INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

I.371.1

(06/97)

SERIES I: INTEGRATED SERVICES DIGITAL NETWORK

Overall network aspects and functions – General network requirements and functions

Traffic control and congestion control in B-ISDN: conformance definitions for ABT and ABR

ITU-T Recommendation I.371.1

(Previously CCITT Recommendation)

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.

ITU-T RECOMMENDATION I.371.1

TRAFFIC CONTROL AND CONGESTION CONTROL IN B-ISDN: CONFORMANCE DEFINITIONS FOR ABT AND ABR

Summary

Recommendation I.371 states that conformance definitions for ABT and ABR are required.

This Recommendation specifies the ABT/DT and ABT/IT conformance definitions and the ABR conformance definition for the explicit rate mode. Appendices provide information on source, destination and network reference behaviours for ABR explicit cell rate and binary modes.

Source

ITU-T Recommendation I.371.1 was prepared by ITU-T Study Group 13 (1997-2000) and was approved under the WTSC Resolution No. 1 procedure on the 20th of June 1997.

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.

CONTENTS

Recommendation I.371.1

TRAFFIC CONTROL AND CONGESTION CONTROL IN B-ISDN: CONFORMANCE DEFINITIONS FOR ABT AND ABR

(Geneva, 1997)

1 Scope

This Recommendation completes the specification of ATM Transfer Capabilities provided in Recommendation I.371. It provides conformance definitions for the following ATM Transfer Capabilities: ABT/DT, ABT/IT, ABR for the explicit rate mode.

– The main body specifies conformance definitions for ABT/DT, ABT/IT, ABR for the explicit rate mode.

– Annex A describes how multiple renegotiations in ABT should be handled.

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.371 (1996), *Traffic control and congestion control in B-ISDN*.

[2] ITU-T Recommendation I.356 (1996), *B-ISDN ATM layer cell transfer performance.*

[3] ITU-T Recommendation I.610 (1995), *B-ISDN operation and maintenance principles and functions*.

3 Abbreviations

This Recommendation uses the following abbreviations:

ABR	Available Bit Rate
ABT	ATM Block Transfer
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
BECN	Backward Explicit Congestion Notification
CBR	Constant Bit Rate
CCR	Current Cell Rate

CDV	Cell Delay Variation
CEQ	Customer Equipment
CI	Congestion Indication
CLP	Cell Loss Priority (bit)
CLR	Cell Loss Ratio
CTD	Cell Transfer Delay
DBR	Deterministic Bit Rate
DGCRA	Dynamic GCRA
DIR	Direction
DT	Delayed Transmission
ECR	Explicit Cell Rate
EFCI	Explicit Forward Congestion Indication
FIFO	First-In First-Out
FRM	Fast Resource Management
GCRA	Generic Cell Rate Algorithm
GFC	Generic Flow Control
IACR	Initial Allowed Cell Rate
IBT	Intrinsic Burst Tolerance
INI	Inter-Network Interface
IT	Immediate Transmission
ITT	Ideal Transmission Time
LVMT	Last Virtual Modification Time
LVST	Last Virtual Schedule Time
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
NE	Network Element
NI	No Increase
NPC	Network Parameter Control
OAM	Operation And Maintenance
PACR	Potential Allowed Cell Rate
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PEI	Peak Emission Interval
PHY	Physical Layer
PTI	Payload Type Indicator
QOS	Quality of Service

RDF	Rate Decrease Factor
RIF	Rate Increase Factor
RM	Resource Management
SAP	Service Access Point
SBR	Statistical Bit Rate
SCR	Sustainable Cell Rate
SDU	Service Data Unit
SN	Sequence Number
TAT	Theoretical Arrival Time
TBE	Transient Buffer Exposure
UNI	User-Network Interface
UPC	Usage Parameter Control
VBR	Variable Bit Rate
VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VD	Virtual Destination
VPC	Virtual Path Connection
VPI	Virtual Path Identifier
VS	Virtual Source
VSA	Virtual Scheduling Algorithm

4 High level description of ATM transfer capabilities (to complete 5.5.2/I.371)

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

Deterministic transfer capability – DBR

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

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

For a complete specification of the DBR ATC, refer to 5.5.3/I.371.

Statistical bit rate transfer capability – SBR

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

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

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

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

For a complete specification of the SBR ATCs, refer to 5.5.4/I.371.

ATM block transfer – ABT

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

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

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

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

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

SCR is specified, QOS commitments at the ATM block level are in terms of block loss ratio. In this respect, ABT/IT implements frame discard.

Selective cell discard on the basis of the CLP bit and tagging do not apply to ABT.

Subclause 5.5.5/I.371 describes the ABT/DT and ABT/IT service models, specifies the ABT RM cell format and the types of messages exchanged at standardized interfaces. This Recommendation specifies the ABT conformance definition in clauses 5 and 6.

Available Bit Rate – ABR

The ABR transfer capability is intended to support non-real time elastic applications that may adapt to the instantaneous bandwidth available within the network. In such a case, the network may share the available resources between connections supporting such applications. ABR uses static parameters declared at connection set-up and dynamic parameters renegotiable via resource management procedures based on RM cells.

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

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

In ABR, QOS commitments are in terms of cell loss ratio for CLP=0 cells. Subclause 5.5.6/I.371 describes the ABR service model, specifies the ABR RM cell format and the types of messages exchanged at standardized interfaces.

Clause 6 specifies the conformance definition for ABR in the explicit rate mode only. Reference source and destination behaviours to network indications are provided for both the explicit rate mode and the binary mode in Appendix III. In the binary mode, no QOS commitments can be made, but QOS indications in terms of cell loss ratio may be provided to connections respecting source and destination reference behaviours.

5 Conformance for ABT/DT

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

Both conformance definitions depend on RM cells that pass the interface. General principles for the ABT/DT conformance definition are described in 5.1.

5.1 General principles of conformance definition for ABT/DT

The control messages that define the ATM block for ABT/DT at an interface are described in Annex C/I.371.

RM cells delineating ATM blocks in the forward direction are:

1) either BCR decrease RM cells sent by the source (TM=0);

2) or acknowledgement RM cells sent by the source in response to:

   - a positive acknowledgement sent by the network following a BCR increase request by the source;
   - a BCR modification initiated by the destination or by the network.

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

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

Conformance for an ABT connection is tested against:

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

5.2 Conformance of RM cells

5.2.1 Conformance of user-generated RM cells

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

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

1) It is the response of the source to either an acknowledgement RM cell or a request RM cell sent by the network to the source (see Annex C/I.371).
2) It arrives within a time-out interval after the RM cell sent by the network to the source it responds to has crossed the interface. The time-out value depends on the round-trip time from the interface to the source. This value is either determined by the network operator or, if applying to an INI, negotiated between network operators. It may be specified on a subscription basis or on a per-connection basis.
3) It conveys information (BCR values, Sequence Number, CI bit, etc.) consistent with the message sent by the network. In particular, valid BCR values are BCR values less than or equal to the BCR values conveyed by the RM cell sent by the network to the source.

