

Superseded by a more recent version



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

I.356

(10/96)

SERIES I: INTEGRATED SERVICES DIGITAL
NETWORK

Overall network aspects and functions – Performance
objectives

B-ISDN ATM layer cell transfer performance

ITU-T Recommendation I.356
Superseded by a more recent version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Superseded by a more recent version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
SERVICE CAPABILITIES	I.200–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
INTERNETWORK INTERFACES	I.500–I.599
MAINTENANCE PRINCIPLES	I.600–I.699
B-ISDN EQUIPMENT ASPECTS	I.700–I.799

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

Superseded by a more recent version

ITU-T RECOMMENDATION I.356

B-ISDN ATM LAYER CELL TRANSFER PERFORMANCE

Summary

This Recommendation defines parameters for quantifying the ATM cell transfer performance of a broadband ISDN connection. It includes adjusted parameter definitions that might be used when cells do not conform with the negotiated traffic contract. This Recommendation includes provisional performance objectives for cell transfer, some of which depend on the user's selection of Quality of Service (QOS) class. This Recommendation includes the definitions for those QOS classes. Finally, each performance objective is allocated to the individual national portions involved in providing the international connection.

Annexes A, B and C provide information about ATM adaptation layer performance, information about factors contributing to cell transfer delay and cell delay variation, and information about performance measurement methods.

Source

ITU-T Recommendation I.356 was revised by ITU-T Study Group 13 (1993-1996) and was approved by the WTSC (Geneva, 9-18 October 1996).

Keywords

AAL performance, ATM, ATM performance, ATM transfer capabilities, B-ISDN, B-ISDN performance, cell block, Cell Delay Variation (CDV), Cell Error Ratio (CER), Cell Loss Ratio (CLR), Cell Misinsertion Rate (CMR), Cell Transfer Delay (CTD), cell transfer outcome, errored cell, hypothetical reference connection (HRX), international portion, lost cell, misinserted cell, national portion, Network Performance (NP), performance, performance allocation, performance measurement, performance monitoring, performance objectives, Quality of Service (QOS), quality of service class, quality of service negotiation, severely errored cell block, Severely Errored Cell Block Ratio (SECBR), successfully transferred cell, tagged cell, unbounded performance, unspecified performance.

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..... 3
4	Performance model..... 4
5	ATM cell transfer outcomes 7
5.1	Successful cell transfer outcome..... 8
5.2	Tagged cell transfer outcome..... 9
5.3	Errored cell outcome..... 9
5.4	Lost cell outcome..... 9
5.5	Misinserted cell outcome 9
5.6	Severely errored cell block outcome 9
6	ATM performance parameters..... 10
6.1	Cell error ratio..... 10
6.2	Cell loss ratio 10
6.2.1	CLR ₀ 11
6.2.2	CLR ₀₊₁ 11
6.2.3	CLR ₁ 11
6.3	Cell misinsertion rate..... 12
6.4	Severely errored cell block ratio..... 12
6.5	Cell transfer delay..... 12
6.5.1	Mean cell transfer delay 12
6.5.2	Cell delay variation..... 12
6.6	Cell flow related parameters..... 15
7	Network performance when some cells are non-conforming..... 15
7.1	A method for computing the number of non-conforming cells..... 15
7.2	Upper bounding the number of non-conforming cells..... 18
7.2.1	Bounding cell non-conformance for a single peak cell rate specification..... 18
7.2.2	Bounding cell non-conformance for independent peak cell rate specifications 19
7.2.3	Bounding cell non-conformance for dependent, coordinated cell rate specifications defined on the aggregate cell stream 19
7.2.4	Bounding cell non-conformance for dependent, coordinated cell rate specifications defined on the aggregate and on the CLP = 0 substream 19

Superseded by a more recent version

	Page
7.3	Adjusted CLR performance when there are non-conforming cells 20
7.3.1	CLR when the conformance specification applies to a single cell stream 20
7.3.2	CLR when the conformance specification applies to the aggregate and the CLP = 0 substream 21
7.4	Non-conforming cells and the severely errored cell block outcome 21
8	Network performance objectives 21
8.1	General discussion of per-connection QOS 22
8.2	QOS classes 22
8.2.1	Nature of the network performance objectives..... 22
8.2.2	Statistical estimation issues 23
8.2.3	Unbounded (unspecified) performance 23
8.2.4	Default values for cell error ratio, cell misinsertion rate, and severely errored cell block ratio..... 23
8.2.5	Association of QOS classes with ATM transfer capabilities 25
8.3	Alternative QOS negotiation procedures 25
9	Allocation of the performance objectives 26
9.1	General principles of allocation 26
9.2	Route length calculation 26
9.3	Allocation of the QOS class 1 CTD objective 27
9.4	Allocation of the QOS class 1 CDV objective 27
9.5	Allocation of the SECBR and CER objectives 28
9.6	Allocation of the QOS class 1 CLR objective 29
9.7	Allocation of the QOS class 2 and QOS class 3 CLR objectives 29
9.8	Allocation of the CMR objective 30
9.9	Concatenation of QOS values 30
Annex A	Relationship between ATM layer NP and the NP of AAL type 1 for CBR services 30
A.1	Possible AAL functions and their effects 30
A.1.1	Lost cell outcome and misinserted cell outcome 30
A.1.2	Errored cell outcome 31
A.1.3	Cell transfer delay 31
A.2	Bounding relationships between NP parameters and binary errors 31

Superseded by a more recent version

	Page
Annex B – Cell transfer delay, 1-point CDV and 2-point CDV characteristics	31
B.1 Components of delay associated with ATM-based user information transfer	31
B.2 Relationship between cell clumping and cell queues	32
B.3 1-point CDV and non-conformance	32
B.4 Relationship between cell transfer delay and cell loss in a single shared buffer	33
Annex C – Cell transfer performance measurement methods	34
C.1 General aspects of performance monitoring using OAM cells.....	34
C.2 Cell error ratio.....	36
C.3 Cell loss ratio	36
C.4 Cell misinsertion rate	37
C.5 Severely errored cell block ratio	37
C.6 Cell transfer delay	38
C.7 Cell delay variation.....	38
C.8 Estimation of CLR and SECBR in case of lost FM OAM cells	39
Appendix I – Assessing the performance of a UPC/NPC mechanism.....	42
Appendix II – Hypothetical reference connections for validating the ATM performance objectives	42
II.1 Number ATM nodes in the HRXs	43
II.2 Switching speeds in the HRXs.....	44
II.3 Loading within the HRXs	44
II.4 Geostationary satellites within the HRXs.....	44
II.5 Other aspects of the HRXs.....	45
Appendix III – Example applications of the allocation rules of 9.5, 9.6 and 9.7	45

Superseded by a more recent version

Recommendation I.356

B-ISDN ATM LAYER CELL TRANSFER PERFORMANCE

(revised in 1996)

1 Scope

This Recommendation defines speed, accuracy, and dependability performance parameters for cell transfer in the ATM layer of a broadband ISDN. The defined parameters apply to end-to-end ATM connections and to specified portions of such connections. The parameters are defined on the basis of ATM cell transfer reference events which may be observed at physical interfaces between ATM networks and associated customer equipment, and at physical interfaces between ATM networks.

NOTE 1 – The parameters defined in this Recommendation may be augmented or modified based upon further study of the requirements of the services to be supported on broadband ISDNs.

NOTE 2 – The parameters defined in clause 6 apply to cell streams in which all cells conform with the negotiated I.371 traffic contract. Clause 7 illustrates one way of extending the definitions and measurement methods to cell streams in which some cells do not conform with the traffic contract. It is recognized that further work is needed in this area.

NOTE 3 – The defined parameters are intended to characterize ATM connections in the available state. Availability decision parameters and associated availability parameters and their objectives are the subject of Recommendation I.357.

Clause 8 recommends ATM performance values to be achieved internationally for each of the defined parameters. Some of these values depend on which Quality of Service (QOS) class the end-users and network providers agree on for the connection. Clause 8 defines four different QOS classes. Clause 9 provides guidance on the allocated performance levels that each specified portion should provide in order to achieve the recommended end-to-end international performance.

The material in clauses 4 to 7 is equally applicable to international Virtual Channel Connections (VCCs) and international Virtual Path Connections (VPCs). The international objectives and allocations in clauses 8 and 9 are also applicable to both VCCs and VPCs. However, the end users (customers) of an international VPC will often be two networks that use the VPC to support individual VCCs. In order to meet the end-to-end objectives on each VCC, the performance of the supporting VPC must be better. The degree to which the VPC performance must be better is for further study.

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 listed below. A list of the currently valid ITU-T Recommendations is regularly published.

- [1] ITU-T Recommendation G.114 (1996), *One-way transmission time*.
- [2] ITU-T Recommendation G.826 (1996), *Error performance parameters and objectives for international, constant bit rate digital paths at or above the primary rate*.

Superseded by a more recent version

- [3] ITU-T Recommendation I.113 (1993), *Vocabulary of terms for broadband aspects of ISDN.*
- [4] ITU-T Recommendation I.150 (1995), *B-ISDN asynchronous transfer mode functional characteristics.*
- [5] ITU-T Recommendation I.311 (1996), *B-ISDN general network aspects.*
- [6] CCITT Recommendation I.321 (1991), *B-ISDN protocol reference model and its application.*
- [7] ITU-T Recommendation I.350 (1993), *General aspects of quality of service and network performance in digital networks, including ISDNs.*
- [8] ITU-T Recommendation I.351 (1993), *Relationships among ISDN performance Recommendations.*
- [9] ITU-T Recommendation I.353 (1996), *Reference events for defining ISDN and B-ISDN performance parameters.*
- [10] ITU-T Recommendation I.357 (1996), *B-ISDN semi-permanent connection availability.*
- [11] ITU-T Recommendation I.363.1 (1996), *B-ISDN ATM adaptation layer specification: Type 1 (AAL).*
- [12] ITU-T Recommendation I.371 (1996), *Traffic control and congestion control in B-ISDN.*
- [13] ITU-T Recommendation I.413 (1993), *ISDN user-network interface.*
- [14] ITU-T Recommendation I.610 (1995), *B-ISDN operation and maintenance principles and functions.*
- [15] ITU-T Recommendation I.361 (1995), *B-ISDN ATM layer specification.*
- [16] ITU-T Recommendation O.191 (1997), *Equipment to assess ATM layer cell transfer performance.*
- [17] ITU-T Recommendation Q.2761 (1995), *Function description of the B-ISDN User Part (B-ISUP) of Signalling System No. 7.*
- [18] 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.*
- [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 (DSS 2) – User-Network Interface (UNI) layer 3 specification for basic call/connection control.*
- [21] ITU-T Recommendation Q.2961 (1995), *Broadband Integrated Services Digital Network (B-ISDN) – Digital subscriber Signalling System No. 2 (DSS 2) – Additional traffic parameters.*
- [22] ITU-T Recommendation Q.2962 (1996), *Digital subscriber Signalling System No. 2 (DSS 2) – Connection characteristics negotiation during call/connection establishment phase.*
- [23] ITU-T Recommendation Q.2963.1 (1996), *Digital subscriber Signalling System No. 2 (DSS 2) – Connection modification: Peak cell rate modification by the connection owner.*

Superseded by a more recent version

3 Abbreviations

This Recommendation uses the following abbreviations:

AAL	ATM Adaptation Layer
ABR	Available Bit Rate ATC
ABT	ATM Block Transfer ATC
ABT/DT	ABT Delayed Transmission
ABT/IT	ABT Immediate Transmission
ATC	ATM Transfer Capability
ATM	Asynchronous Transfer Mode
B-ISDN	Broadband ISDN
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CEQ	Customer equipment/customer network
CER	Cell Error Ratio
CLP	Cell Loss Priority bit
CLR	Cell Loss Ratio
CMR	Cell Misinsertion Rate
CRE	Cell Reference Event
CTD	Cell Transfer Delay
DBR	Deterministic Bit Rate ATC
FM	Forward Monitoring
GCRA	Generic Cell Rate Algorithm
HEC	Header Error Control
HRX	Hypothetical Reference connection
IIP	International Interoperator Portion
INI	Inter-Network Interface
ISDN	Integrated Services Digital Network
ITP	International Transit Portion
MCSN	Monitoring Cell Sequence Number
MP	Measurement Point
MPI	International Measurement Point
MPT	Measurement Point at T _B
NP	Network Performance
NPC	Network Parameter Control
OAM	Operations and Maintenance
PCR	Peak Cell Rate

Superseded by a more recent version

PDH	Plesiochronous Digital Hierarchy
PL	Physical Layer
PM	Performance Monitoring
QOS	Quality of Service
SBR	Statistical Bit Rate ATC
SDH	Synchronous Digital Hierarchy
SECB	Severely Errored Cell Block
SECBR	Severely Errored Cell Block Ratio
SN	Sequence Number
SSN	Switching/Signalling Node
T	Nominal cell interarrival time
T _{max}	Timer for declaring a cell lost
TUC	Total User Cell
U	Unspecified/Unbounded
UNI	User-Network Interface
UPC	Usage Parameter Control
VC	Virtual Channel
VCC	Virtual Channel Connection
VP	Virtual Path
VPC	Virtual Path Connection

4 Performance model

ATM cell transfer performance is measured by observing the reference events created as ATM cells cross Measurement Points (MPs). Recommendation I.353 defines the MPs and the associated reference events that provide the basis for ISDN performance description.

According to Recommendation I.353, the only cells that create reference events are those with a payload type field indicating a user information cell. Thus, the performance recommendations of I.356 only apply to the transfer of user information cells.

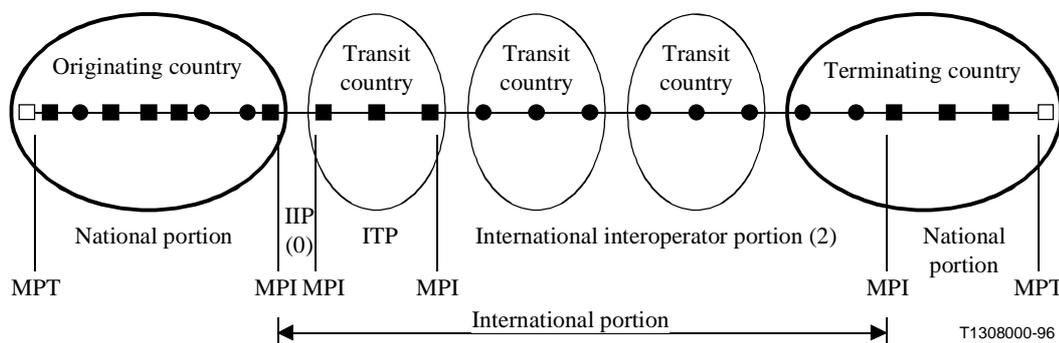
NOTE 1 – Network providers should endeavour to deliver the same QOS to other end-to-end user cells such as end-to-end OAM cells. It is expected that networks will transport these cells similarly to the way they transport user information cells, so it is appropriate to assume that their transfer performance will be similar. Performance recommendations for cell types other than user information cells are for further study.

