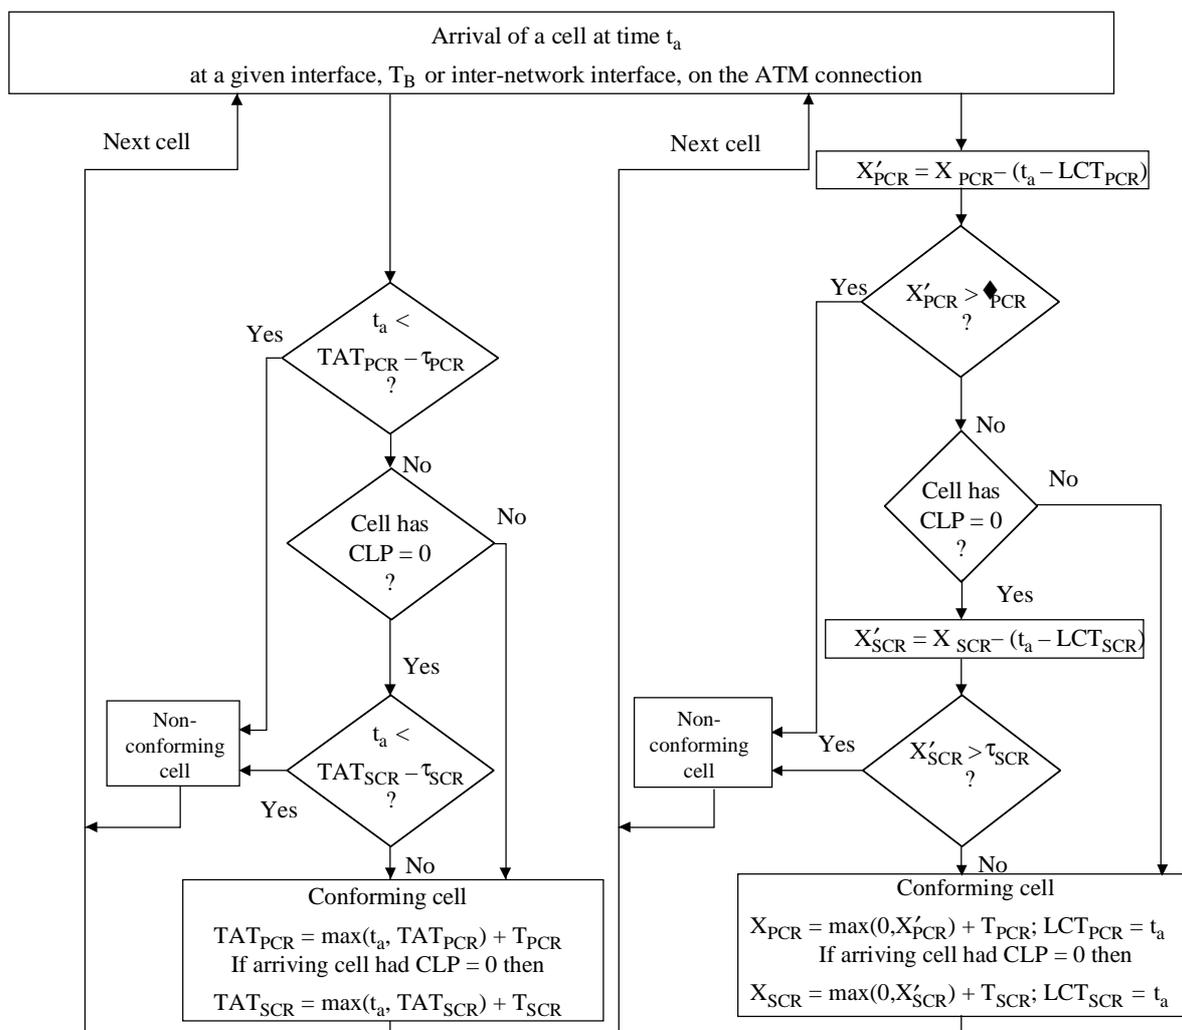
LCT Last Conformance Time

At the time of arrival t_a of the first cell of the connection to cross the given interface,

$$X_{SCR} = X_{PCR} = 0 \text{ and } LCT = t_a$$

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

Superseded by a more recent version



Virtual scheduling algorithm

TAT_{SCR}, TAT_{PCR} Theoretical Arrival 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,
 $TAT_{SCR} = TAT_{PCR} = t_a$