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

5.2.2 Conformance of network-generated RM cells

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

5.3 Dynamic GCRA for ABT/DT (to complete 5.5.5.1.4/I.371)

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

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

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

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

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

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

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

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

The cell conformance algorithm is depicted in Figure 1.

The following notation is used in Figure 1:

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

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

$\tau(\lambda(x))$ CDV tolerance used to test conformance of cell flow x for the allocated BCR $\lambda(x)$, the function $\tau(\lambda)$ is specified at connection establishment for user data cell flows, a unique value may be specified; for OAM traffic $\tau(\lambda)$ should be consistent with the standardized default rule specified for OAM traffic (see Appendix II/I.371); if $\lambda=0$, τ takes a default value;

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

$T(\lambda)$ peak emission interval corresponding to the BCR λ in the standardized ATM layer peak cell rate granularity list given in 5.4.1.2/I.371; if $\lambda=0$, T takes a default value equal to the maximum value supported by the network;

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

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

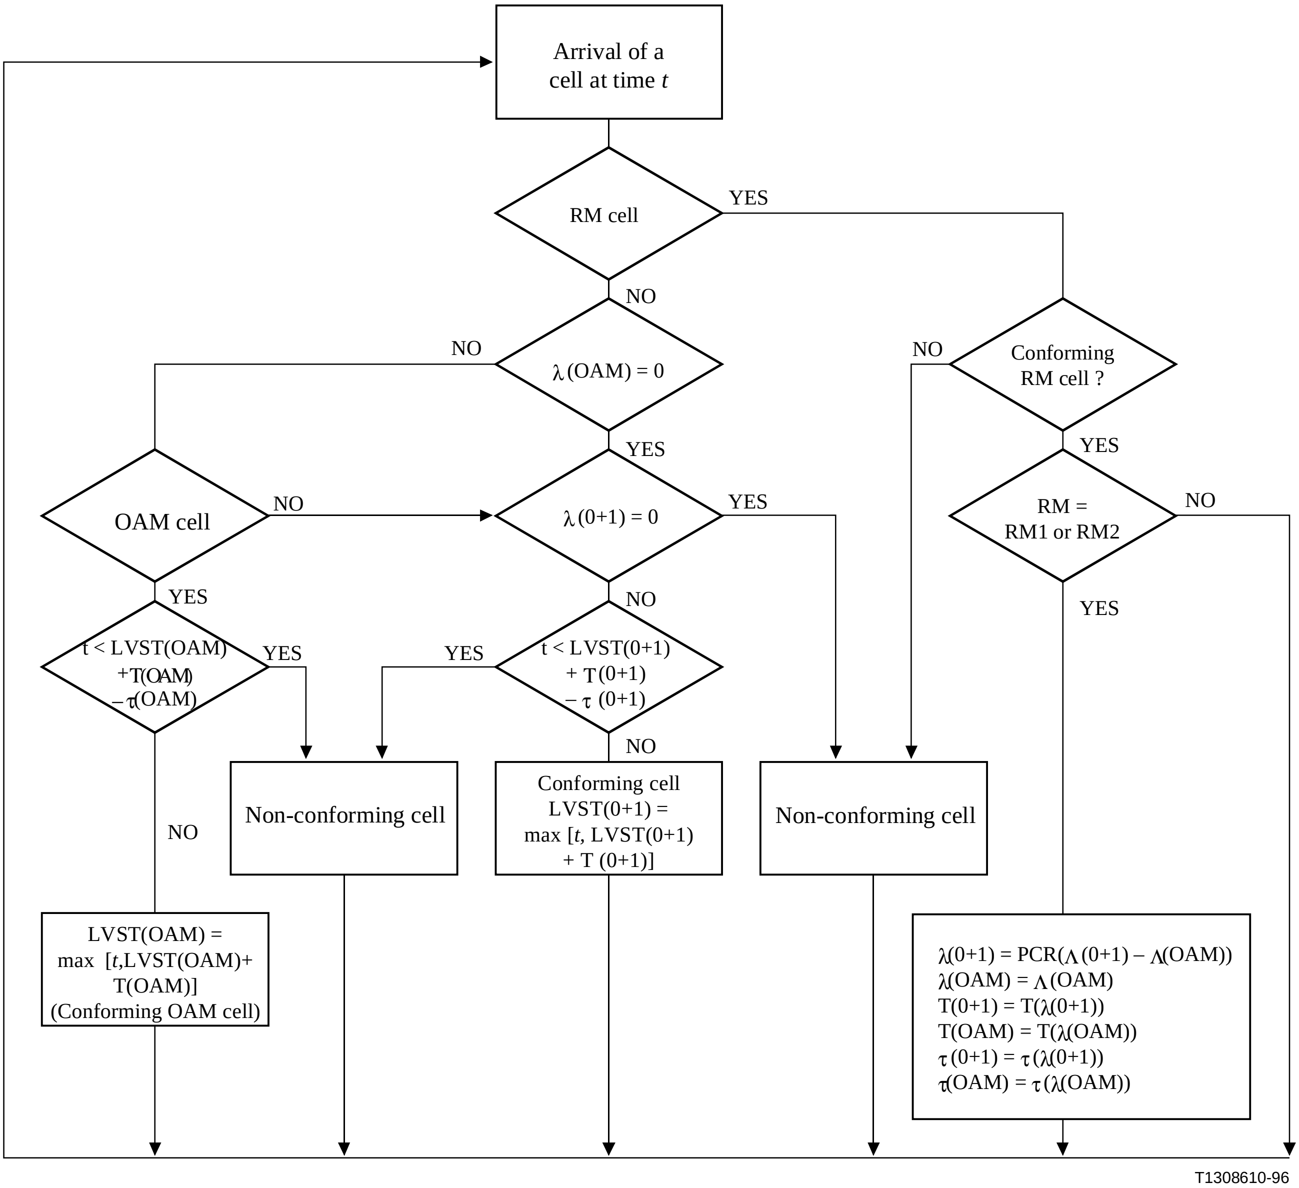

NOTE 1 – LVST(0+1) and LVST(OAM) are initialized to $-\infty$, which is a default value for identifying the first cell of an ATM connection; $\lambda(0+1)$ and $\lambda(\text{OAM})$ are initialized to 0.

NOTE 2 – By definition, $\Lambda(0 + 1) \geq \Lambda(\text{OAM})$; otherwise, the peak cell rate re-negotiation would be invalid.

Figure 1/I.371.1 – Cell conformance for an ABT/DT connection

5.4 ATM block conformance for ABT/DT

ATM block conformance is tested against the sustainable cell rate, if greater than 0, that is specified for the CLP=0+1 cell flow. ATM block conformance testing relies on an algorithm which computes a number of credits. ATM blocks are non-conforming when the number of credits is null. Moreover, the ATM block conformance algorithm makes use of a virtual time u defined at the arrival time of a cell as the maximum between this arrival time and the LVST of the last conforming CLP=0+1 data cell, which is computed by the cell level conformance algorithm (see 5.3). More precisely, $u = \max\{\text{LVST}, t\}$ where t is the current time.