For broadband ISDN, the MPs are located at interfaces where the ATM layer is accessible. The exact location within the protocol stack depends on whether the connection is a Virtual Channel (VC) or a Virtual Path (VP).

- For VCs: MPs are in the protocol stack above the VC multiplexing and demultiplexing functions, but below any other VC functions such as cell rate policing.
- For VPs: MPs are in the protocol stack above the VP multiplexing and demultiplexing functions, but below any other VP functions such as cell rate policing.

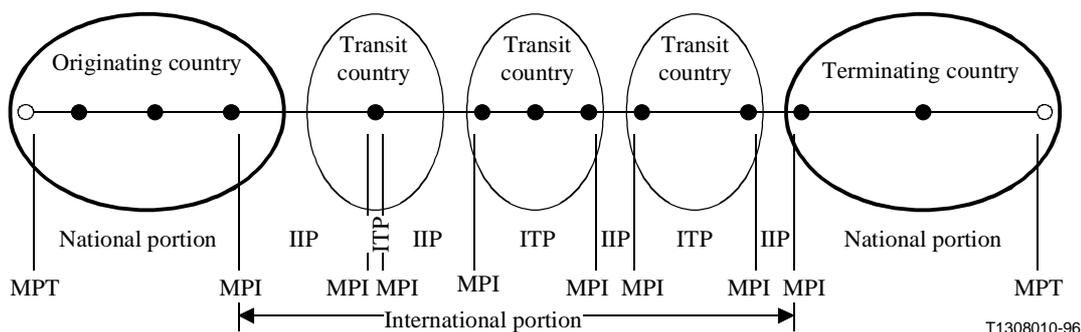
Superseded by a more recent version

As stated in Recommendation I.353, there are two types of MPs. MPTs are measurement points at or near T_B reference points and thus are at or near customer equipment/customer networks (CEQ). MPIs are measurement points established at the international switching/signalling nodes (SSN) before and after the connection crosses the national border.



- Customer equipment/customer network (CEQ) – VC source/sink
- VC Switching/Signalling Node (SSN)
- VP Switching/Signalling Node (SSN)

Figure 1/I.356 – Example VCC and its connection portions



- Customer equipment/customer network (CEQ) – VP source/sink
- VP Switching/Signalling Node (SSN)

Figure 2/I.356 – Example VPC and its connection portions

For the purpose of performance management, ATM connections are consequently divided into three types of connection portions:

- National portions: connection portions between the MPT and the MPI both within the originating (or terminating) country.
- International Transit Portions (ITPs): connection portions between two MPIs in a single transit country. For VCs there are VC switching or cross-connect elements between the two MPIs. For VPs, there are VP switching or cross-connect elements between the two MPIs.
- International Interoperator Portions (IIPs): connection portions between two MPIs in different countries. For VPs there are no ATM switching or cross-connect elements between these two MPIs. For VCs there may be VP switching or cross-connect elements, but there

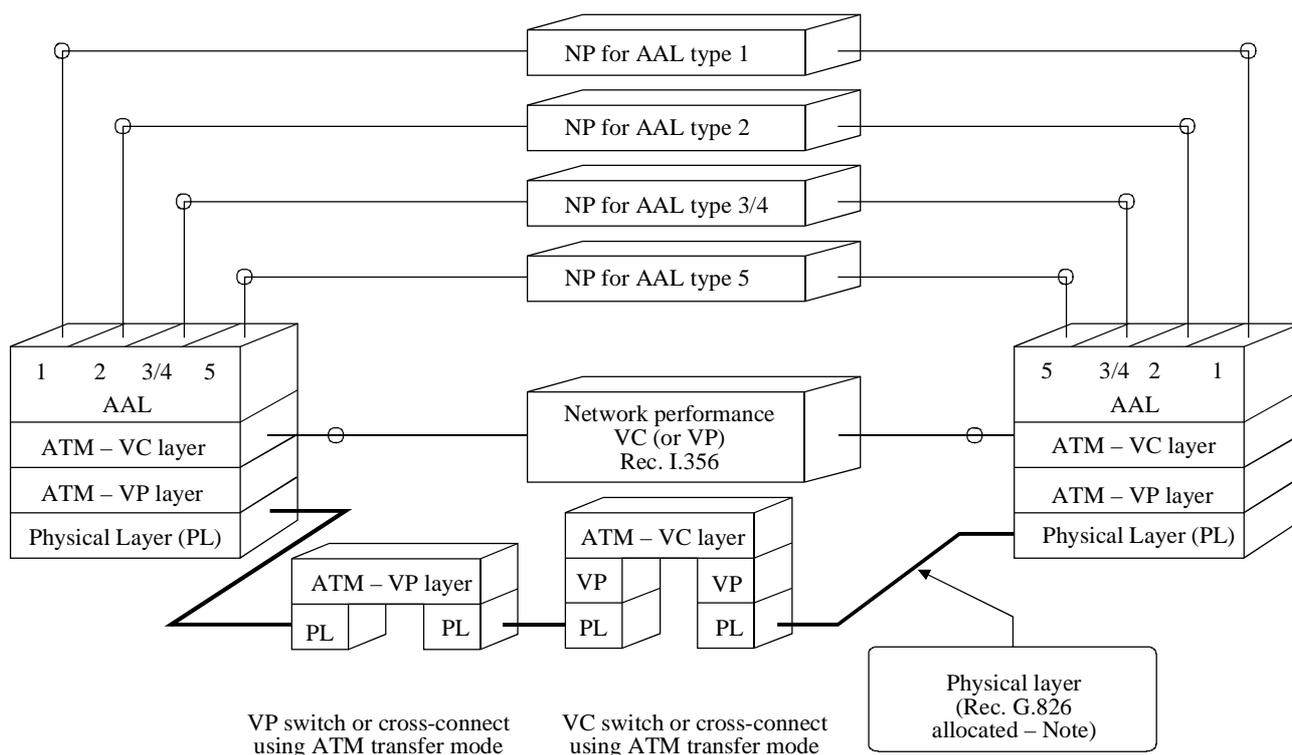
Superseded by a more recent version

are no VC switching or cross-connect elements between the MPIs. The abbreviation IIP(x) (x = 0,1,2,...) is used to indicate a VC IIP with "x" intervening transit countries, each providing VP switching or cross-connect functions.

The complete set of ITP and IIP is the international portion of the connection. Figures 1 and 2 illustrate these concepts.

NOTE 2 – Using the terminology of Recommendation I.353, private ATM networks are considered to be CEQ. Private networks may connect the end users to this public network model at one or both MPTs. The quantitative impact of CEQ on the end-to-end quality of service is an issue for further study and is not currently addressed in this Recommendation.

More information about measurement points and cell transfer reference events can be found in Recommendation I.353.



T1308020-96

NOTE – The need for additional physical layer performance parameters and objectives is under study.

Figure 3/I.356 – Layered model of performance for B-ISDN

Figure 3 illustrates the layered nature of B-ISDN performance issues. The Network Performance (NP) provided to B-ISDN users depends on the performance of three layers:

- The physical layer, which may be based on Plesiochronous Digital Hierarchy (PDH), Synchronous Digital Hierarchy (SDH), or cell-based transmission systems. This layer is terminated at points where the connection is switched or cross-connected by equipment using the ATM technique, and thus the physical layer has no end-to-end significance when such switching occurs.
- The ATM layer, which is cell-based. The ATM layer is physical media and application independent and has end-to-end significance. Recommendation I.356 specifies network performance at the ATM layer.

Superseded by a more recent version

- The ATM Adaptation Layer (AAL), which may enhance the performance provided by the ATM layer to meet the needs of higher layers. The AAL supports multiple protocol types, each providing different functions and different performance.

Qualitative relationships between ATM layer Network Performance (NP) and the NP provided by the type 1 AAL are described in Annex A. It is intended that quantitative relationships between ATM layer network performance and the performance of the physical layer and AALs will be developed.

5 ATM cell transfer outcomes

In the following, it is assumed that the sequence of ATM cells on a virtual channel or virtual path is preserved (see Recommendation I.150). Two cell reference events are said to be corresponding if they are created by the "same" cell at a predefined pair of boundaries. The practical determination as to whether two cell reference events were caused by the "same" cell is usually *ad hoc* and will rely on some combination of VP/VC identification, cell sequencing, and cell content.

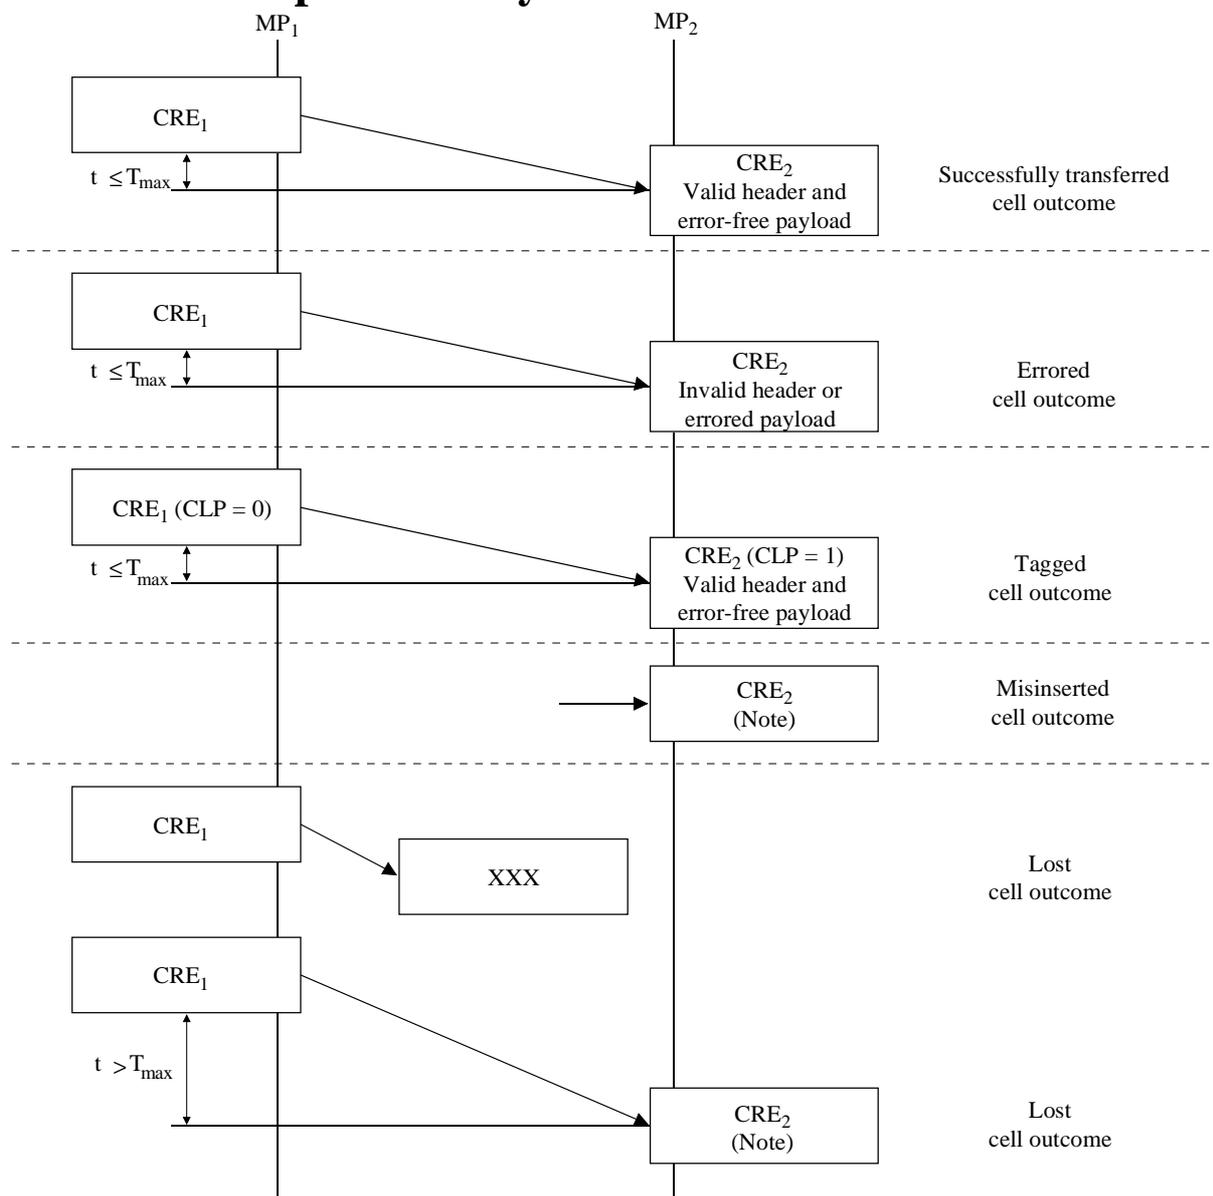
By considering two corresponding transfer reference events¹, CRE₁ and CRE₂ at MP₁ and MP₂ respectively, a number of possible cell transfer outcomes may be defined. A transmitted cell is either successfully transferred, errored, tagged, or lost. A received cell for which no corresponding transmitted cell exists is said to be misinserted. Figure 4 illustrates the cell transfer outcome definitions.

¹ Recommendation I.353 establishes that cell reference events only occur when:

- a) the cell's VPI/VCI field corresponds to the VPI/VCI of the monitored connection (after HEC processing); and
- b) the payload type field indicates a user information cell.

No other cells create reference events.

Superseded by a more recent version



NOTE – Outcome occurs independent of cell content.

T1308030-96

Figure 4/I.356 – Cell transfer outcomes

5.1 Successful cell transfer outcome

A successful cell transfer outcome occurs when a CRE_2 corresponding to CRE_1 happens within a specified time T_{max} of CRE_1 , and:

- 1) the binary content of the received cell information field conforms exactly with that of the corresponding transmitted cell; and
- 2) the cell is received with a valid header field.