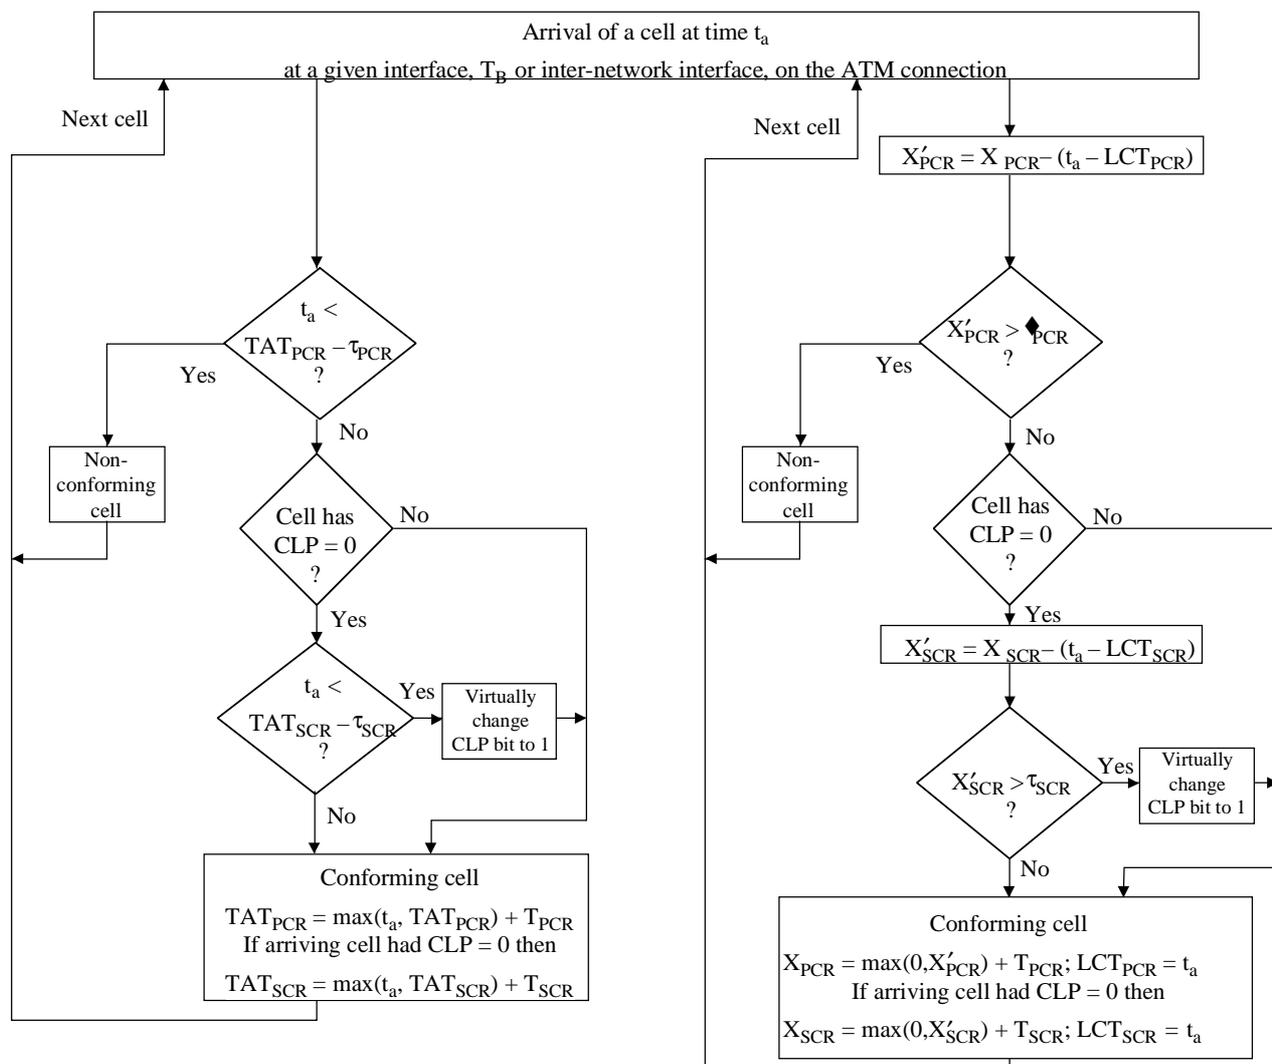
Continuous-state leaky bucket algorithm

X_{SCR}, X_{PCR} Values of the Leaky Bucket counters
 X'_{SCR}, X'_{PCR} Auxiliary variables
 LCT_{SCR}, LCT_{PCR} Last Conformance Times

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 option is not applied)

Superseded by a more recent version



T1307720-96

Virtual scheduling algorithm

TAT_{SCR}, TAT_{PCR} Theoretical Arrival 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,

$$TAT_{SCR} = TAT_{PCR} = t_a$$

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

At the time of arrival t_a of the first cell of the connection to cross the given interface,

$$X_{SCR} = X_{PCR} = 0 \text{ and } LCT_{SCR} = LCT_{PCR} = t_a$$

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 option is applied)

Superseded by a more recent version

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).

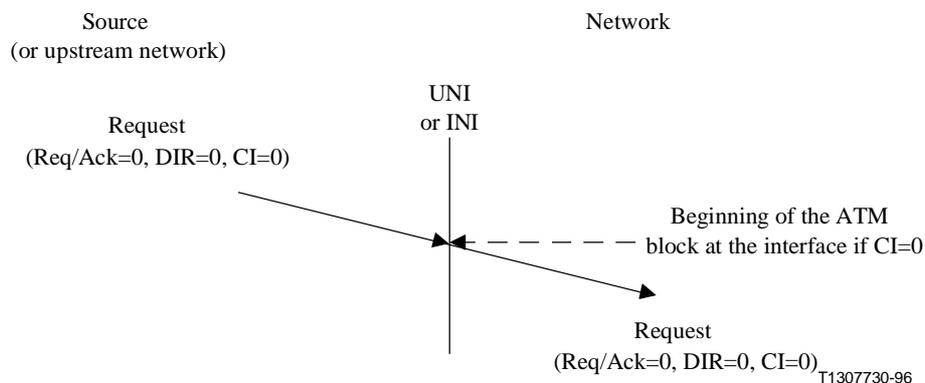


Figure C.1/I.371 – BCR decrease initiated by the source
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