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

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

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

where:

1) $1/T$ is the peak cell rate of the connection and T_{SCR} is the emission interval corresponding to Λ_{SCR};
2) it is assumed that forward and backward user request RM cell flows issued by both users of the ABT/DT communication are at the interface considered conforming to GCRA(T_{RM}, τ_{RM}) and GCRA(T'_{RM}, τ'_{RM}), respectively;
3) τ''_{SCR} is the difference between the maximum and the minimum (or equivalently remote quantiles) of the virtual transfer delays for RM cells delineating ATM blocks. The virtual transfer delay for an RM cell delineating an ATM block is defined as the difference between the time when the RM cell is transmitted at the PHY-SAP of the equivalent terminal and the virtual time *u* when it is received at the interface;
4) Δ is the cell transmission time (in seconds) at the interface link speed.

ATM block conformance is tested by taking into account the volume of reserved resources. The ATM block conformance testing algorithm is depicted in Figure 2/I.371. The principles of ATM block conformance are as follows (see Figure 3/I.371):

– Specific RM cells are RM cells of RM_1 and RM_2 types defined above.
– ATM block conformance is tested slot by slot by comparing the current time *t* with a non-conformance date t_{exp}.
– The non-conformance date is relevant only when the allocated BCR λ is greater than the sustainable cell rate $\Lambda_{SCR} = 1/T_{SCR}$.
– The non-conformance time t_{exp} is computed by using a variable *X*, the sustainable cell rate Λ_{SCR}, and the BCR λ allocated to the cell flow.
– *X* is updated at each specific RM cell arrival time and represents the number of credits for the new BCR reservation (*X* is computed by using the BCR allocated to the previous BCR reservation).
– Upon reception of a forward bandwidth acknowledgement RM cell with the Traffic Management and Maintenance bits set equal to 1 and 0, respectively, the number *X* of credits is reset to 0. This is intended to realign ATM block conformance algorithms when a policing procedure is run.
– The non-conformance time t_{exp} and variable *X* are computed from the maximum of the current time and the Last Virtual Schedule Time (LVST) of the cell flow (LVST is computed by the cell conformance test for the CLP=0+1 cell flow.

The following relations hold:

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

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

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

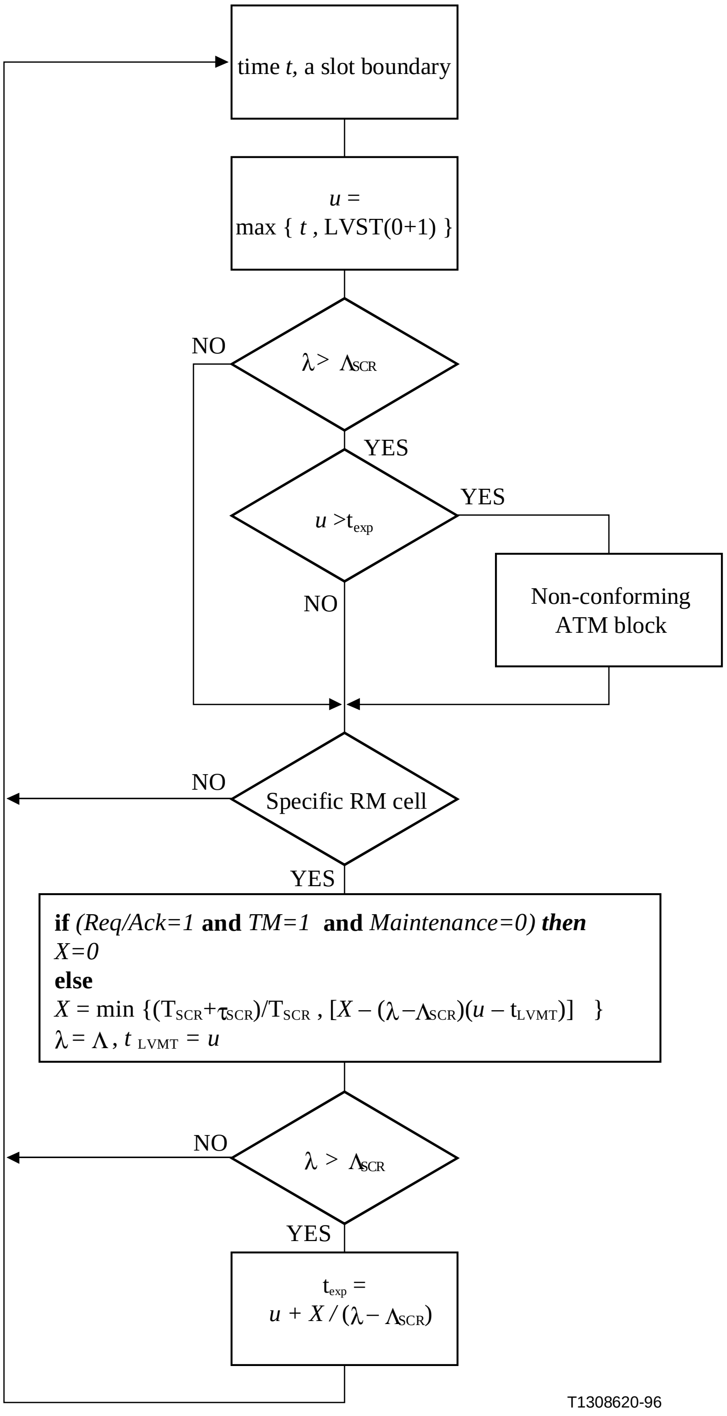

NOTE 1 – t_{LVMT} and u are initialized to $-\infty$, which is a default value for identifying the first cell of an ATM connection, X is initialized to $\frac{T_{SCR} + \tau_{SCR}}{T_{SCR}}$, λ is initialized to 0.

NOTE 2 – BCRs and LVST are those of the cell flow considered, LVST is given by the cell conformance algorithm (see Figure 1).