Superseded by a more recent version

5.2 Tagged cell transfer outcome

A tagged cell outcome occurs when a CRE₂ corresponding to CRE₁ happens within a specified T_{max} of CRE₁, and:

- 1) the binary content of the received cell information field conforms exactly with that of the corresponding transmitted cell; and
- 2) the cell is received with a valid header field; and
- 3) the Cell Loss Priority (CLP) bit is changed from CLP = 0 at MP₁ to CLP = 1 at MP₂.

NOTE – A cell that was both tagged by the network and has errors in its information field creates an errored cell outcome but does not create a tagged cell outcome.

5.3 Errored cell outcome

An errored cell outcome occurs when a CRE₂ corresponding to CRE₁ happens within a specified time T_{max} of CRE₁, but:

- 1) the binary content of the received cell information field differs from that of the corresponding transmitted cell (i.e. one or more bit errors exist in the received cell information field); or
- 2) the cell is received with an invalid header field after Header Error Control (HEC) procedures are completed.

NOTE 1 – Most cells with header errors that are undetected or miscorrected by the HEC will be misdirected by the ATM layer procedures with the result that no CRE₂ occurs. These cell transfer attempts will be classified as lost cell outcomes.

NOTE 2 – An example of an invalid header field is a change in the CLP bit from CLP = 1 at MP₁ to CLP = 0 at MP₂.

NOTE 3 – A cell that was both tagged by the network and has errors in its information field creates an errored cell outcome.

5.4 Lost cell outcome

A lost cell outcome occurs when a CRE₂ fails to happen within time T_{max} of the corresponding CRE₁.

NOTE – Cell losses attributable to customer equipment shall be excluded in assessing the performance of the network. Estimation of cell losses occurring in customer equipment due to network causes is for further study.

5.5 Misinserted cell outcome

A misinserted cell outcome occurs when a CRE₂ happens without a corresponding CRE₁.

5.6 Severely errored cell block outcome

A cell block is a sequence of N cells transmitted consecutively on a given connection. A severely errored cell block outcome occurs when more than M errored cell, lost cell, or misinserted cell outcomes are observed in a received cell block.

When there is a cell loss ratio commitment for the aggregate cell flow, then all cell loss outcomes are considered in the determination of SECBs. When there are no performance commitments concerning the cell loss ratio for the aggregate cell flow, CLP = 0 + 1, or the CLP = 1 cell flow (as in the bi-level class defined in clause 8), then lost CLP = 1 cells are not considered in the determination of SECBs. In such cases, transmitted CLP = 1 cells are counted in determining the cell blocks, but lost CLP = 1 cells do not count against the SECB threshold, M Errored and misinserted CLP = 1 cells do count against the threshold. (See Annex C for more information.)

Superseded by a more recent version

The value of N is uniquely determined by the Peak Cell Rate (PCR) of the aggregate cell flow, $CLP = 0 + 1$. N is constructed so that there are between 12.5 and 25 cell blocks transmitted per second whenever the connection is operating at its aggregate PCR. Cell block sizes smaller than 128 cells are for further study. The value of M is fixed to be 1/32 of N.

$$N = \frac{PCR}{25}, \text{ where } N \text{ is then rounded to the next larger power of } 2,$$

$$M = \frac{N}{32}$$

Table 1/I.356 – Computation of cell block sizes and SECB threshold

PCR (cells/second)	(User information rate in Mbit/s)	N (block size)	M (threshold)
$0 < x \leq 3\,200$	$(0 < y \leq 1.23)$	128	4
$3\,200 < x \leq 6\,400$	$(1.23 < y \leq 2.46)$	256	8
$6\,400 < x \leq 12\,800$	$(2.46 < y \leq 4.92)$	512	16
$12\,800 < x \leq 25\,600$	$(4.92 < y \leq 9.83)$	1\,024	32
$25\,600 < x \leq 51\,200$	$(9.83 < y \leq 19.66)$	2\,048	64
$51\,200 < x \leq 102\,400$	$(19.66 < y \leq 39.32)$	4\,096	128
$102\,400 < x \leq 202\,800$	$(39.32 < y \leq 78.64)$	8\,192	256
$202\,800 < x \leq 409\,600$	$(78.64 < y \leq 157.29)$	16\,384	512
$409\,600 < x \leq 819\,200$	$(157.29 < y \leq 314.57)$	32\,768	1\,024
<p>NOTE 1 – The equation for N is valid for peak cell rates up to 819 200 cells per second. The values of N and M for PCR > 819 200 are for further study.</p> <p>NOTE 2 – For practical measurement purposes, a cell block can be approximated by an OAM cell block. The size of OAM cell blocks may vary from block to block, but if the SECB ratio (see 6.4) is to be approximated, the OAM block sizes should average to the specific value of N appropriate for the aggregate PCR.</p>			

6 ATM performance parameters

This clause defines a set of ATM cell transfer performance parameters using the cell transfer outcomes defined in clause 5. All parameters may be estimated on the basis of observations at the MPs. Cell transfer performance measurement methods are described in Annex C.

6.1 Cell error ratio

Cell Error Ratio (CER) is the ratio of total errored cells to the total of successfully transferred cells, plus tagged cells, plus errored cells in a population of interest. Successfully transferred cells, tagged cells, and errored cells contained in severely errored cell blocks are excluded from the calculation of cell error ratio.

6.2 Cell loss ratio

Cell Loss Ratio (CLR) is the ratio of total lost cells to total transmitted cells in a population of interest. Lost cells and transmitted cells in severely errored cell blocks are excluded from the calculation of cell loss ratio. Three special cases are of interest, CLR_0 , CLR_{0+1} , and CLR_1 .

Superseded by a more recent version

The definitions in 6.2, 6.2.1, 6.2.2 and 6.2.3 are comprehensive, including cell losses occurring in UPC/NPC mechanisms. Thus, they include cell losses, if any, due to non-conforming traffic and cell losses, if any, due to misbehaving UPC/NPCs. Defined in this way, the parameters are representative of the quality of service observed and they are suitable for evaluating network performance when all cells conform to the traffic contract.

NOTE 1 – Clause 7 provides adjusted definitions of CLR that may be used to evaluate network performance when some cells are not conforming with the traffic contract.

NOTE 2 – In order to meet its CLR commitments, a network provider will need to evaluate the performance of its UPC/NPC mechanism. Appendix I provides information about assessing UPC/NPC mechanisms.

6.2.1 CLR_0

Let $N_i(0)$ represent the number of $CLP = 0$ cells transmitted and let $N_l(0)$ represent the number of corresponding lost cell outcomes plus the number of corresponding tagged cell outcomes. The cell loss ratio for high priority cells (CLR_0) is the ratio of $N_l(0)$ to $N_i(0)$.

NOTE – With this definition, cells that are tagged by the network (possibly due to overpolicing) are considered lost from the stream of high priority cells.

6.2.2 CLR_{0+1}

Let $N_i(0 + 1)$ represent the total number of cells transmitted and let $N_l(0 + 1)$ represent the number of corresponding lost cell outcomes. The cell loss ratio for the aggregate cell stream (CLR_{0+1}) is the ratio of $N_l(0 + 1)$ to $N_i(0 + 1)$.

NOTE 1 – Tagged cells are not considered lost from the aggregate stream.

NOTE 2 – When all cells are $CLP = 1$, CLR_{0+1} equals CLR_1 .

6.2.3 CLR_1

Let $N_i(1)$ represent the number of $CLP = 1$ cells transmitted and let $N_l(1)$ represent the number of corresponding lost cell outcomes. The cell loss ratio for low priority cells (CLR_1) is the ratio of $N_l(1)$ to $N_i(1)$.

NOTE 1 – With this definition, cells that are tagged by the network (but are still conforming to the aggregate traffic contract) are not considered in either the numerator or the denominator of the expression for CLR_1 .

NOTE 2 – As defined, CLR_1 quantifies the user's perception of the cell loss ratio for their low priority traffic.

Superseded by a more recent version

6.3 Cell misinsertion rate

Cell Misinsertion Rate (CMR) is the total number of misinserted cells observed during a specified time interval divided by the time interval duration² (equivalently, the number of misinserted cells per connection second). Misinserted cells and time intervals associated with severely errored cell blocks are excluded from the calculation of cell misinsertion rate.

6.4 Severely errored cell block ratio

Severely Errored Cell Block Ratio (SECBR) is the ratio of total severely errored cell blocks to total cell blocks in a population of interest.

NOTE – The severely errored cell block outcome and parameter provide a means of quantifying bursts of cell transfer failures and preventing those bursts from influencing the observed values for cell error ratio, cell loss ratio, cell misinsertion rate, and the associated availability parameters.

6.5 Cell transfer delay

The definitions in 6.5, 6.5.1, and 6.5.2 can only be applied to successfully transferred, errored, and tagged cell outcomes.

Cell Transfer Delay (CTD) is the time, $t_2 - t_1$, between the occurrence of two corresponding cell transfer events, CRE_1 at time t_1 and CRE_2 at time t_2 , where $t_2 > t_1$ and $t_2 - t_1 \leq T_{max}$. The value of T_{max} is for further study, but should be larger than the largest practically conceivable cell transfer delay.

6.5.1 Mean cell transfer delay

Mean cell transfer delay is the arithmetic average of a specified number of cell transfer delays.

6.5.2 Cell delay variation

Two cell transfer performance parameters associated with Cell Delay Variation (CDV) are defined. The first parameter, 1-point cell delay variation, is defined based on the observation of a sequence of consecutive cell arrivals at a single MP. The second parameter, 2-point cell delay variation, is defined based on the observations of corresponding cell arrivals at two MPs that delimit a virtual connection portion. The 1-point CDV parameter describes variability in the pattern of cell arrival (entry or exit) events at an MP with reference to the negotiated peak cell rate $1/T$ (see Recommendation I.371); it includes variability present at the cell source (customer equipment) and the cumulative effects of variability introduced (or removed) in all connection portions between the cell source and the specified MP. It can be related to cell conformance at the MP, and to network queues. It can also be related to the buffering procedures that might be used in AAL 1 to compensate for cell delay variation. The 2-point CDV parameter describes variability in the pattern of cell arrival events at the output of a connection portion (e.g. measurement point MP_2) with reference to the pattern of corresponding events at the input to the portion (e.g. measurement point MP_1); it includes only the delay variability introduced within the connection portion. It provides a direct measure of portion performance and an indication of the maximum (aggregate) length of cell queues that may exist within the portion. Additional information on relationships of these CDV-related parameters to cell queues and their application in ATM network performance specification is provided in Annex B.

² By definition, a misinserted cell is a received cell that has no corresponding transmitted cell on the considered connection. Cell misinsertion on a particular connection is caused by impairments either on physical layer unassigned cells or on cells being transmitted on a different connection. Since the mechanisms that cause misinserted cells have nothing to do with the number of cells transmitted on the observed connection, this performance parameter cannot be expressed as a ratio, only as a rate.

Superseded by a more recent version

6.5.2.1 1-point CDV at an MP

The 1-point CDV (y_k) for cell k at an MP is the difference between the cell's reference arrival time (c_k) and actual arrival time (a_k) at the MP [see Figure 5 a)]: $y_k = c_k - a_k$. The reference arrival time pattern (c_k) is defined as follows: 0

$$c_0 = a_0 = 0,$$
$$c_{k+1} = \begin{cases} c_k + T & \text{when } c_k \geq a_k ; \\ a_k + T & \text{otherwise.} \end{cases}$$

Positive values of 1-point CDV ("early" cell arrivals) correspond to cell clumping; negative values of 1-point CDV ("late" cell arrivals) correspond to gaps in the cell stream. The reference pattern defined above eliminates the effect of gaps in the specification and measurement of cell clumping.³

6.5.2.2 Cell delay variation between two MPs (2-point CDV)

The 2-point CDV (v_k) for cell k between MP_1 and MP_2 is the difference between the absolute cell transfer delay (x_k) of cell k between the two MPs and a defined reference cell transfer delay ($d_{1,2}$) between those MPs [see Figure 5 b)]: $v_k = x_k - d_{1,2}$.

The absolute cell transfer delay (x_k) of cell k between MP_1 and MP_2 is the difference between the cell's actual arrival time at MP_2 (a_{2k}) and the cell's actual arrival time at MP_1 (a_{1k}): $x_k = a_{2k} - a_{1k}$ ⁴. The reference cell transfer delay ($d_{1,2}$) between MP_1 and MP_2 is the absolute cell transfer delay experienced by cell 0 between the two MPs.

NOTE – The specification of cell 0 is for further study.

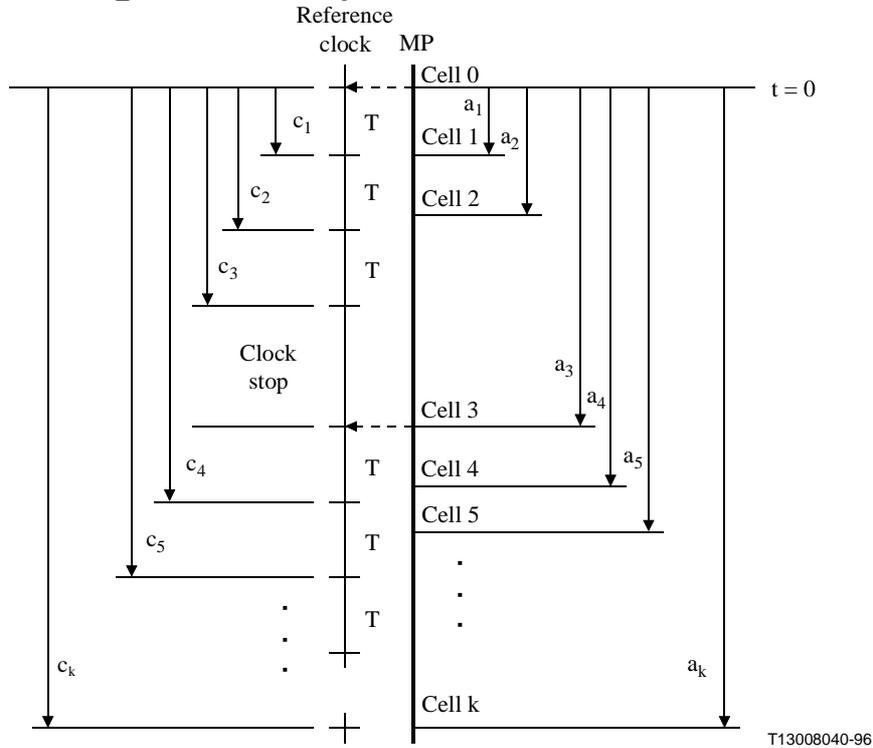
Positive values of 2-point CDV correspond to cell transfer delays greater than that experienced by the reference cell; negative values of 2-point CDV correspond to cell transfer delays less than that experienced by the reference cell. The distribution of 2-point CDVs is identical to the distribution of absolute cell transfer delays displaced by a constant value equal to $d_{1,2}$.