Superseded by a more recent version

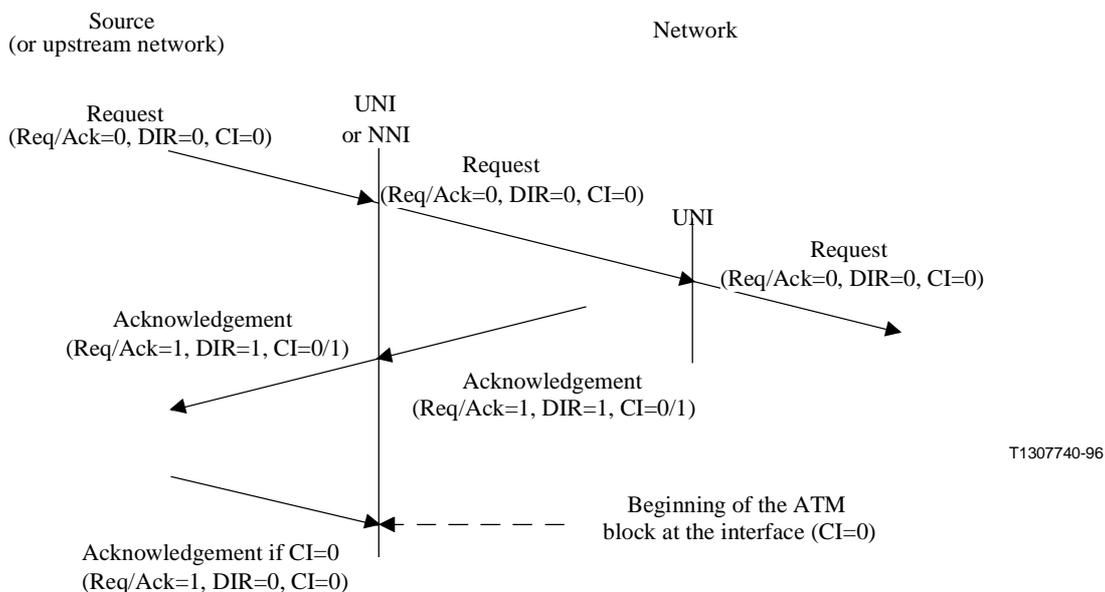


Figure C.2/I.371 – BCR increase initiated by the source
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

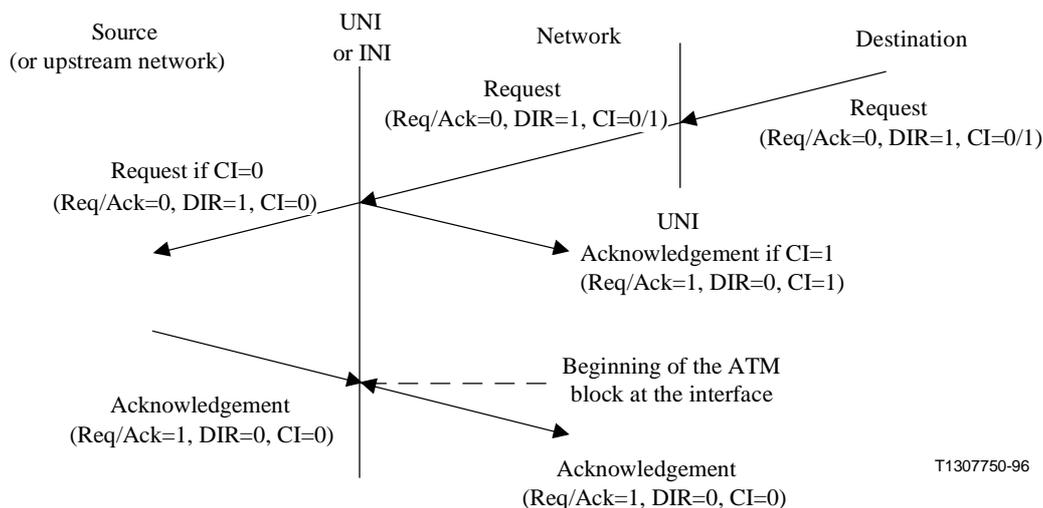


Figure C.3/I.371 – BCR modification initiated by the destination
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)