Figure 2/I.371.1 – ATM block conformance for a cell flow of an ABT/DT connection

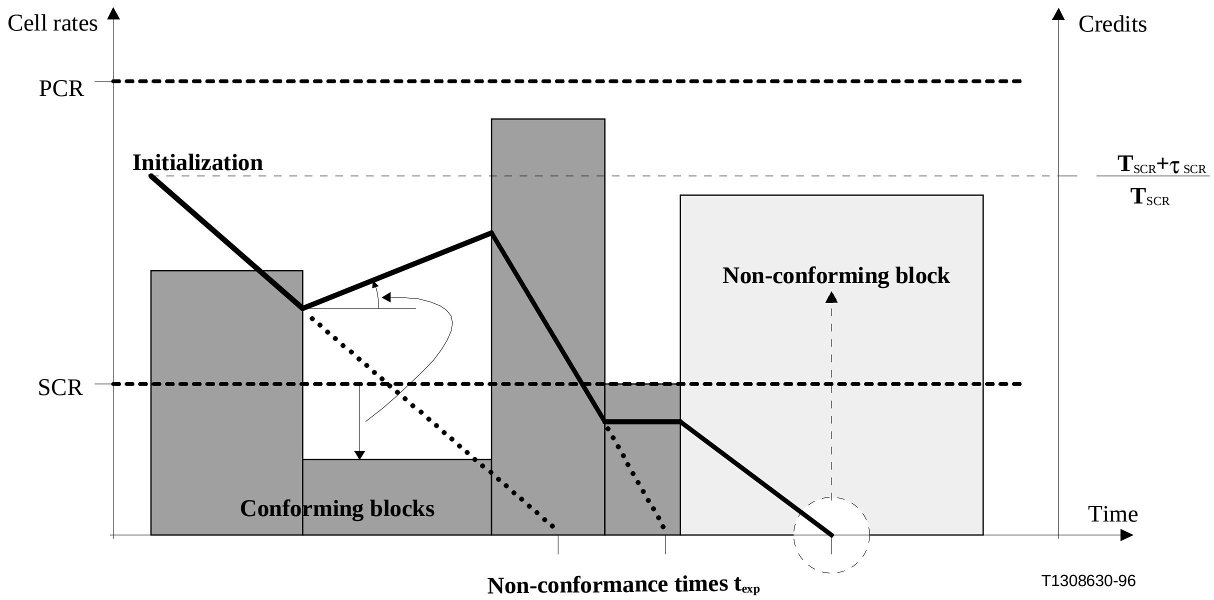

NOTE – This figure is for illustration purposes. The bold solid line represents the instantaneous number of credits. The slope of this curve is equal to the difference between the SCR and the BCR. Computation of t_{exp} for a given ATM block is based on the number of credits available at the ATM block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

Figure 3/I.371.1 – Examples of evolution of block conformance variables

Loss of RM cells may unduly produce block non-conformance and, in the case of lost RM_1 or RM_2 cells, may require recovery or reinitialization of conformance variables. ATM block conformance algorithms are realigned by policing procedures (see 6.2.3.6/I.371).

6 Conformance for ABT/IT

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

6.1 Cell conformance for ABT/IT

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

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

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

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

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

6.2 ATM block conformance for ABT/IT

The ATM block level conformance algorithm for ABT/IT is identical to that for ABT/DT (given by Figure 2/I.371), except that the specific RM cells to take into account are conforming bandwidth increase or decrease request RM cells sent by the source (TM=0) and acknowledgement RM cells (TM=1) sent by the source on the forward direction. Furthermore, the sustainable cell rate Λ_S and the tolerance τ_{SCR} taken into account in the ATM block conformance definition, are those valid at the interface considered and deduced from the sustainable cell rate Λ_S^0 and the maximum burst size MBS^0 negotiated at connection establishment (see Appendix II) as:

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

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

where the notation of 5.3 is used.

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

7 Conformance for ABR (to complete 5.5.6.4/I.371)

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

NOTE – Although RM cells with CLP=0, including BECN cells, are included in the current allowed cell rate (according to 5.5.6.1/I.371), BECN cells are excluded from the flow tested by the conformance definition. It follows that, where equipment emit BECN cells as part of the current allowed cell rate flow, these BECN cells will not cause any loss of conformance.

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

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

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

If some cells are non-conforming to some of the relevant conformance tests, the network may consider the connection as non-compliant (see 5.3.2/I.371). If the network chooses to offer QOS commitments to a connection with some non-conforming cells, the ATM layer QOS is only assured to a volume of cells that is conforming to all relevant conformance tests. The precise definition of a compliant ABR connection is left to the network operator. Any definition of a compliant ABR connection shall find a connection compliant if all cells on the connection are conforming and if the RM cells on the connection satisfy the requirements, if any, of the mechanism implemented by the network operator(s).

For compliant connections at the UNI or inter-network interface, the agreed QOS class shall be supported for at least the number of cells equal to the conforming cells according to the conformance definition.

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

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

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

Note that in the ABR Capability, a source is not required to send RM cells. However, if there is not a flow of user-generated, backward RM cells, and if the network wishes to convey feedback to the user, the network can make use of the capability of itself generating backward (BECN) RM cells (see 5.5.6.3.1/I.371).

7.1 Definitions of ABR delays used in the conformance definition

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

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

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

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

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

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

- The delay t_1 denotes the time from a cell's transmission time by the traffic source to its receipt at the interface in question.

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

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

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

The parameters τ_1, τ_2 and τ_3 are specified at the given interface for the given connection. (Note that for simplicity, τ_3 could be set to zero, though with the consequence of a less tight conformance definition.) The conformance definition in 7.3 makes use of these parameters, as well as the ACRs determined by backward RM cells on the companion backward connection.

7.2 Requirements on the ABR conformance definition

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

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

7.3 ABR conformance algorithm

7.3.1 Dynamic Generic Cell Rate Algorithm (DGCRA) for ABR

The conformance definition is based on the dynamic GCRA. The dynamic GCRA (DGCRA) is an extension of the GCRA defined in Annex A/I.371. The DGCRA differs from the GCRA in that the increment T changes with time, as determined by ABR feedback information conveyed on the corresponding backward connection.

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

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

At the arrival time $t_a(k)$ of the k^{th} cell, the DGCRA first calculates $T(k)$ (see 7.3.2) and then checks the cell's conformance and updates its own Last Virtual Scheduling Time (LVST) as follows:

Initialize:

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

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

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

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

else *# cell is non-conforming*

do not update algorithm state.

$T_{old} = T(k)$

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

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

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

The sequence $\{T(k),\ k \geq 1\}$ of increments, which are successively used at arrival times $\{t_a(k),\ k \geq 1\}$ of cells at the interface, depends on the feedback information in the backward RM cell sent across the interface at departure times $\{t_b(j),\ j \geq 1\}$ on the backward connection (see 7.3.2). Each backward RM cell determines an allowed cell rate that could apply to some future cells on the forward direction.

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

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

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

In mode 1 (explicit rate mode), the ECR field is the only field in the pertinent backward RM cells (see 7.3.2 for a definition of the set of pertinent backward RM cells) that is used in the calculation of $T(k)$. Conformance to mode 1 is specified in this Recommendation.

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

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

Subclause 7.3.2 presents the algorithm for determining the sequence of increments $\{T(k),\ k \geq 1\}$ for the case of explicit rate mode. This algorithm is recognized to be less than optimal with respect to its tightness in order to reduce complexity. Improved algorithms for explicit rate mode are for further study.

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

The algorithm is written in a format that determines the ACR at the interface as a continuous-time variable whose reciprocal at time $t_a(k)$ is *T(k)*; if ACR happens to be computed to a value less than 1 cell/s, ACR is set to 1 cell/s. The increment for the DGCRA at each cell arrival in the forward direction is thus determined.

NOTE – At a given time instant, the ACR that is valid at the interface may differ from the ACR considered by the source to be valid. This may be for example because of a time lag, or because some backward RM cells considered at the interface have not reached the source.