Annex C illustrates one method of estimating the range of the 2-point CDV distribution based on observations of 1-point CDV values (y_k) for connections providing CBR services. Annex B relates the probability distribution for 2-point cell delay to the cell loss ratio.

³ The reference clock "skips" by an amount equal to the difference between the actual and expected arrival times immediately after each "late" cell arrival.

⁴ Variables a_{2k} and a_{1k} are measured with reference to the same reference clock.

Superseded by a more recent version

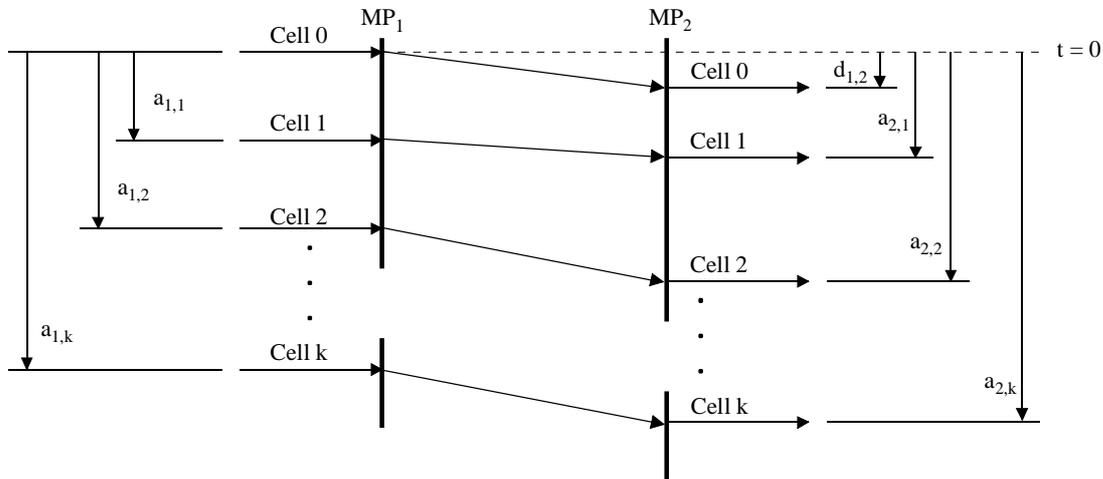


Variables:

- a_k Cell k actual arrival time at MP
- c_k Cell k reference arrival time at MP
- y_k 1-point CDV

$$y_k = c_k - a_k$$

a) Cell delay variation – 1-point definition

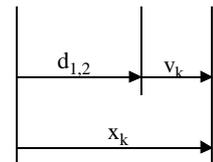


Variables:

- $a_{1,k}$ Cell k actual arrival at MP_1
- $a_{2,k}$ Cell k actual arrival time at MP_2
- $d_{1,2}$ Absolute cell 0 transfer delay between MP_1 and MP_2
- x_k Absolute cell k transfer time between MP_1 and MP_2
- v_k 2-point CDV value between MP_1 and MP_2

$$x_k = a_{2,k} - a_{1,k}$$

$$v_k = x_k - d_{1,2}$$



b) Cell delay variation – 2-point definition

T1308050-96

Figure 5/I.356 – Cell delay variation parameter definitions

Superseded by a more recent version

6.6 Cell flow related parameters

Network performance parameters quantifying an ATM connection's total capacity for transporting cells are for further study. New parameters may be needed for the specific ATM transfer capabilities defined in Recommendation I.371. For the ABT/DT transfer capability, it may be appropriate to quantify the number of times a request for a new block cell rate is denied by a network. For the ABT/IT transfer capability, it may be appropriate to quantify the number of times block transfers fail. For the ABR transfer capability, it may be appropriate to quantify the network's support or use of flow control mechanisms.

7 Network performance when some cells are non-conforming

This clause addresses the issue of defining network performance parameters applicable when some cells do not conform with the negotiated traffic contract.

It is assumed that the user has negotiated a traffic contract as described in Recommendation I.371. Such a contract specifies one or several traffic parameters and the Quality of Service (QoS) requirements. If non-conforming cells are observed on the connection, the network is allowed to discard a number of cells equal to the number of non-conforming cells obtained by an ideal UPC/NPC mechanism that implements the I.371 cell conformance definition. Such discarded cells are not counted as lost cells in assessing the network's CLR performance.

It is a network privilege to define its criteria for non-compliant connections, possibly based on the number of non-conforming cells observed. When a connection is judged to be non-compliant, no network performance commitments need be honoured. However, if the connection includes non-conforming cells, but it is not judged non-compliant, the network may choose to offer modified network performance commitments. This clause adjusts the network performance parameter definitions of Clause 6 to compensate for the occurrence of non-conforming cells and to provide a method that might be used in evaluating the modified network commitments.

7.1 A method for computing the number of non-conforming cells

It is assumed here that the user has negotiated a single cell rate that applies to its entire cell flow. Let T and τ respectively denote the negotiated emission interval and associated CDV tolerance.

Define (y'_k) and an associated theoretical time (c'_k) as follows:

$$\begin{aligned} c'_0 &= a_0 \\ y'_k &= c'_k - a_k \\ c'_{k+1} &= \begin{cases} c'_k & \text{when } c'_k > a_k + \tau \\ a_k + T & \text{when } c'_k \leq a_k \\ c'_k + T & \text{otherwise} \end{cases} \end{aligned}$$

These equations are a modification of the 1-point CDV equations presented in 6.5.2.1. The modified CDV parameter y'_k differs from 1-point CDV, y_k , only if for some cell $j < k$, y'_j is larger than τ (equivalently, if some cell j is non-conforming).

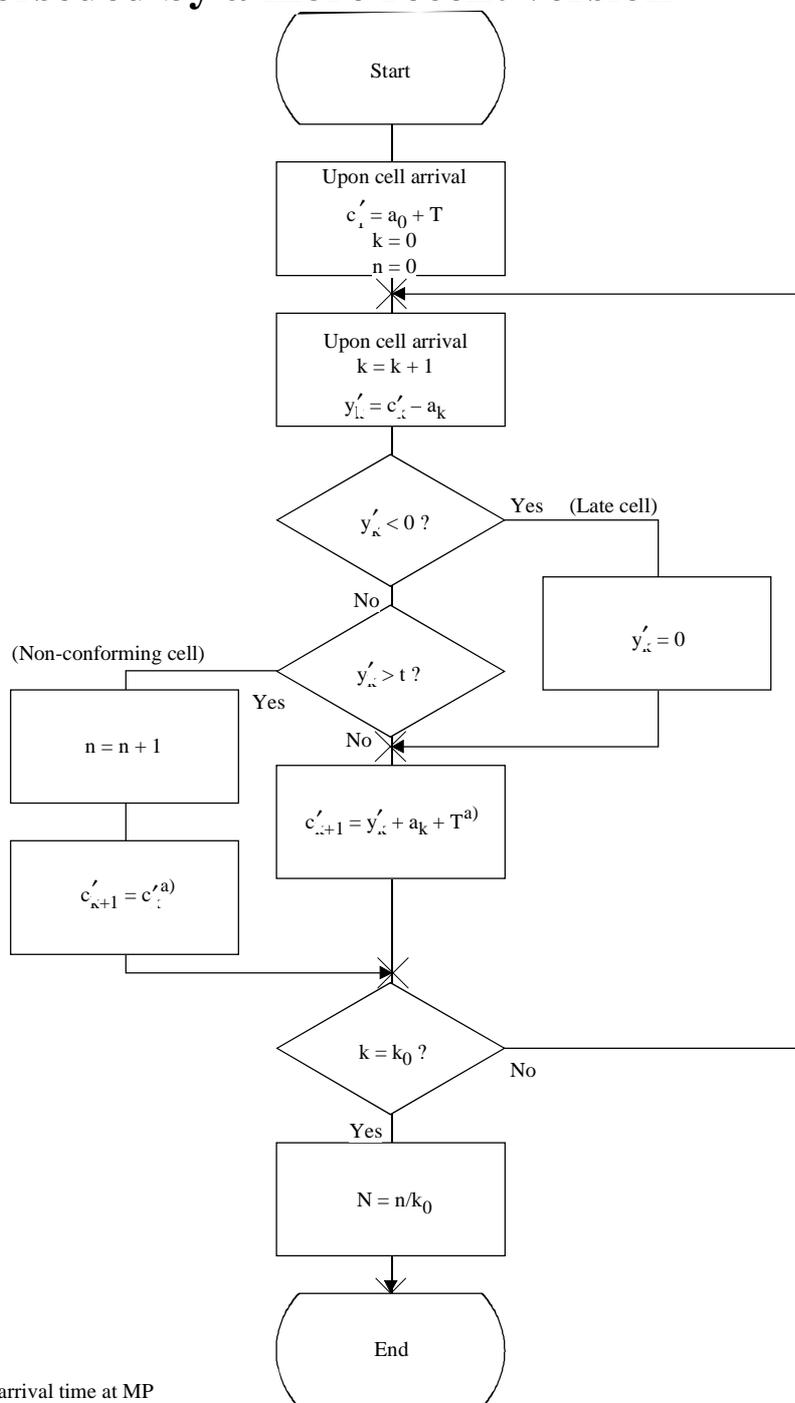
The equations mirror the behaviour of the Generic Cell Rate Algorithm (GCRA) as defined in Recommendation I.371: cell k is non-conforming, as specified by the GCRA formalism, if and only if $y'_k > \tau$.

Superseded by a more recent version

Figure 6 illustrates one measurement method that calculates, for a cell stream received at an MP, the number of cells (n) that do not conform with a specified peak cell rate ($1/T$) and CDV tolerance (τ).⁵ To calculate the cell non-conformance ratio (n/k_0), (n) is divided by the number of cells (k_0) arriving at the MP during the observation period.

⁵ Other methods of calculating non-conforming cell totals are possible, see Annex B.

Superseded by a more recent version



Variables:

c'_k	Cell k reference arrival time at MP
a_k	Cell k actual arrival time at MP
y'_k	Modified 1-point CDV for cell k at MP ($c'_k - a_k$)
t	CDV tolerance at MP
T	Negotiated cell interarrival time
n	Non-conforming cell count
k_0	Measurement size

T1308060-96

$$C'_{K+1} = \begin{cases} Y'_k + a_k + T = a_k + T & \text{if } y'_k < 0 \\ y'_k + a_k + T = c'_k + T & \text{if } 0 \leq y'_k \leq \tau \end{cases}$$

^{a)} Additional updating is required for the continuous state leaky bucket:

a_k if k^{th} cell is conforming

LCT_k otherwise

Figure 6/I.356 – One method of calculating non-conforming cell total for a given CDV tolerance and peak cell rate

Superseded by a more recent version

7.2 Upper bounding the number of non-conforming cells

The set of cells that are identified by the measurement process defined in 7.1 and Figure 6 depends on the starting point of this process (i.e. on the choice of the first cell to be observed). In particular, it is not possible to identify the non-conforming cells independently of the first cell observed by the measurement process. Furthermore, in cases where multiple conformance tests are applied to substreams with non-empty intersections (as is the case with traffic contracts on the aggregate and CLP = 0 streams), the limiting values for the ratios of non-conforming cells may depend on the starting point of the measurement processes.

This subclause addresses the definition of a "maximally constraining test" that yields an upper bound for the maximum number of cells that can be possibly considered as non-conforming within a finite set of consecutive cells. This maximally constraining test does not replace the I.371 definition of cell conformance.

Recommendation I.371 specifies the traffic parameters that are negotiated in a traffic contract. Four basic cases are considered here:

- One single peak cell rate that applies to the aggregate cell stream.
- Two peak cell rate specifications that independently apply to two independent substreams of the connection.
- One peak cell rate specification and one sustainable cell rate specification that both apply to the aggregate cell stream.
- One peak cell rate specification and one sustainable cell rate specification that respectively apply to the aggregate cell stream and to CLP = 0 cells.

The first two cases correspond to the two possible cases identified for the deterministic bit rate (DBR) ATM transfer capability defined in Recommendation I.371. The third case corresponds to the first version of the statistical bit rate ATM transfer capability (SBR1) capability. The fourth case corresponds to the second and third versions of the statistical bit rate ATM transfer capability (SBR2 and SBR3) capability.

Other cases will be developed as the traffic descriptors in Recommendation I.371 evolve.

7.2.1 Bounding cell non-conformance for a single peak cell rate specification

Two conformance tests starting with different values for (c'_0) may yield different numbers of non-conforming cells. However, the difference between the theoretical time and the arrival time is always upper bounded by $(T + \tau)$, and taking $(c'_0 = a_0 + T + \tau)$ will lead to the maximum possible number of non-conforming cells, minus 1. Thus the maximally constraining test for this case is obtained by using the method in 7.1, but with $(c'_0 = a_0 + T + \tau)$, and adding 1 to the resulting number of non-conforming cells.

For a finite set of consecutive cells, let N_{nc} denote the number of cells that are determined to be non-conforming using the I.371 cell conformance definition started with an arbitrarily chosen (possibly earlier) cell. Let also N_u denote the number of cells that are considered as non-conforming by the maximally constraining test applied to the given finite set. The following inequality then applies:

$$N_{nc} \leq N_u$$

Although this maximally constraining test is defined based on a conformance test for a peak cell rate, it can be generalized to any cell rate specified using the GCRA formalism.

NOTE – The development of a tighter upper bound on the number of non-conforming cells is for further study.

Superseded by a more recent version

7.2.2 Bounding cell non-conformance for independent peak cell rate specifications

When independent conformance tests are defined for non-overlapping components of a given cell stream, separate maximally constraining tests (7.2.1) will yield valid upper bounds for the numbers of non-conforming cells that will be observed in each component.

7.2.3 Bounding cell non-conformance for dependent, coordinated cell rate specifications defined on the aggregate cell stream

If one peak cell rate and one sustainable cell rate are negotiated for the aggregate cell flow, Recommendation I.371 provides the definition of cell conformance in the specification of SBR1 capability.