Superseded by a more recent version

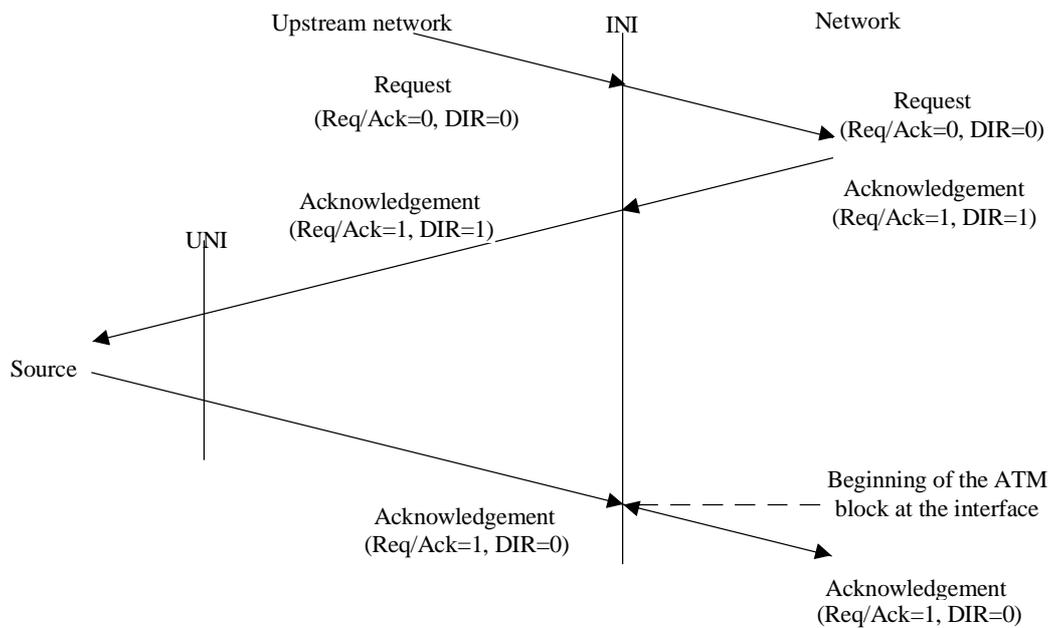


Figure C.4/I.371 – BCR modification initiated by the network in the forward direction
(RM cells have maintenance = 0, traffic management = 1, CI = 0, elastic/rigid = 1)

ANNEX D

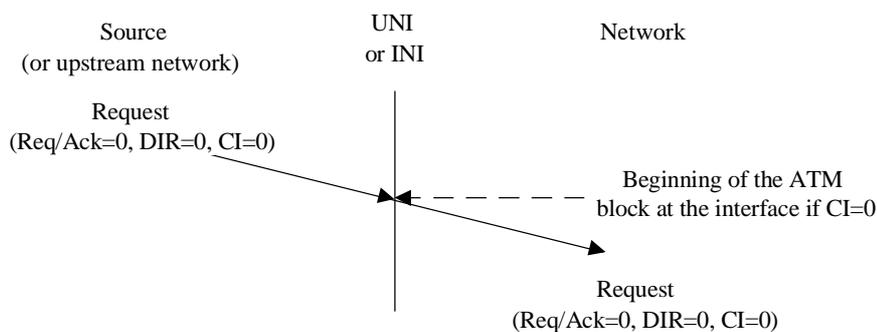
ABT/IT control messages across a standardized interface

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.

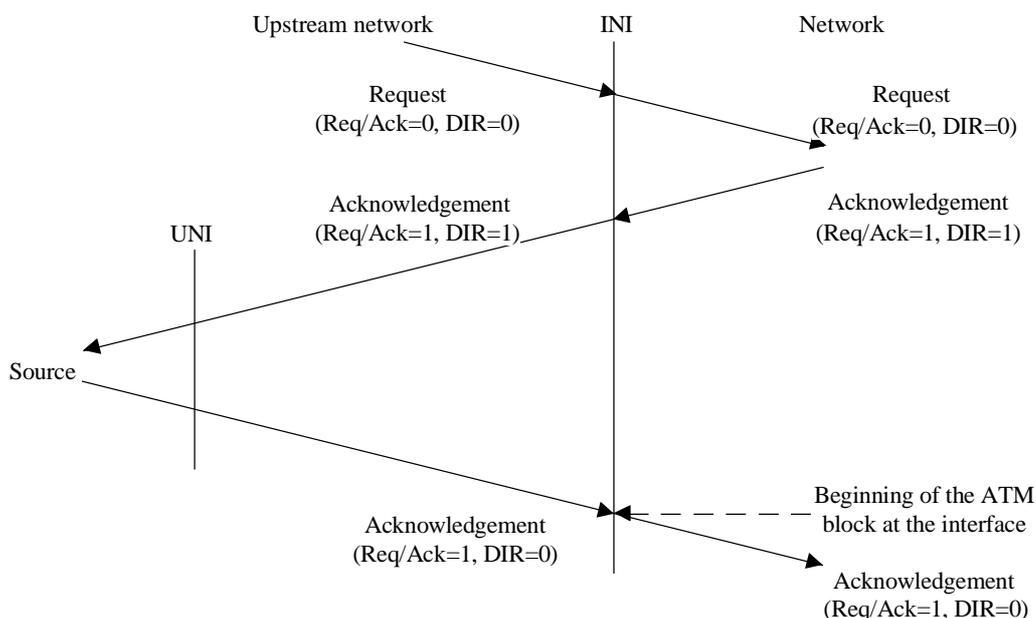
The beginning of an ATM block as illustrated in Figures D.1 and D.2 is also the end of the previous ATM block, if not the first. These figures currently only apply to the rigid mode (elastic/rigid bit set to 1).

Superseded by a more recent version



T1307770-96

Figure D.1/I.371 – BCR modification initiated by the source
(RM cells have maintenance = 0, traffic management = 0, elastic/rigid = 1)



T1307780-96

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

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 E.1 for the case of one RM cell outstanding on the bi-directional connection and Figure E.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 E.3.