The following algorithm computes two sets of counters (t_first,PACR_first) and (t_last,PACR_last). PACR_max is an auxiliary variable defined as Max(PACR_first,PACR_last).

- t_first is the time at which *T*(*k*) is scheduled to be set to 1/PACR_first.
- If different from t_first, t_last is the planned update for t_first, at expiration of t_first; at that time PACR_first is updated to PACR_last.

PACR_first and PACR_last are determined on the basis of the value PACR(j) of the ECR field carried in pertinent RM cells. Pertinent RM cells are backward RM cells with correct CRC-10 in the EDC field (see 7.1/I.371) that are either non-BECN cells, or BECN cells with ECR<PACR_last.

The algorithm described below has the following characteristics:

- At most, two rate modifications can be scheduled, which can be either increases or decreases from the current ACR.
- Since t_first, t_last, PACR_first and PACR_last are potentially updated every time a backward RM cell is observed in the backward direction, a given value of PACR_first or PACR_last may never be used in the DGCRA since, prior to its scheduled time to be applied, it can be revised by another backward RM cell.
- If less than two rate updates are scheduled, t_first=t_last and PACR_first=PACR_last.
- If no rate update is scheduled, PACR_first=PACR_last=ACR and t_first=t_last <tb(j).
- If at least one rate update is scheduled (PACR_first≠ACR), t_first cannot be delayed by a later rate update, and PACR_first can only be increased.
- At any time, PACR_last carries the ECR value of the last pertinent cell that has crossed the interface.
- If the ECR of a new pertinent cell is equal to PACR_last, no update takes place.
- $MCR \leq PACR\_first \leq PCR$ and $MCR \leq PACR\_last \leq PCR$.
- $tb(j) \leq t\_first \leq t\_last \leq tb(j) + \tau_2$ if at least one rate update is scheduled.
- If ACR<PACR_first, $t\_first \leq tb(j) + \tau_3$.
- If PACR_first<PACR_last, $t\_last \leq tb(j) + \tau_3$.

Adjustment of ACR(t) based on the ECR field in backward RM cells

- **Initialization:**
  t_first=t_last=0
  PACR_max=PACR_first=PACR_last=IACR

```
•       At each tb(j) which is the arrival time of a pertinent RM cell:
compute PACR(j) = min( PCR, max(MCR, ECR in backward RM cell)
if PACR(j) ≠ PACR_last:                                   # else no update takes place
        if (t_first>tb(j))                                # is the scheduling list non-empty?
        # start update of a non-empty scheduling list
                if (PACR(j)≥PACR_max)                     # PACR(j) is an increase
                                                          # over the current PACR_max
                # start processing an increase
                        PACR_max=PACR(j)                  # update PACR_max

                        if (tb(j)+τ3>t_first)             # t_first and PACR_first unchanged
                                if ((t_first=t_last) or (t_last>tb(j)+τ3))
                                        t_last=tb(j)+τ3   # else t_last in unchanged
                                endif
                        endif                             # endif (tb(j)+τ3>t_first)

                        else                              # tb(j)+τ3≤t_first
                                PACR_first=PACR(j)        # update PACR_first
                                if (PACR(j)≥ACR)          # PACR(j) is an increase over ACR
                                        t_first=tb(j)+τ3  # else t_first is unchanged.
                                endif                     # endif (PACR(j)≥ACR)
                                t_last=t_first            # a single rate update is scheduled
                        endelse                           # endelse (tb(j)+τ3≤t_first)
                endif
                # end processing an increase
                else                                      # PACR(j) is a decrease
                                                          # over PACR_max
                # start processing a decrease
                        PACR_first=PACR_max               # schedule highest rate at t_first
                        if (PACR(j)<PACR_last)            # PACR(j) is a decrease
                                                          # over PACR_last
                                t_last=tb(j)+τ2           # t_last is delayed
                        endif                             # else t_last is unchanged
                endelse
                # end processing a decrease
                PACR_last=PACR(j)                         # store new rate in PACR_last
        endif
        # end update of a non-empty scheduling list
        else                                              # the scheduling list is empty
        # start update of an empty scheduling list
                if (PACR(j)>ACR)
                        t_first=tb(j)+τ3                  # an increase is scheduled (τ3 lag)
                else
                        t_first=tb(j)+τ2                  # a decrease is scheduled (τ2 lag)
                t_last=t_first                            # a single rate update is scheduled
                PACR_max=PACR_first=PACR_last=PACR(j)
        endelse
        # end update of an empty scheduling list
endif                                                     # endif for PACR(j)≠PACR_last
•       At expiration of t_first:
                ACR=PACR_first                            # update ACR
                t_first=t_last                            # update t_first
                PACR_first=PACR_last                      # update PACR_first
                PACR_max=PACR_last
```

End of Adjustment of ACR(t) based on the ECR field in backward RM cells.

ANNEX A

Avoidance of multiple outstanding BCR negotiations

To avoid multiple outstanding network-generated BCR negotiations in a network, the following priority principles between different network-generated BCR negotiations within a given network are introduced as follows:

1) A BCR negotiation request initiated by an upstream network has priority over any BCR negotiation initiated by the network considered or a downstream network. Given this priority principle, if a BCR negotiation with lower priority is pending in the network considered, then this network should interrupt the BCR negotiation with low priority and enable the higher priority level BCR negotiation to be processed.

2) If a BCR negotiation has been initiated by the network considered or an upstream network, then the network considered should deny any BCR negotiation request issued by a downstream network.

For the implementation of the above priority principles between BCR negotiations, it is desirable that two BCR negotiations processed by a given network are not identified by the same sequence number. The sequence number of the response given by a network following a BCR request should be compatible with the sequence number of the request and the priority principles between network-generated BCR negotiations. Different methods that achieve this requirement are described in Appendix I.

The interruption or denial of BCR negotiations is performed by discarding physically request or acknowledgement RM cells so that these cells do not cross a standardized interface. Examples where the above priority principles apply for ABT/DT are depicted in Figures A.1 and A.2.