Basically, a cell is considered to be conforming if and only if it is considered to be conforming by both GCRA's defined for the aggregate cell stream. The two GCRA's are specified to be coordinated, which means that the internal variables of the GCRA's are updated only when a cell is found to be conforming to both GCRA's.

The method defined in 7.2.1 for obtaining a maximally constraining test for a single cell conformance test extends to the present case.

Consider a finite set of consecutive cells on the connection. Let N_{nc} denote the number of $CLP = 0 + 1$ cells that are non-conforming for a coordinated cell conformance test applied to this set of cells.

In order to define a maximally constraining test for this conformance definition, one single coordinated conformance test is considered: for each component of this test, the initial theoretical time is taken to be the maximum possible value ($c'_0 = a_0 + T + \tau$). Let N_u denote the number of cells that are considered as non-conforming by this coordinated test. The following inequality then applies:

$$N_{nc} \leq N_u$$

The above method for obtaining a maximally constraining test can be generalized to any coordinated conformance definition using the GCRA formalism and applied to a single cell stream.

7.2.4 Bounding cell non-conformance for dependent, coordinated cell rate specifications defined on the aggregate and on the $CLP = 0$ substream

If a peak cell rate is negotiated for the aggregate cell flow and a sustainable cell rate is negotiated for the $CLP = 0$ flow, Recommendation I.371 provides the definition of cell conformance in the specification of SBR2 and SBR3 capabilities.

Basically, a $CLP = 0$ cell is considered to be conforming if and only if it is considered to be conforming by both the GCRA for the aggregate and the GCRA for the $CLP = 0$ substream. The two GCRA's are specified to be coordinated, which means that the internal variables of the GCRA's are updated only when a cell is found to be conforming in both GCRA's. If cell tagging is permitted (SBR3 capability), a $CLP = 0$ cell that is non-conforming to the $CLP = 0$ peak cell rate may be considered conforming as a $CLP = 1$ cell if it is tagged and conforming to the aggregate peak cell rate.

A $CLP = 1$ cell is considered to be conforming if and only if it is conforming to the aggregate peak cell rate.

Consider first a supposedly infinite cell stream. The above definition of cell conformance is subject to measurement phasing: in some cases, depending on the choice of the first cell to be observed by the conformance process, there may exist different limiting values for the proportions of non-conforming aggregate cells and non-conforming $CLP = 0$ cells.

Superseded by a more recent version

Consider now a finite set of consecutive cells on the connection. Let $N_{nc}(0)$ and $N_{nc}(0 + 1)$ respectively denote the numbers of CLP = 0 and CLP = 0 + 1 cells that are non-conforming for a coordinated cell conformance test applied to this set of cells.

In order to define a maximally constraining test for the coordinated cell conformance definition, two independent maximally constraining tests (7.2.1) are considered:

- the first maximally constraining test is specified by the traffic parameters that define the aggregate peak cell rate, and is applied to the complete set of considered cells;
- the second maximally constraining test is specified by the traffic parameters that define the CLP = 0 sustainable cell rate and is applied only to the CLP = 0 subset of considered cells.

The maximally constraining tests are independently applied to the set of cells in an uncoordinated manner.

For the given set of cells, let $N_u(0)$ denote the number of CLP = 0 cells that are considered as non-conforming by the maximally constraining test performed on the CLP = 0 stream, and let $N_u(0 + 1)$ denote the total number of cells that are considered as non-conforming by the maximally constraining test performed on the aggregate stream. The following inequalities apply:

$$\begin{aligned} N_{nc}(0) &\leq N_u(0) + N_u(0 + 1) \\ N_{nc}(0 + 1) &\leq N_u(0) + N_u(0 + 1) \end{aligned}$$

Thus the upper bound for the numbers of CLP = 0 and CLP = 0 + 1 cells that can be considered as non-conforming is $[N_u = N_u(0) + N_u(0 + 1)]$. In the case where tagging is applied to all CLP = 0 cells found to be non-conforming (SBR3 capability) with the CLP = 0 peak cell rate, a smaller upper-bound $[N_u(0 + 1)]$ can be considered for the aggregate cell stream.

The above method for obtaining a maximally constraining test can be generalized to any cell conformance definition for multiple overlapping substreams each specified using the GCRA formalism.

7.3 Adjusted CLR performance when there are non-conforming cells

In order to take account of the cells that can be discarded in case of cell non-conformance:

- the performance objectives for CLR would not apply to the total set of transmitted cells. The number of cells that should be considered is the number of cells that are conforming to the maximally constraining test, excluding cells in SECBs;
- the lost cell (and tagged cell) outcomes used in the definition of CLR should only include those lost (and tagged) cells that are in excess of the number of cells that have been identified as non-conforming by the maximally constraining test, excluding lost and tagged cells in SECBs.

The maximally constraining conformance tests derived in 7.2 allow for the definition of adjusted cell loss ratio parameters.

7.3.1 CLR when the conformance specification applies to a single cell stream

For a given set of consecutive cells of size N_t , let N_l and N_u respectively denote the number of cells that have been lost (or tagged, if applicable), and the upper bound for the number of non-conforming cells computed by the maximally constraining conformance tests defined in 7.2.1 or 7.2.3. The adjusted cell loss ratio parameter CLR_{mod} is defined as:

$$CLR_{mod} = \frac{\max(0, N_l - N_u)}{\max(0, N_t - N_u)}$$

Superseded by a more recent version

7.3.2 CLR when the conformance specification applies to the aggregate and the CLP = 0 substream

For a given set of consecutive cells of size $N_t(0 + 1)$, let $N_t(0)$, $N_t(0 + 1)$, $N_l(0)$ and N_u respectively denote the number of CLP = 0 cells that have been transmitted, the total number of cells that have been lost, the number of CLP = 0 cells that have been lost or tagged, and the upper bound for the number of non-conforming cells computed by the maximally constraining conformance test. The adjusted cell loss ratio parameters are defined as:

$$\text{CLR}_{0+1,\text{mod}} = \frac{\max(0, N_t(0+1) - N_u)}{\max(0, N_r(0+1) - N_u)}$$

$$\text{CLR}_{0,\text{mod}} = \frac{\max(0, N_t(0) - N_u)}{\max(0, N_r(0) - N_u)}$$

NOTE – For the SBR2, ATM transfer capability the N_u in these two equations are the same. For the SBR3 ATC the N_u in these two equations are different. (See 7.2.)

The definition of an adjusted CLR_1 parameter is for further study.

7.4 Non-conforming cells and the severely errored cell block outcome

As defined in 5.6, a Severely Errored Cell Block (SECB) is a sequence of N cells, transmitted consecutively on a given connection, for which more than M cells are errored, lost or misinserted. Since the network is allowed to discard non-conforming cells, a cell block could wrongly be considered as severely errored if the effect of non-conforming cells is not excluded when assessing whether more than M cells have been lost. Therefore, if some cells in the block are non-conforming, the comparison with M should consider only those lost cell outcomes that are in excess of the cells identified as non-conforming by the maximally constraining tests of 7.2.

Since traffic control functions may not be synchronized with the maximally constraining test, the cells that the maximally constraining test considers non-conforming might not be those discarded by the traffic control functions. An equivalent number of cells could be discarded, but those could belong to a different cell block. The ambiguity created by this situation is for further study.

NOTE – It is anticipated but not required that the UPC/NPC will take action on the cell stream as soon as non-conformance is detected.

8 Network performance objectives

This clause discusses objectives for the user information transfer performance of public B-ISDNs. These objectives are stated in terms of the ATM layer performance parameters defined in Clause 6⁶. A summary of the objectives can be found in Table 2 together with its associated general notes. All values in Table 2 are provisional and they need not be met until they are revised (up or down) based on real operational experience.

This Recommendation differs from other ITU-T performance Recommendations because:

- the user has the option of requesting a different Quality of Service (QOS) for each new VP and VC connection; and

⁶ Some network providers may support performance objectives even when some cells are non-conforming. In these cases, the adjusted parameter definitions of clause 7 are one way of comparing network performance with the numeric objectives of clause 8.

Superseded by a more recent version

- for some QOS classes and certain performance parameters, the ITU-T will not recommend any minimum level of quality.

8.1 General discussion of per-connection QOS

While establishing a new connection (VP or VC), calling users can signal their preferred quality of service (QOS) class from among those presented in Table 2. As the connection is established, the network providers sequentially commit to support the requested QOS class by progressing the connection request toward the called user. If one of the networks is unable to support the requested QOS class, that network will clear the connection request using an appropriate message. If all of the network providers agree to the requested QOS class, the QOS class definition in Table 2 presents bounds on the end-to-end network performance. As long as the users comply with their traffic contract, the network providers should collaboratively support these end-to-end bounds for the duration of the connection.

The actual QOS offered to a given connection will depend on its distance and complexity. It will often be better than the bounds included with the QOS class definitions in Table 2. See clause 9 for relevant additional information.

On a connection-by-connection basis, users may request and receive different QOS classes. In this fashion, the distinct performance needs of different services and applications can be met.

8.2 QOS classes

This subclause describes the currently defined QOS classes. Each QOS class creates a specific combination of bounds on the performance values. This subclause includes guidance as to when each QOS class might be used, but it does not mandate the use of any particular QOS class in any particular context.

8.2.1 Nature of the network performance objectives

The objectives in Table 2 apply to public B-ISDNs, MPT-to-MPT. The objectives are believed to be achievable on complex 27 500 km connections. These objectives do not account for the performance of private networks or other CEQ performance. CEQ performance is a subject for further study.

The left-hand part of Table 2 indicates the statistical nature of the performance objectives that appear in the subsequent rows. Statistical estimation issues are discussed below.

The performance objectives for cell transfer delay are upper bounds on the underlying mean CTD for the connection. Although many individual cells may have transfer delays that exceed this bound, the average CTD for lifetime of the connection (a statistical estimator of the mean) should normally be less than the CTD bounds.

Recommendation Q.2931 enables the user to signal a maximum acceptable CTD. If the user takes advantage of this facility, signalling nodes will deliver an estimate of end-to-end CTD to the called user. Taking account of this CTD estimate, the users may or may not accept the call. The relationship between the signalled CTD values and the network providers' commitment to CTD performance remains to be studied.

The performance objectives for 2-point cell delay variation are upper bounds on the difference between the 10^{-8} and the $1-10^{-8}$ quantiles of the underlying CTD distribution for the connection. Thus within the connection it should be very difficult to find any two cells with a difference in CTD larger than the CDV bounds. 10^{-8} was chosen because it allows for the proper engineering of delay buildout buffers when the overall CLR objective is 10^{-8} . The use of other quantiles for 2-point CDV specification is for further study.

Superseded by a more recent version

The CDV objectives only apply to connections that have negotiated appropriately small CDV tolerances in conjunction with their PCRs. A network's CDV objective does not include the 2-point CDV resulting from actions taken at the network ingress to reduce the amount of 1-point CDV. These network actions are not considered a network induced degradation.

The performance objectives for the cell loss ratios and the cell error ratio are upper bounds on the cell loss and cell error probabilities for the connection. Although individual cells will be lost or errored, the underlying probability that any individual cell is lost or errored during the connection should be less than the bounds presented in Table 2. When small numbers of cells are observed, it is possible that the computed CLR_{0+1} , CLR_0 , and CER will be greater than the bounds for cell loss and cell error probabilities.

The performance objective for the cell misinsertion rate is an upper bound on the underlying mean rate at which misinserted cell outcomes occur. For a set of connections of sufficient total duration, the computed CMR should be less than the CMR bound.

The performance objective for the severely errored cell block ratio is an upper bound on the SECB probability. Although individual cell blocks will be severely errored, the underlying probability that any individual cell block is severely errored should be less than the bounds presented in Table 2. When small numbers of cell blocks are observed, it is possible that the computed SECBR is greater than the bound for SECBR.

8.2.2 Statistical estimation issues

The evaluation of the per-connection quality of service commitment, including measurement requirements, statistical issues, and caveats is for further study. The following statistical questions are to be considered:

- With what precision must the performance parameters be measured in order to compare the observed performance with the QOS class commitment?
- How will a per-connection QOS commitment be verified when the total number of cells transferred during the life of the connection is small?
- How are short-term variations in performance (e.g. hourly, daily, and weekly variations) addressed when comparing the observed performance with the QOS class commitment?
- How can the 10^{-8} and $1-10^{-8}$ quantiles of the CTD distribution be estimated?

8.2.3 Unbounded (unspecified) performance

For some QOS classes the values for some performance parameters is designated "U". In these cases, the ITU-T sets no objectives regarding these parameters and any default I.356 objectives for these parameters can be ignored. Network operators may unilaterally elect to assure some minimum quality level for the unspecified parameters, but the ITU-T will not recommend any such minimum.

Users of these QOS classes should be aware that the performance of unspecified parameters can, at times, be arbitrarily poor.

NOTE – The word "unspecified" may have a different meaning in Recommendations concerning B-ISDN signalling.

8.2.4 Default values for cell error ratio, cell misinsertion rate, and severely errored cell block ratio

The CER, CMR, and SECBR values cannot easily be adjusted on connection-by-connection basis. Therefore the performance commitments for these parameters do not differ among the QOS classes. The exception is that no commitment will be made to these parameters in the U class.

Superseded by a more recent version

Table 2/I.356 – Provisional QOS class definitions and network performance objectives

	Nature of the network performance objective	Default objectives	QOS Classes			
			Class 1 (stringent class)	Class 2 (tolerant class)	Class 3 (bi-level class)	U class
CTD	Upper bound on the mean CTD	No default	400 ms (Notes 4, 5)	U	U	U
2-pt. CDV	Upper bound on the difference between upper and lower 10^{-8} quantiles of CTD	No default	3 msec (Note 6)	U	U	U
CLR₀₊₁	Upper bound on the cell loss probability	No default	3×10^{-7} (Note 7)	10^{-5}	U	U
CLR₀	Upper bound on the cell loss probability	No default	None	None	10^{-5}	U
CER	Upper bound on the cell error probability	4×10^{-6} (Note 1)	Default	Default	Default	U
CMR	Upper bound on the mean CMR	1/day (Note 2)	Default	Default	Default	U
SECBR	Upper bound on the SECB probability	10^{-4} (Note 3)	Default	Default	Default	U

General Notes to Table 2:

All values are provisional and they need not be met by networks until they are revised (up or down) based on real operational experience.