Superseded by a more recent version

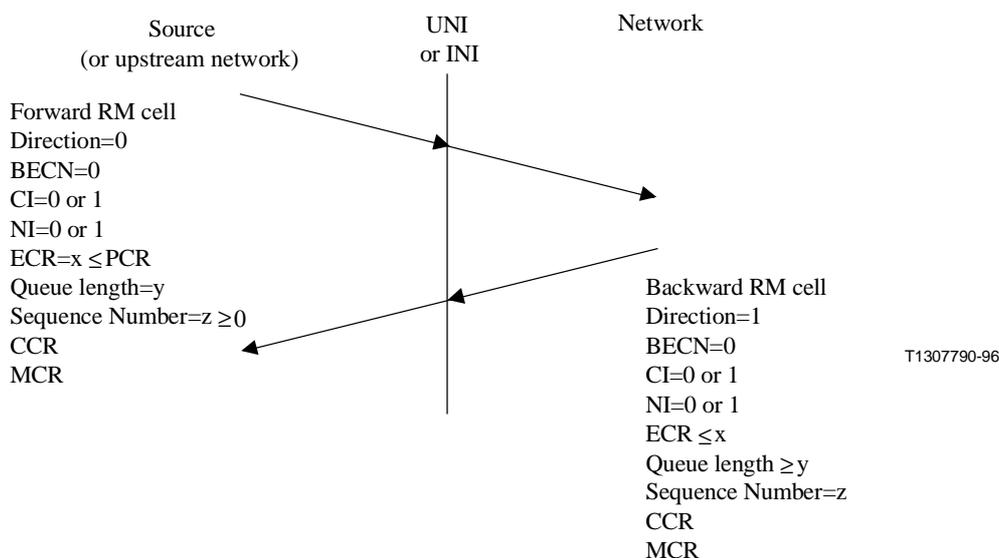
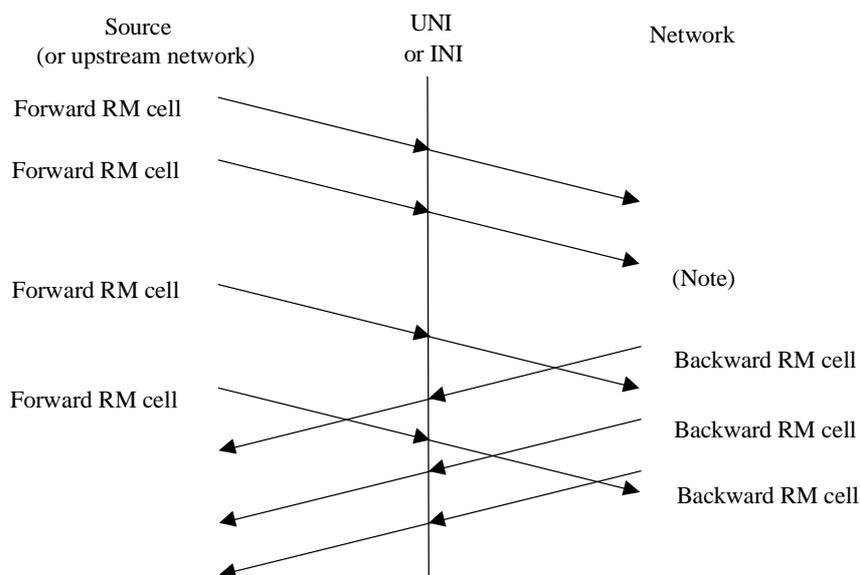


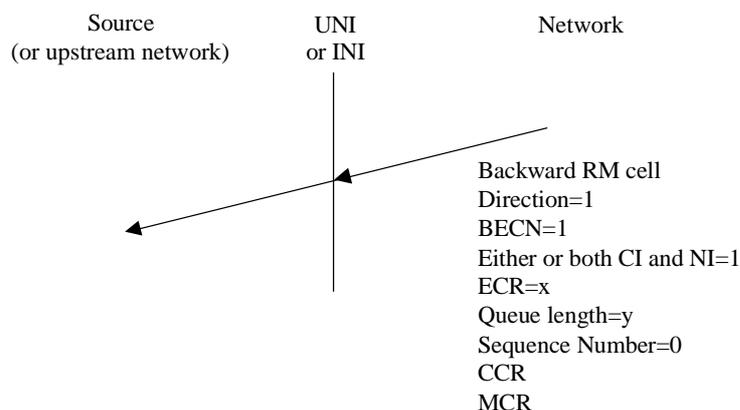
Figure E.1/I.371 – One RM cell outstanding in the bi-directional connection



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 E.2/I.371 – Multiple RM cells are outstanding in the bi-directional connection, and where, for illustration purposes, the destination overwrites an RM cell

Superseded by a more recent version



T1307810-96

Figure E.3/I.371 – RM cell on the backward connection originated by a network element or by the destination

ANNEX F

Traffic characteristics relevant to ATCs

Table F.1 lists traffic characteristics (source traffic descriptor and CDV tolerances) that are relevant to each ATM transfer capabilities as defined in 5.5.

An X in the table indicates that the traffic characteristic is negotiated for the corresponding ATC.