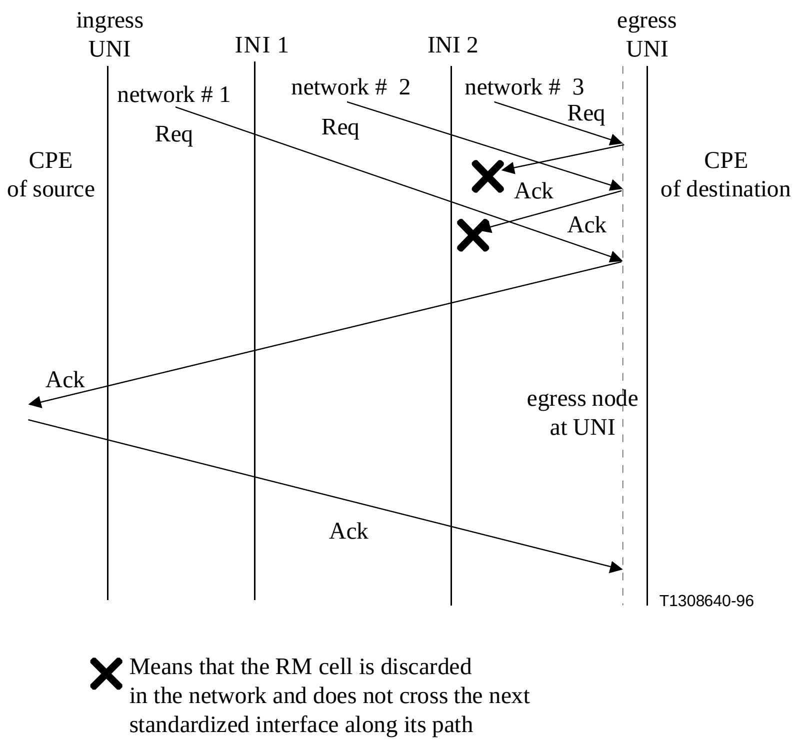

Figure A.1/I.371.1 – BCR negotiations denial in the originating network, or a downstream network (RM cells have different sequence numbers)

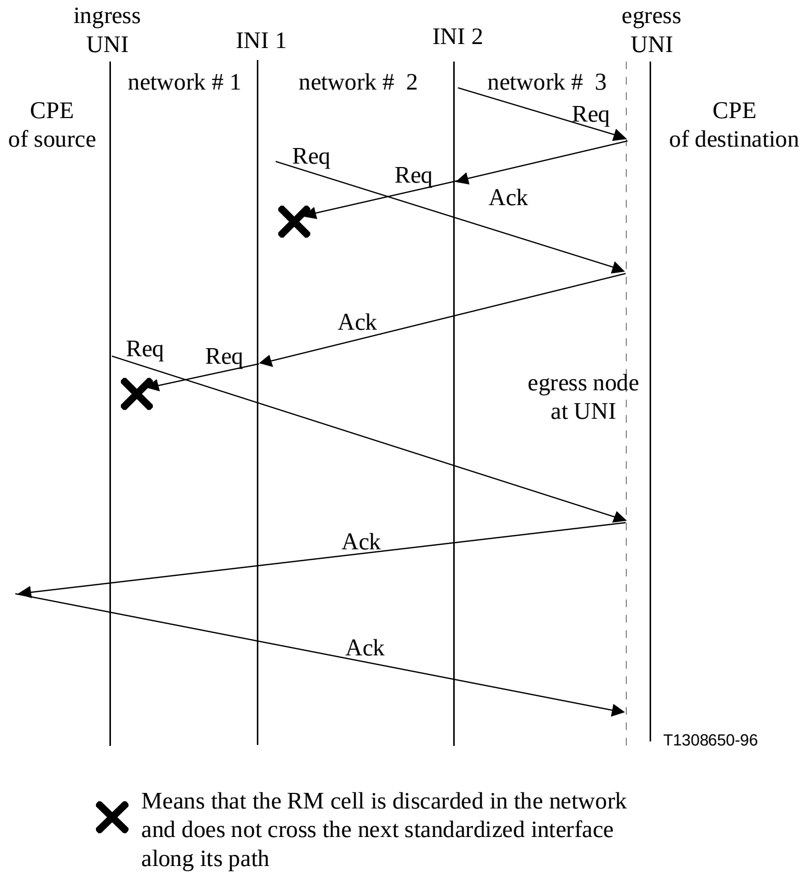

Figure A.2/I.371.1 – BCR negotiations denial in an upstream network (RM cells have different sequence numbers)

APPENDIX I

Examples of methods ensuring unicity of RM cell numbering in ABT

In order to implement the priority scheme (described in Annex A), that allows to discriminate between potentially conflicting BCR requests initiated by the network, it is necessary in some cases to rely on the Sequence Number (SN) value; this can be done only if cells corresponding to different BCR negotiations are allocated different SN values. However, Request RM cells that are generated by different networks may carry identical SN values unless a specific scheme is implemented. No scheme is currently recommended to ensure this property. Three possible methods are described in this appendix.

I.1 Segmentation of the SN field between different networks

It is possible to segment the coding of the 4-octet SN field between the networks along the connection. This would naturally prevent two different BCR negotiations being identified with the same SN, since a given network should not initiate a new BCR renegotiation while one initiated by itself is pending.

I.2 Proprietary handling of SN field

For instance, if a network is processing a BCR negotiation identified by a given sequence number and if this network receives a BCR request with higher priority but with the same SN value, the network may change the sequence number of this latter BCR transaction for processing in this network and the downstream networks but the network considered should also restore the initial sequence number value in the response to upstream networks. The actions taken when different RM cells have the same sequence number values are depicted in Figure I.1.

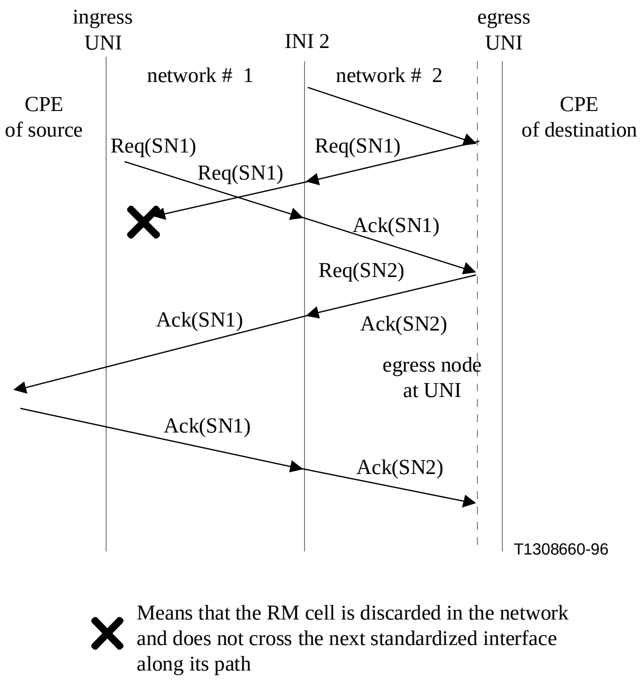

Figure I.1/I.371.1 – Avoidance of conflicts between sequence numbers

I.3 Segmentation of the SN field for indicating relative location of RM cell

The following scheme may be envisaged: out of the four available octets in the SN field of the ABT RM cell, three are used to assign a number (NA) to each RM cell that is generated, and one (RL) is used to identify the location of the network where it is observed at a given time, relative to the network that generated it.

- A cell is generated by the network, in the forward direction, as a Req cell, with a given NA and RL=0;
- NA is not modified when the RM cell crosses an interface;
- when the request crosses a standardized interface, RL is increased by 1;
- when the cell is turned back as an Ack at the destination UNI, RL is unchanged;
- when the Ack crosses a standardized interface, RL is decreased by 1 until RL=0;
- the Ack cell is then passed across the interface, as a Req, with RL=1;
- when the request crosses a standardized interface, RL is increased by 1.