The objectives apply to public B-ISDNs, MPT-to-MPT. The objectives are believed to be achievable on the 27 500 km hypothetical reference connections presented in Appendix II. The network providers' commitment to the user is to attempt to build end-to-end connections achieving each of the applicable objectives. The vast majority of public network connections should meet those objectives. When the MPTs are separated by large geographic distances, the probability of not meeting the applicable objectives is increased. For some parameters, performance on shorter and/or less complex connections may be significantly better.

Individual network providers may choose to offer performance commitments better than their allocated objectives.

"U" means "unspecified" or "unbounded". When the performance relative to a particular parameter is identified as being "U" the ITU-T establishes no objective for this parameter and any default I.356 objective can be ignored. When the objective for a parameter is set to "U", performance with respect to that parameter may, at times, be arbitrarily poor.

NOTE 1– It is possible that in the near future, networks will be able to commit to a CER of 4×10^{-7} . This subject is for further study.

NOTE 2 – Some network phenomena have been observed that tend to increase the CMR as the cell rate of the virtual connection increases. More complete analyses of these phenomena may ultimately suggest a larger CMR objective for high bit rate connections.

NOTE 3 – The SECBR is sensitive to short interruptions in the cell stream (i.e. 2 to 9 seconds in duration) which will result in many SECBs and may make the SECBR objective difficult to meet.

Superseded by a more recent version
Table 2/I.356 – Provisional QOS class definitions and network performance objectives (concluded)

NOTE 4 – See Recommendation G.114 for further guidance on the delay requirements of some applications.

NOTE 5 – Some applications may require performance similar to QOS class 1, but do not require a CTD commitment. These applications can make use of QOS class 1, but the need for a new QOS class is a subject for further study.

NOTE 6 – Applies when there are no more than 9 ATM nodes in the connection with 34 to 45 Mbit/s output links and all other ATM nodes are operating at 150 Mbit/s or higher. 2-pt. CDV will generally increase as transport rates decrease. High bit rate DBR connections may need and may receive less CDV. This is for further study.

NOTE 7 – It is possible that in the near future, networks will be able to commit to a CLR for class 1 of 10^{-8} . This subject is for further study.

8.2.5 Association of QOS classes with ATM transfer capabilities

Recommendation I.371 defines ATM transfer capabilities (ATC). Table 3 recommends that certain ATC be associated with certain QOS classes. This material does not mandate the use of any particular QOS class in any particular context.

Table 3/I.356 – Association of ATCs with QOS classes

ATM transfer capabilities (ATC)	Applicable QOS class
DBR, SBR1, ABT/DT, ABT/IT	Class 1 (stringent class)
DBR, SBR1, ABT/DT, ABT/IT	Class 2 (tolerant class)
SBR2, SBR3, ABR	Class 3 (bi-level class)
Any ATC	U class

The QOS class commitments for ABT/DT, ABT/IT, and ABR only apply when users adhere to the relevant conformance definitions for these ATC. Performance commitments when users do not follow the conformance definitions are for further study.

No CDV commitment can be made to a connection using ABT/IT when the elastic/rigid bit is set to zero (see Recommendation I.371).

8.3 Alternative QOS negotiation procedures

Other, more complex methods of negotiating and supporting quality of service needs are under study. In the future there may be more complete protocols for "negotiating" quality of service between users and networks. The ability to negotiate "connection blocking" and "connection cutoff" probabilities on a per-connection basis is for further study.

This Recommendation does not preclude the future possibility of negotiating QOS by signalling individual parameter values instead of using QOS classes to select groups of parameter values. Recommendation Q.2931 enables the user to signal a maximum acceptable CTD. This subject is for further study.

Superseded by a more recent version

9 Allocation of the performance objectives

An analysis of several hypothetical reference connections (HRX) demonstrated that the objectives in Table 2 are achievable on long (27 500 km), complex connections. In order to cooperatively achieve those objectives, allocation rules are needed for each standardized portion of the end-to-end connection. The following subclauses list the allocation rules for each parameter. Network providers should attempt to build their connection portions in such a way that the vast majority of their connection portions achieve their allocated objectives for every performance parameter. Thus the performance of a shorter and less complex end-to-end connection will often be better than that represented by Table 2.

The rules used in computing allocations should not be interpreted as recommendations for implementation. For example, the CTD that the computation allows for route length could be used for additional ATM nodes instead. The goal is to achieve the allocated objectives using whatever strategies the network operator deems appropriate.

9.1 General principles of allocation

The allocation rules for several of the objectives are based on the G.826 rules for allocating physical layer performance. Physical layer impairments contribute strongly to the ATM layer performance parameters SECBR, CER, and CLR.

Performance impairments for each ATM layer parameter grow with increasing "distance" and increasing "complexity". In this context, the term "complexity" refers to impairments that increase with additional switching and queueing stages and/or increase as more international and jurisdictional boundaries are crossed. The term "distance" refers to impairments that are not directly tied to switching or queueing stages and are less directly controllable with ATM network design. In the following allocation rules, block allocations are given to connection portions to allow for impairments due to "complexity", and route length allocations are given to connection portions to allow for impairments due to "distance". Appendix III presents examples illustrating the use of some of these allocation rules.

Portions that contain geostationary satellites receive relatively large block allocations for several parameters. However, a geostationary satellite is normally expected to span significant terrestrial distances and eliminate the need for multiple ATM switching nodes and/or transit country portions. It is not expected that a connection would include more than one geostationary satellite hop when providing QOS class 1. Geostationary satellite systems that include ATM switching and processing, such as on-board processing satellites, are for further study, but it is reasonable to assume that they would receive an allocation of performance, including CDV, to accommodate their ATM functions. Performance allocations for portions that contain low- and medium- Earth orbit satellites are also for further study.

9.2 Route length calculation

Several of the parameters have a part of their allocation proportional to terrestrial route length. The route length calculation is taken from Recommendation G.826. If D_{km} is the air-route distance between the two MPs that bound the portion, then the route length calculation is:

- if $D_{km} < 1000$ km, $R_{km} = 1.5 \times D_{km}$
- if $1000 \leq D_{km} \leq 1200$ km, $R_{km} = 1500$ km
- if $D_{km} > 1200$, $R_{km} = 1.25 \times D_{km}$

The above rule does not apply when the portion contains a satellite hop.

Superseded by a more recent version

9.3 Allocation of the QOS class 1 CTD objective

This subclause calculates the maximum CTD allocation for any connection portion supporting a QOS class 1 connection.

When a connection portion does not contain a satellite hop, its computed CTD allocation is:

$$\text{CTD (in microseconds)} \leq (R_{\text{km}} \times 6.25) + (N_{\text{sw}} \times 300)$$

In this formula:

- R_{km} represents the route length assumption computed in 9.2.
- $(R_{\text{km}} \times 6.25)$ is an allowance for "distance" within the portion.
- N_{sw} is taken from Table 4.
- $(N_{\text{sw}} \times 300)$ is an allowance for the "complexity" of the portion.

Table 4/I.356 – N_{sw} : number of ATM switching and cross-connect stages assumed in the computation of CTD allocations

	National portion	IIP(0)	IIP(1)	IIP(2)	IIP(3)	ITP
VCC	8 nodes (VC or VP)	0 nodes	3 VP nodes	6 VP nodes	9 VP nodes	3 nodes (VC or VP)
VPC	4 VP nodes	0 VP nodes	Not applicable	Not applicable	Not applicable	3 VP nodes

The 300 μ sec value is considered as an approximate worst-case value for ATM nodes providing class 1 service. A corresponding value for other classes is for further study.

When a connection portion contains a satellite hop, this portion is allocated a fixed CTD. Although most portions that include a geostationary satellite are not expected to exceed 290 msec of CTD, all portions containing a geostationary satellite are allocated 320 msec of CTD to account for, e.g. low earth station viewing angle and low rate TDMA systems.

It is expected that in most cases, the end-to-end CTD that results when each connection portion complies with its allocation will be below 400 msec. However, it may happen in some cases that the value of 400 msec is exceeded. For very long connections to remote areas, network providers may need to make additional bilateral agreements to improve the probability of achieving the 400 msec objective.

9.4 Allocation of the QOS class 1 CDV objective

This subclause indicates the maximum CDV allocation for any connection portion supporting a QOS class 1 connection.

- A national portion of the international connection is allowed 1.5 msec of CDV. This allocation applies when there are no more than 3 ATM nodes in the national portion with 34 to 45 Mbit/s output links and all other ATM nodes in the portion are operating at 150 Mbit/s or higher.
- The international portion of the international connection is allowed 1.5 msec of CDV. This allocation applies when there are no more than 3 ATM nodes in the international portion with 34 to 45 Mbit/s output links and all other ATM nodes in the portion are operating at 150 Mbit/s or higher.
- An IIP(0) receives essentially no allocation of CDV.

Superseded by a more recent version

- For an ITP take a block allowance of 0.7 msec. This allocation applies when there is no more than one ATM node in the portion with 34 to 45 Mbit/s output links and all other ATM nodes in the portion are operating at 150 Mbit/s or higher.
- For an IIP(1) take a block allowance of 0.7 msec. This allocation applies when there is no more than one ATM node in the portion with 34 to 45 Mbit/s output links and all other ATM nodes in the portion are operating at 150 Mbit/s or higher.
- For an IIP(2) take a block allowance of 0.9 msec. This allocation applies when there are no more than two ATM nodes in the portion with 34 to 45 Mbit/s output links and all other ATM nodes in the portion are operating at 150 Mbit/s or higher.
- For an IIP(3) take a block allowance of 1.1 msec. This allocation applies when there are no more than three ATM nodes in the portion with 34 to 45 Mbit/s output links and all other ATM nodes in the portion are operating at 150 Mbit/s or higher.

The allocated CDVs sum to more than the end-to-end CDV because CDV accumulates similarly to the standard deviation of roughly independent random variables. When independent random variables are summed, the resulting standard deviation is roughly the square root of the sum of the squares.

9.5 Allocation of the SECBR and CER objectives

This subclause calculates the maximum SECBR and CER allocations for any connection portion. These allocations begin with the end-to-end objectives found in Table 2. This method is based on the allocation rules of Recommendation G.826.

- Round the calculated route length, R_{km} , for the portion (national portion, IIP, ITP) up to the nearest 500 km.
- For a national portion take a block allowance of 17.5% plus 1% per 500 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 42% block allowance replaces this computation.
- For an IIP(0) take a block allowance of 1% plus 1% per 500 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 35% block allowance replaces this computation.
- For an ITP take a block allowance of 2% plus 1% per 500 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 36% block allowance replaces this computation.
- For an IIP(1) take a block allowance of 4% plus 1% per 500 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 38% block allowance replaces this computation.
- For an IIP(2) take a block allowance of 7% plus 1% per 500 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 42% block allowance replaces this computation.
- For an IIP(3) take a block allowance of 10% plus 1% per 500 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 48% block allowance replaces this computation.

NOTE 1 – It is expected that portions will achieve their allocations for CER even for QOS class U. CER is primarily governed by transmission performance. However, no commitments are made to the CER because OAM PM capabilities are ineffective in class U.

NOTE 2 – There are no QOS commitments for the SECBR in the QOS class U.

Superseded by a more recent version

NOTE 3 – Using these computations, the 27 500 km HRXs in Appendix II are allocated 100% of the end-to-end objective.

9.6 Allocation of the QOS class 1 CLR objective

This subclause calculates the maximum CLR allocation for any connection portion in support of a QOS class 1 connection. Both physical layer impairments and ATM network complexity play a significant role in the end-to-end performance of class 1 CLR, and hence its allocation is different from the physical layer allocation rules of Recommendation G.826. The allocation begins with the end-to-end objective found in Table 2.

- Round the calculated route length, R_{km} , for the portion (national portion, IIP, ITP) up to the nearest 1000 km.
- For a national portion take a block allowance of 23% plus 1% per 1000 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 35% block allowance replaces this computation.
- For an IIP(0) take a block allowance of 1% plus 1% per 1000 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 25% block allowance replaces this computation.
- For an ITP take a block allowance of 7% plus 1% per 1000 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 30% block allowance replaces this computation.
- For an IIP(1) take a block allowance of 9% plus 1% per 1000 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 30% block allowance replaces this computation.
- For an IIP(2) take a block allowance of 17% plus 1% per 1000 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 33% block allowance replaces this computation.
- For an IIP(3) take a block allowance of 25% plus 1% per 1000 km, if no satellite hop. If there is a geostationary satellite hop within the portion, a single 42% block allowance replaces this computation.

NOTE – Using these computations, the 27 500 km HRXs in Appendix II are allocated 100% of the end-to-end objective.

9.7 Allocation of the QOS class 2 and QOS class 3 CLR objectives

This subclause calculates the maximum CLR allocations for any connection portion in support of a QOS class 2 or class 3 connection. In these classes network complexity (in this case, buffer management) is allowed to play the dominant role in the end-to-end performance. One consequence is that this allocation rule does not make use of the calculated route length. The allocation begins with the end-to-end objective found in Table 2.

- For a national portion take 34.5% of the end-to-end objective.
- For an IIP(0) take a block allowance of 1% of the end-to-end objective.
- For an ITP take a block allowance of 9% of the end-to-end objective.
- For an IIP(1) take a block allowance of 11% of the end-to-end objective.
- For an IIP(2) take a block allowance of 21% of the end-to-end objective.
- For an IIP(3) take a block allowance of 31% of the end-to-end objective.

Superseded by a more recent version

NOTE – Using these computations, the 27 500 km HRXs in Appendix II are allocated 100% of the end-to-end objective.

9.8 Allocation of the CMR objective

This subclause indicates the maximum CMR allocation for a national portion or an international portion in support of the end-to-end objective of 1/day.

- Both national portions of the international connection are allowed a CMR of 1 per 72 hours.
- The international portion of the international connection is allowed a CMR of 1 per 72 hours. Allocation to the IIPs and ITPs of the connection is for further study.

NOTE – It is expected that portions will achieve their allocations for CMR even for QOS class U. CMR is primarily governed by transmission performance and the ATM HEC. However, no commitments are made to the CMR because OAM PM capabilities are ineffective in class U.

Most networks should be easily able to achieve these allocated CMR objectives. They are presented here to inform potential users that cell misinsertions are expected to be rare and to remind network designers that the objective is to make the CMR imperceptible.

9.9 Concatenation of QOS values

This subclause addresses the derivation of the end-to-end performance of a connection, knowing the performance of each portion. For all performance parameters except CDV, the end-to-end performance is the sum of the portion values. The rule for deriving the end-to-end CDV performance from the portion values is sub-additive and is for further study.