Table F.1/I.371

	DBR	SBR1	SBR2	SBR3	ABT/DT ABT/IT	ABR
PCR(0+1)	X	X	X	X	X	X (Note 4)
$\tau_{\text{PCR}}(0+1)$	X	X	X	X	X	(Note 5)
SCR(0)			X	X		
$\tau_{\text{IBT}}(0)$			X	X		
$\tau'_{\text{SCR}}(0)$			X	X		
SCR(0+1)		X			X (Note 3)	
$\tau_{\text{IBT}}(0+1)$		X			X (Note 3)	
$\tau'_{\text{SCR}}(0+1)$		X			X (Note 3)	
PCR(RM) + CDVT(RM)					X	(Note 6)
MCR(0+1)						X (Notes 7, 8)
Tagging			(Note 2)	X (Note 2)		
PCR(OAM) + $\tau_{\text{PCR}}(\text{OAM})$	X (Note 1)				X (Note 1)	

Superseded by a more recent version

Table F.1/I.371 (concluded)

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 5.5.3). If this option is not chosen, user OAM cells are aggregated to user data cells when characterizing the offered traffic.

NOTE 2 – SBR2 and SBR3 are identical except for the setting of the tagging indicator (see 5.5.4).

NOTE 3 – Declaration of an (SCR,IBT) parameter set for ABT capabilities is optional. If this parameter set is negotiated, there is a QOS commitment at the block level. If not negotiated, SCR is assumed to be 0 and there are no QOS commitments at the block level (see 5.5.5).

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 5.5.6).

NOTE 5 – It is necessary to define a default value for CDVT(0+1) at each standardized interface; that is valid for any ABR connection that has not implicitly or explicitly negotiated such a value (see 5.5.6).

NOTE 6 – In ABR, user RM cells with CLP bit set to 0 are aggregated to user data. User RM cells with CLP bit set to 1 are not aggregated with user data traffic and this traffic stream may be separately characterized and controlled.

NOTE 7 – User-generated traffic should be transmitted with CLP set to 0. The MCR may be set to 0.

NOTE 8 – Since MCR is a rate at which the user is allowed to continuously transmit, it is necessary to define a default corresponding value for CDV tolerance, for each standardized interface, that is valid for any ABR connection that has not implicitly or explicitly negotiated such a value (see 5.5.6).

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 .

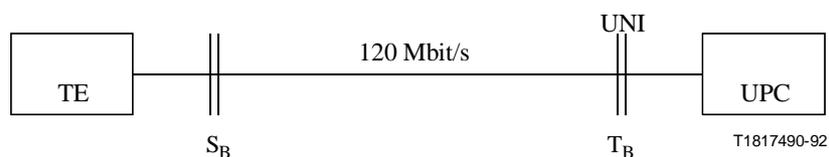
The terminology used here is the one from the virtual scheduling algorithm as shown on Figure A.1.

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 $T = 1.25 \Delta$. This corresponds to a peak bit rate of 120 Mbit/s.

Superseded by a more recent version



Peak emission interval $T = 1.25 \Delta$
 Peak cell rate $= 1/T$
 CDV tolerance τ needed at $T_B = 0.75 \Delta$

Figure I.1/I.371 – Traffic configuration 1

Figure I.2/I.371 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.