The above scheme ensures that, in the network that has originated the BCR renegotiation, an Ack cell corresponding to a given Req cell carries exactly the same sequence number (NA,RL).

Moreover, two RM cells generated in different networks have necessarily different sequence numbers (different RL values).

In order to identify the priority level of a given cell, one stores the number (NA,RL) carried by the RM/Req cell at the given interface. Upon reception of either an Ack, or another Req, it is possible to identify the cell with the highest priority.

APPENDIX II

Derivation of conformance definition parameters for ABT

Consider an ABT connection, which is conforming to the peak cell rate 1/*T*, the sustainable cell rate Λ^0_{SCR}, and the maximum block size MBS^0 at the PHY-SAP of the equivalent terminal. These parameters are specified in the traffic contract. Moreover, the forward and backward user request RM cell flows at the interface considered are conforming to the GCRA(T_{RM}, τ_{RM}) and GCRA(T'_{RM}, τ'_{RM}), respectively. The number *S*(0,*t*) of cells, which can be transmitted at the PHY-SAP of the equivalent terminal over the time interval (0,*t*) satisfies:

$$S(0,t) = \sum_{\substack{\text{number of ATM}\\ \text{blocks in } (0,t)}} \rho_i \left(t_i^d - t_i^f\right) \le \Lambda^0_{SCR} \times t + MBS^0$$

where ρ_i is the BCR in cell/s of ATM block *i* and t_i^d and t_i^f are the starting and finishing times of ATM block *i*, respectively. t_i^d and t_i^f are in fact the transmission times of the leading and the trailing RM cells of ATM block *i*, respectively.

Define t''_{SCR} as in 5.4. To determine the worst case for the amount of resources consumed by the ABT connection, assume that the leading RM cells experience the minimum virtual cell transfer delay and that the trailing RM cells of ATM blocks experience the maximum virtual cell transfer delay. The size of ATM block *i* is then at most increased by $t''_{SCR} \times \rho_i$ cells.

It then follows that under the assumption that the ABT connection is conforming to the peak cell rate 1/*T*, the sustainable cell rate Λ^0_{SCR} and the maximum block size MBS^0 at the PHY-SAP of the equivalent terminal, the number S' of cells which can be transmitted at the interface satisfies, by noting that $\rho_i \le \frac{1}{T}$,

$$\begin{aligned} S'(0,t) &= \sum_{i=1}^{n(t)} \rho_i \left(t_i^d - t_i^f\right) + \sum_{i=1}^{n(t)} \rho_i \times \tau''_{SCR} \\ &\le t \times \Lambda^0_{SCR} + MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times n(t) \end{aligned} \qquad (1)$$

where *n(t)* is the number of ATM blocks over *(*0,*t)*.

The number *n(t)* of ATM blocks actually depends on the transmission mode and on the traffic contract for the user request RM cell flows.

In the immediate transmission mode (ABT/IT), taking into the traffic contract on the user request RM cell flow issued by the source on the forward direction, this number satisfies:

$$n(t) \le \frac{t}{T_{RM}} + \sigma_{RM} \qquad (2)$$

where $\sigma_{RM} = \left\lfloor 1 + \frac{\tau_{RM}}{T_{RM} - \Delta} \right\rfloor$ with Δ denoting the cell transmission time. It follows that:

$$S'(0,t) \le t\left(\Lambda^0_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \frac{1}{T_{RM}}\right) + MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times \sigma_{RM} \quad (3)$$

As a consequence, the cell stream at the interface is characterized by the sustainable Λ_{SCR} and the (fractional) maximum burst size MBS defined by:

$$\begin{aligned} \Lambda_{SCR} &= \min\left(\Lambda^0_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \frac{1}{T_{RM}}, \frac{1}{T}\right) \\ MBS &= MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times \left\lfloor 1 + \frac{\tau_{RM}}{T_{RM} - \Delta} \right\rfloor \end{aligned} \quad (4)$$

The tolerance τ_{SCR} is determined by using the relation:

$$\tau_{SCR} = (MBS - 1)(T_{SCR} - T) \quad (5)$$

In the delayed transmission mode (ABT/DT), since BCR negotiations can be initiated by both the source and the destination, not only the number of ATM blocks due to the source but also those due the destination should be taken into account. The user request RM cell flow of the destination should be conforming at the interface considered to the GCRA(T'_{RM} , τ'_{RM}) (the parameters T'_{RM} and τ'_{RM} are known at connection establishment). The aggregation of the user request RM cell flows generated by both the source and the destination may give rise at most to *n(t)* ATM blocks over the time interval *(*0,*t)* with:

$$n(t) \le t\left(\frac{1}{T_{RM}} + \frac{1}{T'_{RM}}\right) + \sigma''_{RM} \quad (6)$$

where $\sigma''_{RM} = \left\lfloor 2 + \frac{\tau_{RM}}{T_{RM} - \Delta} + \frac{\tau'_{RM}}{T'_{RM} - \Delta} \right\rfloor$.

It follows that the connection is characterized at the interface considered by the sustainable Λ_{SCR} and the (fractional) maximum burst size MBS defined as:

$$\begin{aligned} \Lambda_{SCR} &= \min\left(\Lambda^0_{SCR} + \frac{1}{T} \times \tau''_{SCR} \times \left(\frac{1}{T_{RM}} + \frac{1}{T'_{RM}}\right), \frac{1}{T}\right) \\ MBS &= MBS^0 + \frac{1}{T} \times \tau''_{SCR} \times \left\lfloor 2 + \frac{\tau_{RM}}{T_{RM} - \Delta} + \frac{\tau'_{RM}}{T'_{RM} - \Delta} \right\rfloor \end{aligned} \quad (7)$$

Tolerance τ_{SCR} relevant to the ATM block conformance definition for ABT/DT is deduced by using equation (5).

Remark 1 – The corrective terms in the above formulae giving the parameters to take into account in the block level conformance definition depend on the traffic characteristics of the RM cell flows. In general, T_{RM} is taken large enough (a fraction of the round-trip time through the network for ABT/IT and several times this round-trip time for ABT/DT). Moreover, τ_{RM} should be chosen sufficiently small so as to avoid clumps of RM cells. It thus turns out that the corrective terms are, in general, small when compared to the intrinsic parameters.

APPENDIX III

Source, destination, and network element reference behaviours for ABR

III.1 Source reference behaviour

In order to obtain full use of the dynamic bandwidth of an ABR connection, a source needs to send RM cells in the forward direction (i.e. a Forward RM cell) of the information flow. A source receives RM cells in the backward direction (i.e. a Backward RM cell) unless these cells have been lost in the network. For efficient operation of the close loop control, a source needs to adapt regularly to the changing network conditions. The source will ignore backward RM cells with incorrect CRC-10 in the EDC field (for EDC field, see 7.1/I.371).

User data cells are emitted with the CLP bit set to 0.

A source should normally send an in-rate forward RM cell for every (N_{RM}-1) other in-rate cells. The parameter N_{RM} could be network-specific or set to a default value.