ANNEX A

Relationship between ATM layer NP and the NP of AAL type 1 for CBR services

This Annex describes qualitative relationships between ATM layer network performance and the NP provided by the type 1 AAL.

A.1 Possible AAL functions and their effects

Examples of adaptation layer functions which may compensate for specific performance impairments introduced in ATM cell transfer are provided below.

A.1.1 Lost cell outcome and misinserted cell outcome

A Sequence Number (SN) in the AAL header can be used to detect the lost and misinserted AAL SDUs due to lost cell and misinserted cell outcomes. Detection mechanisms are for further study.

If cell losses are detected, replacement AAL SDUs may be substituted to compensate for the lost cells in order to maintain bit count integrity. However, if there is no error correction in the AAL, this substitution will result in user information bit errors in the AAL SDU. The contents selected for such dummy AAL SDUs (e.g. all "1", all "0", repeat the previous cell, etc.) require further study (see Recommendation I.363.1).

If misinserted cells are detected, they may be discarded, restoring the delivered user information content to that transmitted.

If lost cells and misinserted cells are not detected, they may cause a loss of frame alignment in the delivered user information stream.

Superseded by a more recent version

A.1.2 Errored cell outcome

Error control mechanisms have been identified for some signals transported by AAL type 1. In the absence of such error control, bit errors will be transferred to the AAL user.

A.1.3 Cell transfer delay

To compensate for cell delay variation, arriving cells are buffered in the AAL at the receiving side of a connection. This buffering increases the user information transfer delay. Error control and lost cell detection mechanisms may introduce additional delay.

Excessive cell delay variation that cannot be compensated or excessive delay due to a lost cell detection mechanism can cause the substitution of dummy AAL SDUs for valid AAL SDUs, resulting in bit errors in user information.

A.2 Bounding relationships between NP parameters and binary errors

In the absence of error control covering the cell information field:

- the expected number of binary errors associated with each lost cell is 188 (assuming 47 octets of AAL user information in the ATM cell payload and a bit error ratio of 0.5) if dummy AAL SDUs are inserted;
- an errored cell outcome can theoretically produce any number of errored bits from 1 to 376 (assuming 47 octets of AAL user information in the ATM layer cell payload), with a distribution skewed towards the low end of the theoretical range;
- each misinserted cell delivered to the AAL user – i.e. not dropped by the AAL – results in binary errors. Furthermore, an undetected misinsertion could cause a loss of frame alignment.

ANNEX B

Cell transfer delay, 1-point CDV and 2-point CDV characteristics

B.1 Components of delay associated with ATM-based user information transfer

The overall delay perceived by an end-user of AAL service can be divided into the following components:

- T1 *Coding and decoding delay* (see Note 1).
- T2 *Segmentation and reassembly delay* (see Note 1).
The latter delay can be further subdivided into three:
 - T21 Segmentation delay in the AAL of the sending side.
 - T22 Buffering delay in the AAL of the receiving side to compensate the cell delay variation (see Note 2).
 - T23 Reassembly delay in the AAL of the receiving side.
- T3 *Cell transfer delay (MPT-MPT)*.
This delay is the sum of the following
 - T31 Total inter-ATM node transmission delay (see Note 3).
 - T32 Total ATM node processing (queueing, switching and routing, etc.) delay (see Notes 4 and 5).

Superseded by a more recent version

NOTE 1 – Coding and data segmentation may or may not be performed in the same equipment. Similarly, reassembly and decoding may or may not be performed in the same equipment.

NOTE 2 – The amount of buffering delay consumed in AAL handling equipment will depend on the amount of network cell delay variation.

NOTE 3 – Delay caused by transmission related equipment(s) between two adjacent ATM nodes, e.g. SDH cross-connect systems, is considered to be part of this component.

NOTE 4 – ATM nodes may perform either Virtual Channel (VC) or Virtual Path (VP) switching or cross-connect functions.

NOTE 5 – Due to queuing in ATM nodes, this component is variable on a cell-by-cell basis within one ATM connection.

B.2 Relationship between cell clumping and cell queues

With respect to a particular MP, define a *clump* as a sequence of early cell arrivals between two successive reference clock skips. The corresponding time interval is a *positive queue interval*. Clumps can be considered to increase the aggregate length of cell queues downstream of the MP.

B.3 1-point CDV and non-conformance

A virtual connection provides negotiated values for peak emission interval T (inverse of peak cell rate) and CDV tolerance τ . As long as the y_k value computed as in 6.5.2.1 is smaller than τ , cell k is observed as conforming with the specified peak cell rate ($1/T$) and the CDV tolerance (τ).

However, when some cells are observed as non-conforming (i.e. $y_k > \tau$), it is useful to measure the number of non-conforming cells in a given cell stream. Figure 6 illustrates one measurement method that calculates the number of cells that do not conform with a specified peak cell rate ($1/T$) and CDV tolerance (τ).

The method of Figure 6 is an example, and is not intended to provide any specific implementation or hardware mechanism for measuring the cell non-conformance ratio (n/k_0). The virtual scheduling and leaky bucket algorithms described in Recommendation I.371 as equivalent peak cell rate monitoring algorithms may be used to implement the measurement of non-conformance ratio. To facilitate comparison of such implementations, the mapping between the variables of the equivalent algorithms is summarized in Table B.1.

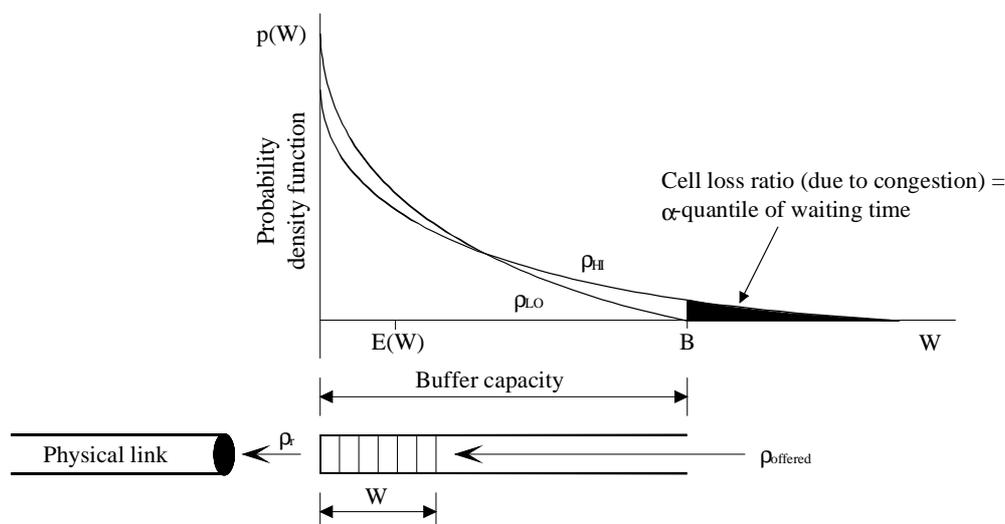
Table B.1/I.356 – Mapping between the variables defined in Figure 6/I.356 and those of virtual scheduling and continuous state leaky bucket algorithms defined in Recommendation I.371

Variables defined in various algorithms	Figure 6/I.356	Virtual scheduling	Leaky bucket
Theoretical arrival time of cell k	c'_k	TAT	$x + \text{LCT}$
Actual arrival time	a_k	t_a	t_a
Modified 1-point CDV parameter for cell k	y'_k	$\text{TAT} - t_a$	x'
Parameter values at first observed arrival time	$c'_0 = a_0$	$\text{TAT} = a_0$	$x = 0$ $\text{LCT} = a_0$

Superseded by a more recent version

B.4 Relationship between cell transfer delay and cell loss in a single shared buffer

Consider the operation of one of the physical links which support a specific ATM connection. All of the cells that are intended to pass through this physical link would be held in a buffer that absorbs momentary surpluses of cells until they are either transmitted over the link, or until this buffer overflows with the resultant loss of some cells. The cells that are intended to pass through this physical link are provided by both the specific ATM connection under consideration and other ATM connections which share this link, and all of these cells combine to establish the link's offered load, which may be characterized by a utilization factor ρ_{offered} . Any cell arriving at this buffer experiences a random waiting time W before it reaches the link and is transmitted. Figure B.1 illustrates this situation, together with some representative probability density functions for W .



T1308070-96

Figure B.1/I.356 – Illustration of random waiting time (W)

With a sufficiently high value of offered load, characterized in Figure B.1 by ρ_{HI} , the tail of the probability density function will place a significant amount of weight beyond the buffer capacity B , as measured in cell emission times⁷. The area under this curve can be interpreted as the cell loss ratio due to congestion. If the buffer were made larger, these cells would not overflow and the shaded area would then represent an upper quantile of cell waiting time.

The maximum waiting time for a cell in this buffer occurs when the cell in question occupies the final available cell space. Thus the maximum delay variation attributable to this buffer is controlled by the buffer size.

With a lower value of offered load, characterized in Figure B.1 by ρ_{LO} , the tail of the probability density function will place less weight beyond B , thereby reducing the resulting value of cell loss ratio.

These effects should be considered in the selection of cell transfer delay timeout T_{max} (a number that should exceed the largest practically conceivable cell transfer delay), and in the specification of 2-point CDV and cell loss ratio values.

⁷ One cell emission time on an STM-1 link is 2.73 microseconds. If, for example, a buffer has 100 cells and feeds an STM-1 link, B would be 273 microseconds.

Superseded by a more recent version

ANNEX C

Cell transfer performance measurement methods

This Annex describes measurement methods which may be used to estimate values for the ATM cell transfer performance parameters defined in this Recommendation. The described methods include in-service methods, which introduce OAM cells into the transmitted user information cell stream, and out-of-service methods, which involve performing measures on a test connection dedicated to measurement. The in-service methods include direct methods, which make use of information derived from the user cell stream (e.g. cell counts), and indirect methods, which rely on the similarity between performance for user cell transfer and performance for OAM cell transfer. The in-service methods allow continued use of the channel under measurement; the out-of-service methods allow greater control of the measurement process and can generally provide better measurement precision.

NOTE 1 – The accuracy of measurement of these events will be no better than about plus or minus 200 microseconds at SDH interfaces if cell entry/cell exit events for cells embedded in SDH frames are approximated by the frame event times.

NOTE 2 – Cell transfer performance parameters are measured during periods of available time (see Recommendation I.357). Cell transfer performance measurement should be disabled during periods of unavailable time.

Measurement methods are described below for cell error ratio, cell loss ratio, cell misinsertion rate, severely errored cell block ratio, cell transfer delay, and 2-point cell delay variation.

A possible approach for out-of-service monitoring is to establish a virtual path on connection at the appropriate measurement point, introduce a cell stream of known content and timing at that point, and then observe the cell stream at the remote measurement point.

Details of OAM functions supporting performance measurement are provided in Recommendation I.610. In-service performance monitoring will likely be performed only on a selected number of Virtual Path Connections/Virtual Channel Connections (VPCs/VCCs) on an on-demand basis.

NOTE 3 – The use of AAL protocol mechanisms in ATM layer performance measurement is for further study.

C.1 General aspects of performance monitoring using OAM cells

Figures C.1 and C.2 illustrate the general approach envisioned for use of OAM cells in performance monitoring. Performance monitoring OAM cells may be introduced into the cell stream at any VP or VC termination or connecting point, and may then be observed or extracted at any similar point downstream.

Measurement methods based on OAM functions use information:

- carried by the performance monitoring OAM cells;
- gathered at the point where performance is estimated.

OAM FM (forward monitoring) cells are inserted into the user information cell stream at suitable intervals in order to delineate blocks of user cells. Each OAM cell carries the running counter values for the numbers of:

- transmitted CLP = 0 + 1 user cells (TUC_{0+1});
- transmitted CLP = 0 user cells (TUC_0).

This information allows the computation of the number of CLP = 0 + 1 cells transmitted in the block (nt_{0+1}), and the number of CLP = 0 cells transmitted in the block (nt_0).

Superseded by a more recent version

Let nr_{0+1} (respectively nr_0) be the number of CLP = 0 + 1 user cells (respectively, the number of CLP = 0 user cells) received in the block at the point where performance is estimated.

The following information is needed to estimate delivered performance:

- the total number Nt_{0+1} of transmitted CLP = 0 + 1 user cells excluding those transmitted in SECBs. For each block not an SECB, nt_{0+1} is added to Nt_{0+1} ;
- the total number Nt_0 of transmitted CLP = 0 user cells excluding those transmitted in SECBs. For each block not an SECB, nt_0 is added to Nt_0 ;
- the total number $N't_{0+1}$ of transmitted CLP = 0 + 1 user cells excluding those transmitted in blocks in which either lost or misinserted cell outcomes are detected. For each block in which neither lost nor misinserted cell outcomes are detected, nt_{0+1} is added to $N't_{0+1}$;
- the total number Nl_{0+1} of CLP = 0 + 1 lost cells excluding those lost in SECBs. For each block not an SECB, the difference $nt_{0+1} - nr_{0+1}$, when it is positive, is added to Nl_{0+1} ;
- the total number Nl_0 of CLP = 0 lost or tagged cells excluding those lost or tagged in SECBs. For each block not an SECB, the difference $nt_0 - nr_0$, when it is positive, is added to Nl_0 ;
- the total number Nm_{0+1} of CLP = 0 + 1 misinserted cells excluding those misinserted in SECBs. For each block not an SECB, the difference $nr_{0+1} - nt_{0+1}$, when it is positive, is added to Nm_{0+1} .