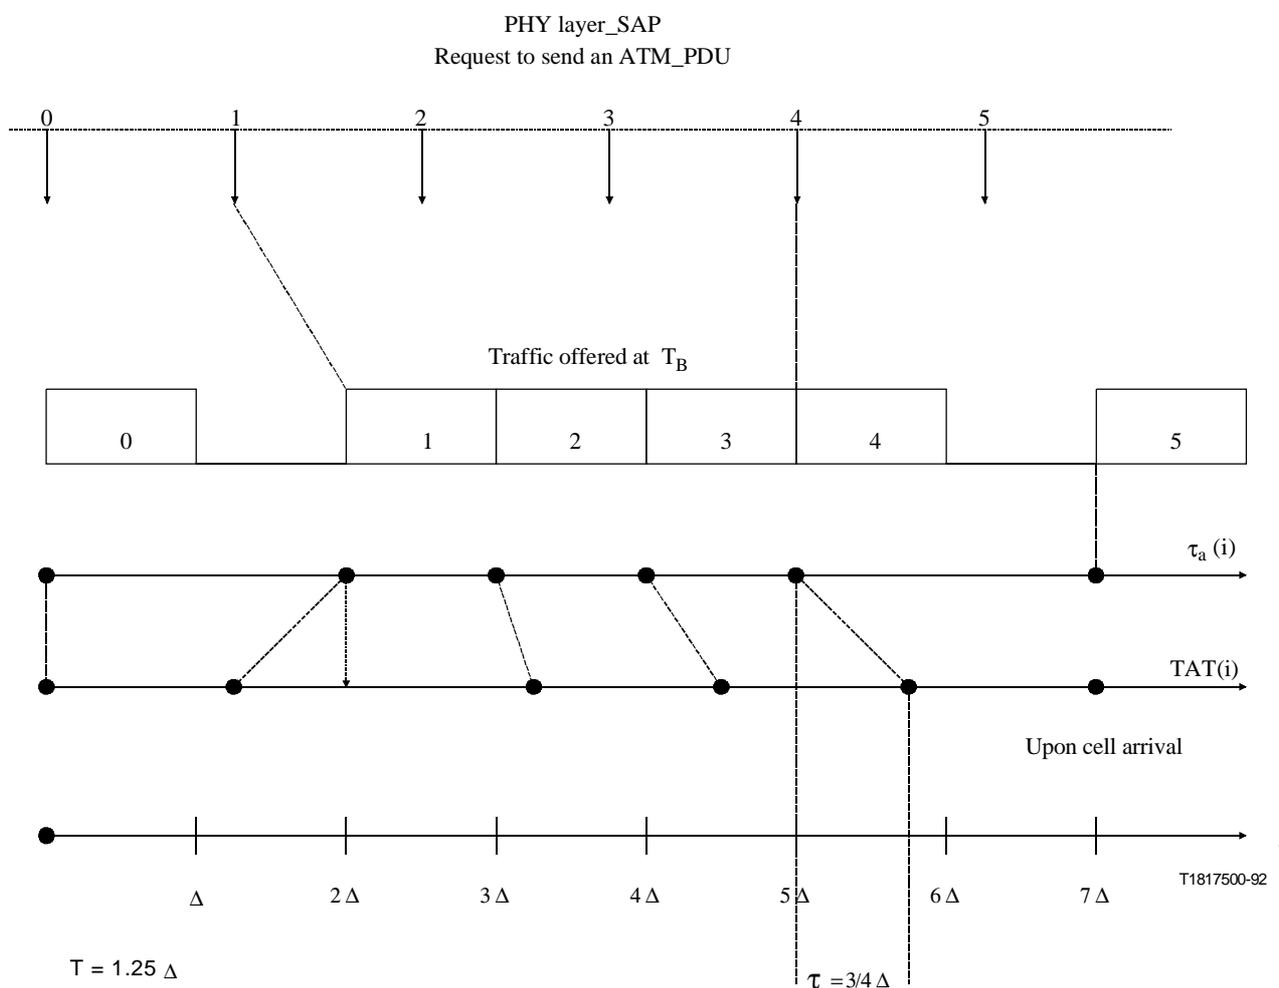


Figure I.2/I.371 – Illustration of CDV tolerance τ for traffic configuration 1

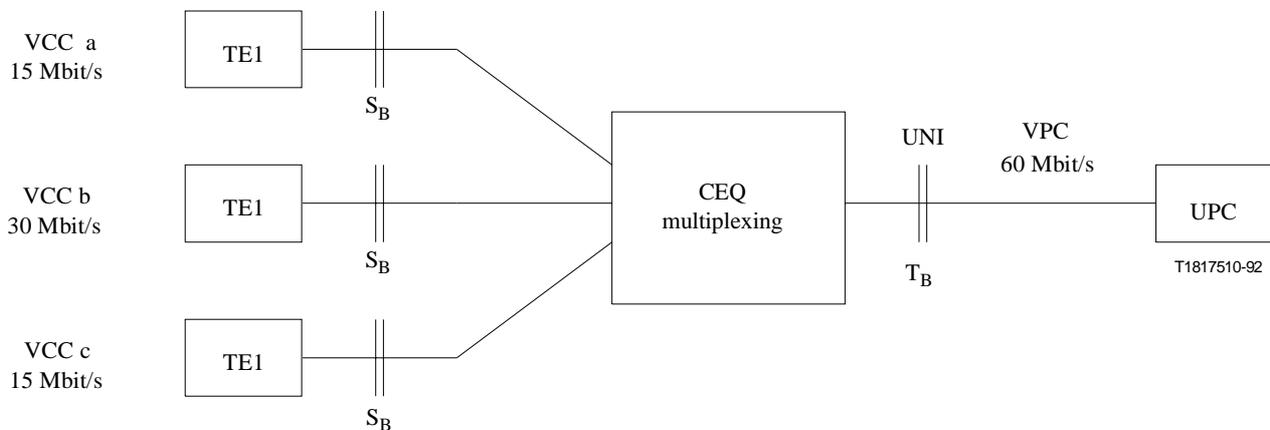
Configuration 2

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.

Superseded by a more recent version

Terminals generate ATM_PDU Data_Requests every 10Δ , 5Δ and 10Δ 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 $T = 2.5 \Delta$, which corresponds to a peak bit rate of 60 Mbit/s.



Peak emission interval $T = 2.5 \Delta$
Peak cell rate = $1/T$
CDV tolerance τ needed at $T_B = 3 \Delta$

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.