In forward RM cells, the source should set the MCR field to MCR, and set the CCR field to be equal to the current ACR.

At the PHY-SAP of the equivalent terminal (see Recommendation I.371), an active source should emit in-rate cells at a rate no more than the current Allowed Cell Rate (ACR). The value of ACR shall never exceed PCR, nor shall it ever be less than MCR.

A source should update its ACR according to the information received in backward RM cells:

1) If the ECR value is less than ACR, then ACR should be reduced to ECR, but not less than MCR.

2) If the ECR value is greater than the ACR, then ACR may be increased (unless the backward RM cell is a BECN, in which case the ACR shall not be increased). The increase in ACR should be limited by a fixed increment RIF*PCR that provides for stepwise convergence to ECR. If the incremented ACR is greater than ECR, it is set to ECR. Setting RIF to 1 would allow an immediate jump to ECR. The Rate Increase Factor (RIF) would be set by default or assigned at connection establishment.

3) A source may make use of the CI and NI bits:

   a) If the source receives a RM cell with CI=1, then the value of ACR (in effect prior to the arrival of the backward RM cell) should be reduced by a multiplicative factor, but not further than MCR. Specifically, ACR should be reduced by at least ACR*RDF where the parameter RDF, the Rate Decrease Factor, can be set by default or assigned at connection establishment by management procedures or by signalling.

   b) If the backward RM cell has CI=0 and NI=0, then the ACR may be increased by at most the additive increment RIF*PCR to a rate not greater than the PCR.

   c) If the backward RM cell has NI=1 then the source should not increase the ACR.

   d) If the value of ACR resulting from steps 3) a) to 3) c) is greater than the value of ECR in the backward RM cell, then the ACR should be reduced to a value that is less than or equal to ECR, but not less than MCR. Otherwise, the source should use the ACR value calculated from the CI and NI bits alone.

4) In addition, if the source makes use of the QueueLength field and if the QueueLength is non-zero, then the sending rate should be further decreased, or no cells should be sent for a period of time, so as to allow the queue length to decrease. Procedures for calculating rate reductions and intervals based on non-zero QueueLengths are under study.

In addition to ACR updates due to reception of backward RM cells, a source should update its ACR according to the following rules:

5) When a source initializes, it should set the allowed cell rate, ACR, to at most the Initial Allowed Cell Rate (IACR) and the first in-rate cell sent should be a forward RM cell. The value of IACR is greater than or equal to MCR. At the beginning of the connection, the user is allowed to send at most the Transit Buffer Exposure (TBE) number of cells at the IACR without receiving a backward RM cell, which explicitly allocates an ACR. When TBE cells have been transmitted without any received RM cell, the source should reduce its rate, in steps or in a single step, to MCR. IACR can be negotiated between the network and the user at connection establishment. The value of TBE is assigned to the connection via management procedures or signalling.

6) A source which has not emitted any in-rate cells for a sufficiently long period should reduce its ACR to IACR if its ACR is above IACR to reflect the reallocation of network resources that may have taken place during the period of inactivity. When the source becomes active again, it should behave as in the previous item 5), using the (possibly reduced) allowed cell rate.

7) A source which has failed to receive a backward RM cell in a sufficiently long period of time should reduce its sending rate but need not reduce it below the MCR. The definitions of when a backward RM cell is considered overdue and the appropriate rate reduction are for further study.

III.2 Destination reference behaviour

A destination enables its corresponding source to estimate the bandwidth available from the network by returning RM cells to the source.

1) The destination should turn all RM cells received around to return them to the source. The direction bit, DIR, should be changed from "Forward" to "Backward".

2) If a destination cannot turn around a forward RM cell before it receives a subsequent forward RM cell to be turned around on the same VC, it may return only the most recent forward RM cell and discard the older forward RM cells. Alternatively, it may emit the older RM cell with the CLP bit set equal 1, and with the contents of the older cell possibly overwritten by the contents of the newer cell. However, loss of a backward CLP=1 RM cell between a standardized interface and the source can cause misalignment between the ACRs at the source and the conformance definition at the interface, which may impact the QOS of the connection. If a destination finds that it does not have an adequate ACR on the backward connection to support the emission of backward RM cells, it should consider itself to be in internal congestion and act as in item 4) below.

3) If an EFCI=1 has been received on the data cell prior to the RM cell, then the destination should mark the backward RM cell. An implementation may either:
   a) reduce the ECR; or
   b) set the CI bit in the RM cell.

4) To declare itself congested, the destination may do one or more of the following:
   a) further reduce ECR to whatever rate it can support;
   b) set the CI bit and/or NI bit;
   c) further increase the value of the QueueLength field in the RM cell.

A destination can also generate a backward RM cell without having received a forward RM cell. These cells are BECN cells. These cells have the following characteristics:

- BECN cells have the CLP bit set to 0.
- The BECN bit in the message field must be set.
- The direction should be "backward".
- Either the CI or the NI bit is set to 1.

Other interactions between forward bandwidth, backward bandwidth and the frequency with which RM cells are sent are for further study.

III.3 Network element reference behaviour

The network element may modify transiting RM cells based on the network element state. The need for network elements to insert forward RM cells is for further study.

A network element is not allowed to update the ABR RM cell fields protected by the EDC field if the CRC-10 code in the EDH field is wrong.

A network element shall implement at least one of the following methods to control congestion at queueing points:

1) The network element may reduce the ER field of forward and/or backward RM cell (Explicit Rate Marking).
2) The network element may set the EFCI flag in the data cell headers (EFCI marking).
3) The network element may set CI=1 or NI=1 in forward and/or backward RM cells (Relative Rate Marking).
4) The network element queueing point may set the queue length field of the RM cell to the maximum of the present value and the number of cells queued for this VC at this queueing point for this connection.

The explicit rate feedback provided by a network element is derived from the defined allocation policy.

In addition, the network element may segment the ABR control loop using a virtual source and destination (VS/VD control).

A network element may generate backward RM cells, called Backward Explicit Congestion Notification (BECN) cells. These cells have the following characteristics:

- BECN cells have the CLP bit set to 0.
- The BECN bit in the message field must be set.
- The direction should be "backward".
- The CI or the NI bit is set to 1.

The maximum rate of network-element-generated BECN cells is currently not specified but should be consistent with the mutual agreement limiting the aggregate rate of BECN cells that applies at standardized interfaces.

Backward RM cells may be serviced out of sequence with respect to data cells. Priority of forward RM cells is for further study. Bounds on the relative delay of RM cells with respect to data cells is under study.

In the special case when 1/ECR becomes large relative to the round-trip time, it is no longer reasonable to send at least one RM cell per round-trip time. This has the effect of increasing the feedback time beyond a round-trip time. As a consequence, one may require additional buffer allocations for those VCs. This is for further study.

ITU-T RECOMMENDATIONS SERIES

Series A	Organization of the work of the ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	**Integrated services digital network**
Series J	Transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communication
Series Z	Programming languages