Superseded by a more recent version

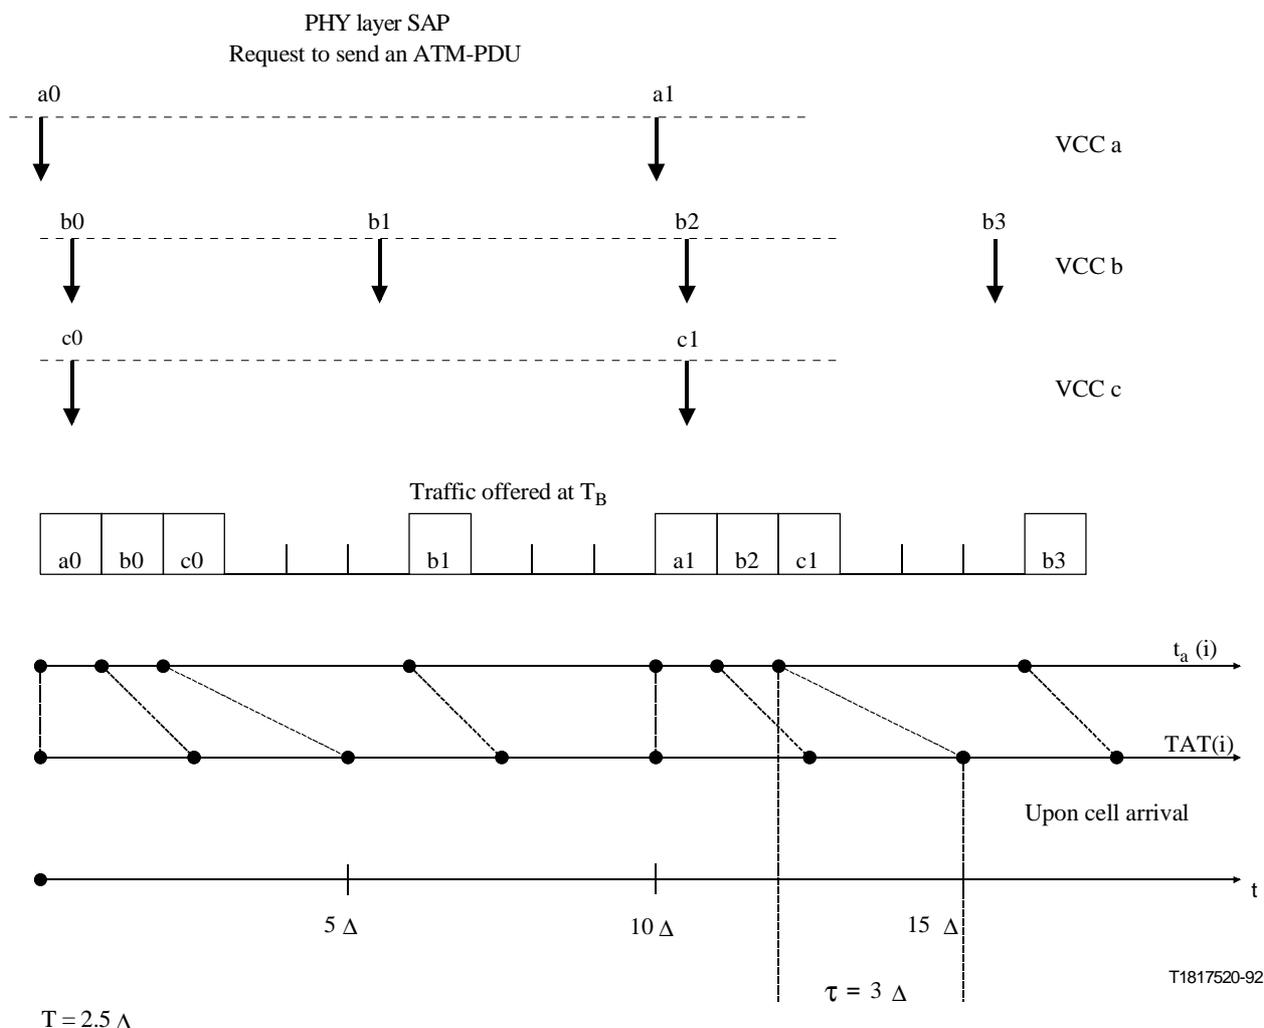


Figure I.4/I.371 – Illustration of CDV tolerance τ for traffic configuration 2

APPENDIX II

Transcoding rules from signalling information onto OAM traffic parameters at the ATM layer

The present release of Recommendation 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, Recommendation 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 $\tau_{PCR}(data)$ 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

Superseded by a more recent version

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 Recommendation I.610.

The three quantities $T_{PCR}(\text{agg})$, $\tau_{PCR}(\text{data})$ 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:

$$T_{PCR}(\text{OAM}) = \tau_{PCR}(\text{OAM}) = \frac{T_{PCR}(\text{agg})}{P_{OAM}}$$

$$T_{PCR}(\text{user}) = \frac{T_{PCR}(\text{agg})}{1 - P_{OAM}}$$

In case of an aggregate conformance definition, the default values for traffic descriptors are given as:

$$\tau_{PCR}(\text{agg}) = T_{PCR}(\text{agg}) + \tau_{PCR}(\text{data})$$

APPENDIX III

Throughput behaviour of the generic cell rate algorithm

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 characterised 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 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:

Superseded by a more recent version

- 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 ($\approx 1\%$). In general, one can show that the following relation approximately holds:

$$\text{for } \Delta \ll 1 \text{ and } \tau \leq T_c, DR \approx \frac{T_c}{\tau} \Delta \text{ and for } \tau \geq T_c, DR = \Delta$$

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 6.2.3.2.1 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.

Superseded by a more recent version

ITU-T RECOMMENDATIONS SERIES

Series A	Organization of the work of the ITU-T
Series B	Means of expression
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Telephone network and ISDN
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media
Series H	Transmission of non-telephone signals
Series I	Integrated services digital network
Series J	Transmission of sound-programme and television signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	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
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminal equipments and protocols for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communication
Series Z	Programming languages