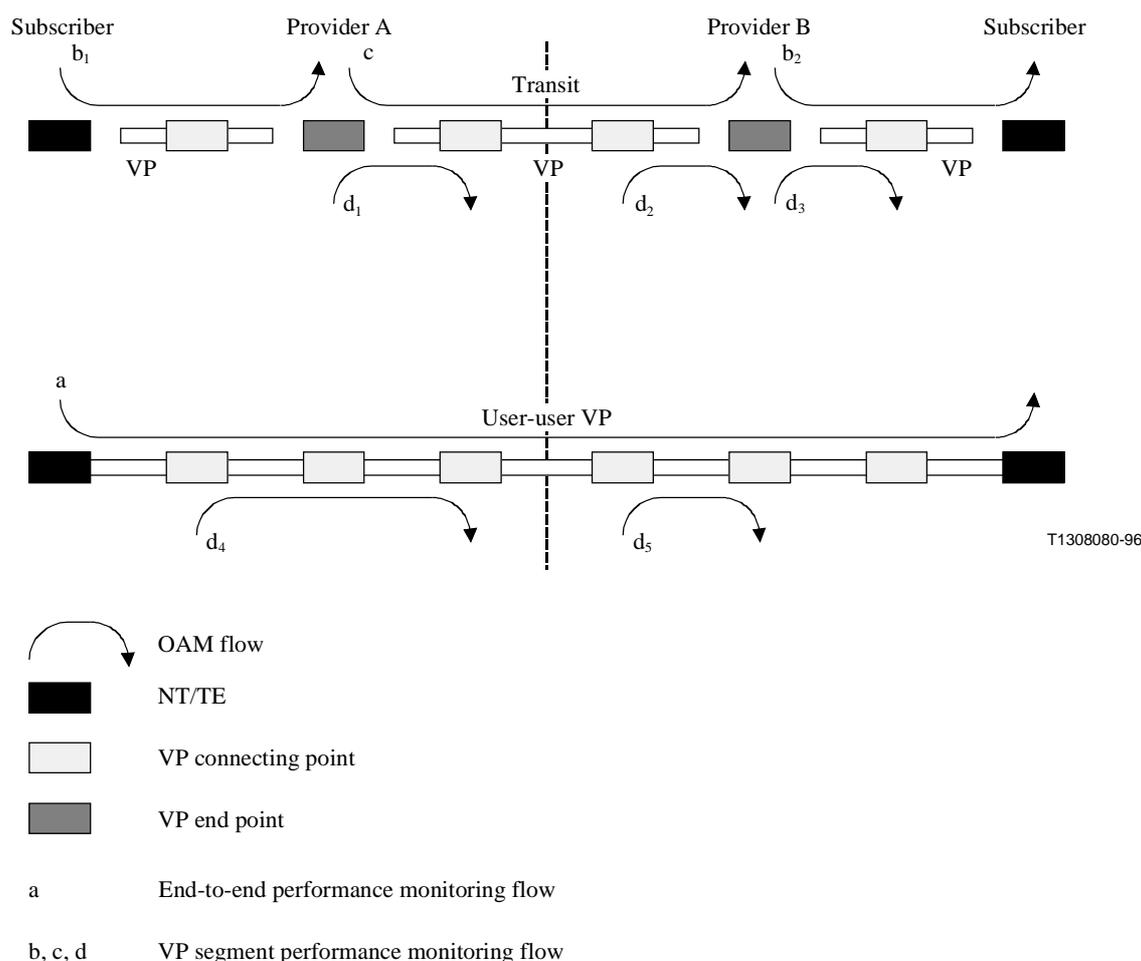
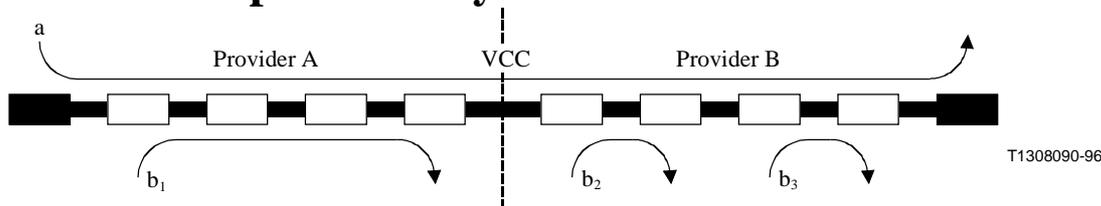


Figure C.1/I.356 – OAM cell flow for VP performance monitoring

Superseded by a more recent version



(An end-to-end performance monitoring flow and a network maintenance flow can be provided at any VC cross-section.)

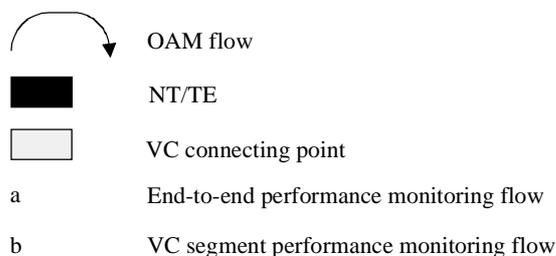


Figure C.2/I.356 – OAM cell flow for VC performance monitoring

C.2 Cell error ratio

Cell error ratio can be measured out-of-service by transferring a known data stream into the network at the source measurement point and comparing the received data stream with the known data stream at the destination MP.

Estimation of cell error ratio by in-service measurement is desirable but difficult. It has been suggested that a BIP16 indicator could be used to estimate the cell error ratio over a block of N cells using the following algorithms:

- If "i" parity violations are observed ($0 \leq i \leq 2$) without any loss of cells, estimate the number of errored cells by i.
- If more than two parity violations are observed without any loss of cells, estimate the number of errored cells by N.

Let N_e be the number of errored cells that are identified using the above procedure. The CER is estimated as the ratio of N_e , the estimated errored cells, over $N't_{0+1}$, the total number of cells that have been transmitted in blocks in which no cell loss or cell misinsertion outcomes have been observed.

The method assumes that the number of cells within a block is small (e.g. less than 200 cells) and that the transmission medium is such that either very few errors are experienced or large bursts of errors occur. The feasibility and accuracy of this and other in-service CER estimation methods are for further study.

C.3 Cell loss ratio

This subclause describes methods for estimating in-service both CLR_0 and CLR_{0+1} . Note that CLR_1 is not considered as a candidate for I.356 performance objectives.

CLR_{0+1} can be estimated by dividing Nl_{0+1} by Nt_{0+1} . CLR_0 can be estimated by dividing Nl_0 by Nt_0 . These methods may undercount cell loss events if cell misinsertions occur during the measurement period.

Superseded by a more recent version

An adjustment to the in-service estimate of CLR_0 (CLR_{0+1}) should be made whenever non-conforming cells are tagged or discarded by the UPC or NPC. These adjustments can be made using discarded and tagged cell counts taken directly from the UPC/NPC and using the modified CLR definitions of 7.3. However, it cannot be expected that the UPC/NPC will be an exact implementation of the conformance definition (see Appendix I). In particular, the number of cells discarded in the UPC/NPC may be different from the number of cells found non-conforming by the maximally stringent test of conformance method applied at the UNI/NNI. Therefore, using UPC/NPC discarded cell counts instead of the number of non-conforming cells will often be inaccurate. Furthermore, the relationship between the performance of each connection portion and the end-to-end performance is not straightforward when cells are identified as non-conforming at either the UNI or at an intermediate NNI. For example, when some cells are identified as conforming by the UPC, but are identified as non-conforming by the NPC, although both connection portions may deliver the committed QOS, the end-to-end QOS may not be delivered (due to apparent non-conformance at the NNI).

C.4 Cell misinsertion rate

Cell misinsertion rate can be estimated in-service by dividing Nm_{0+1} by the duration of the observation period. These methods may undercount cell misinsertion events if cell loss occurs during the measurement period.

A more accurate out-of-service method of estimating cell misinsertion rate is to maintain a VP or VC for a known period of time but transmit no cells on it. Any cells received on the connection are then misinserted cells, and the cell misinsertion rate can be estimated by dividing the number of received cells by the observation time. The likelihood of observing misinserted cells can be increased by increasing the number of idle connections, at a cost of reduced network efficiency.

C.5 Severely errored cell block ratio

Severely errored cell block outcomes can be estimated in-service by computing the number of lost cell or misinserted cell outcomes in each cell block depending on whether a CLR_{0+1} objective is specified.

When an in-sequence FM cell arrives, the number of cells (respectively, the number of $CLP = 0$ cells) transmitted in the corresponding block is nt_{0+1} (respectively, nt_0). These numbers can be compared to the number of received cells nr_{0+1} (respectively, the number of received $CLP = 0$ cells nr_0) in the monitored block. More precisely:

- a) if there is a CLR objective specified for the aggregate $CLP = 0 + 1$ cell stream, a cell block is estimated as severely errored if the absolute value of the number of transmitted cells minus the number of received $CLP = 0 + 1$ cells $|nt_{0+1} - nr_{0+1}|$ is greater than M ;
- b) if the CLR objective for the aggregate $CLP = 0 + 1$ cell stream is U , but there is a CLR objective for the $CLP = 0$ cell stream, a cell block is estimated as severely errored if either the value of the number of transmitted $CLP = 0$ cells minus the number of received $CLP = 0$ cells ($nt_0 - nr_0$) is greater than M , or if the value of the number of received $CLP = 0 + 1$ cells minus the number of transmitted $CLP = 0 + 1$ cells ($nr_{0+1} - nt_{0+1}$) is greater than M ;
- c) if there is no CLR commitment for either the aggregate or the $CLP = 0$ cell stream, there are no QOS commitments for the SECBR. In these cases, network providers may still be interested in evaluating their network's SECBR performance and method a) is suggested.

Severely errored cell block ratio can be estimated in-service for a set of S consecutive or non-consecutive cell blocks by dividing the total number of severely errored cell blocks, as defined

Superseded by a more recent version

above, by S. This in-service measurement method will undercount severely errored cell blocks to some degree, since it does not include delivered errored cells in the estimation of M. A more accurate estimate of severely errored cell block ratio can be obtained by comparing transmitted and received data in an out-of-service measurement.

C.6 Cell transfer delay

Cell transfer delay can be measured in-service by transmitting time-stamped OAM cells through the network on an established connection. The transmitted OAM cell payload contains the time t_1 at which the cell was transmitted. The receiver subtracts t_1 from the time t_2 at which the cell is received to determine the cell transfer delay for that cell. The method requires synchronized clocks at the two MPs or a suitable loopback mechanism at the receiver.

Individual cell transfer delay observations may be combined to calculate statistics of the cell transfer delay distribution. Such statistics also characterize 2-point cell delay variation. The use of OAM cell measurements to develop cell transfer delay and 2-point CDV distributions is possible but may be limited by the OAM cell transmission frequency. This topic is for further study.

C.7 Cell delay variation

Figure C.3 provides a method of estimating the range of the 2-point CDV distribution (or equivalently, the range of the absolute cell transfer delay distribution) for a succession of transferred cells on the basis of observations of 1-point CDV values (y_k). The method assumes that cells are input uniformly at the peak cell rate and is applicable only to connections providing CBR service. At time a_k , when cell k is observed at the measurement point, the value of the 1-point CDV parameter $y_k = c_k - a_k$ is computed to obtain the current value of Q_k (the observed range of cell transfer delays). Then,

- if the y_k is non-negative, the next cell reference time c_{k+1} is computed and the value of Q_k is computed taking into account the observed positive difference between the theoretical emission time and the actual arrival times;
- if y_k is negative, cell k is considered "late" compared to the theoretical time. The next cell reference time c_{k+1} is computed and the value for Q_k is computed taking into account the computed values for Q_{k-1} and y_k .

This method does not provide correct results when cell loss or misinsertion occurs. Methods capable of handling such outcomes are for further study. One such method would count the number of lost or misinserted cells and shift the expected arrival times for subsequent cells accordingly.

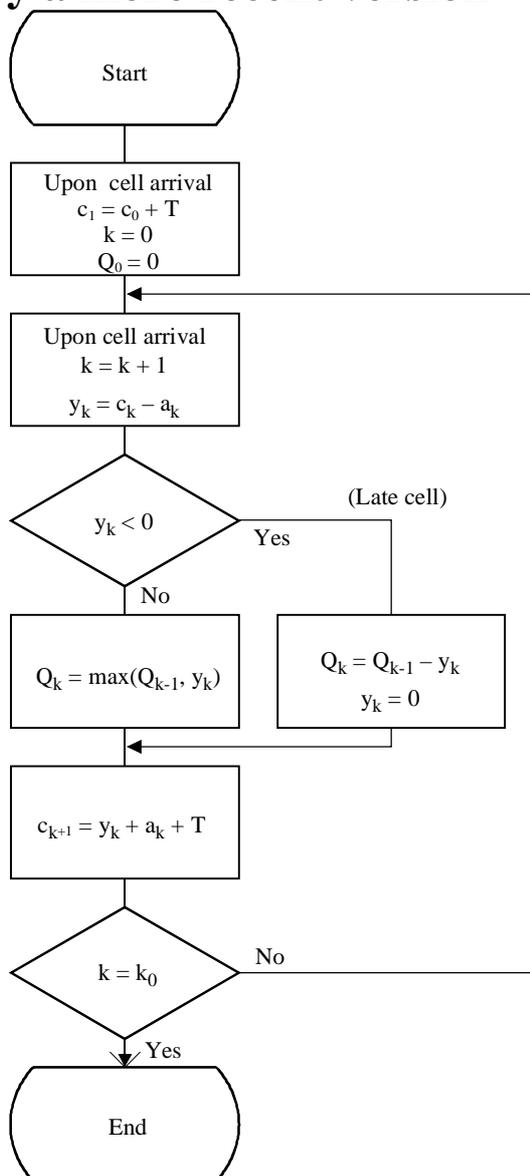
The method described above does not provide an estimate of the quantiles of the cell transfer delay distribution. Such quantiles could be estimated by measuring the 2-point CDV distribution. A more complete measurement process could be elaborated based on the process described here.

When the modified reference arrival pattern $\{c'_k\}$ is defined as follows:

$$c''_0 = a_0 = 0$$
$$c''_{k+1} = c''_k + T$$

and no lost or misinserted cell outcomes occur in the measured cell stream, the distribution of the values of $y''_k = c''_k - a_k$ can be used to estimate 2-point CDV distribution quantiles.

Superseded by a more recent version



T1308100-96

- c_k Reference arrival time for cell k at MP
- a_k Actual arrival time for call k at MP
- y_k 1-point CDV
- Q_k Observed range of cell transfer delay in the set of cells up to cell k

$$c'_{k+1} = \begin{cases} y'_k + a_k + T = a_k + T & \text{if } y'_k < 0 \\ y'_k + a_k + T = c'_k + T & \text{if } 0 \leq y'_k \leq \tau \end{cases} \quad \text{upon cell arrival}$$

Figure C.3/I.356 – Estimation of the range of 2-point CDV from 1-point CDV for connections providing CBR service

C.8 Estimation of CLR and SECBR in case of lost FM OAM cells

This subclause describes an algorithm that allows estimation of the delivered performance, even when one or several forward monitoring (FM) OAM cells are lost.

Superseded by a more recent version

A simple algorithm described below estimates the number of lost cells or misinserted cells for each received OAM FM cell, even when one or several FM cells are lost. This algorithm is based solely on the counts of CLP = 0 + 1 cells (and makes use of the TUC₀₊₁ field in OAM FM cells) although the generalization is straightforward. This algorithm uses the following variables:

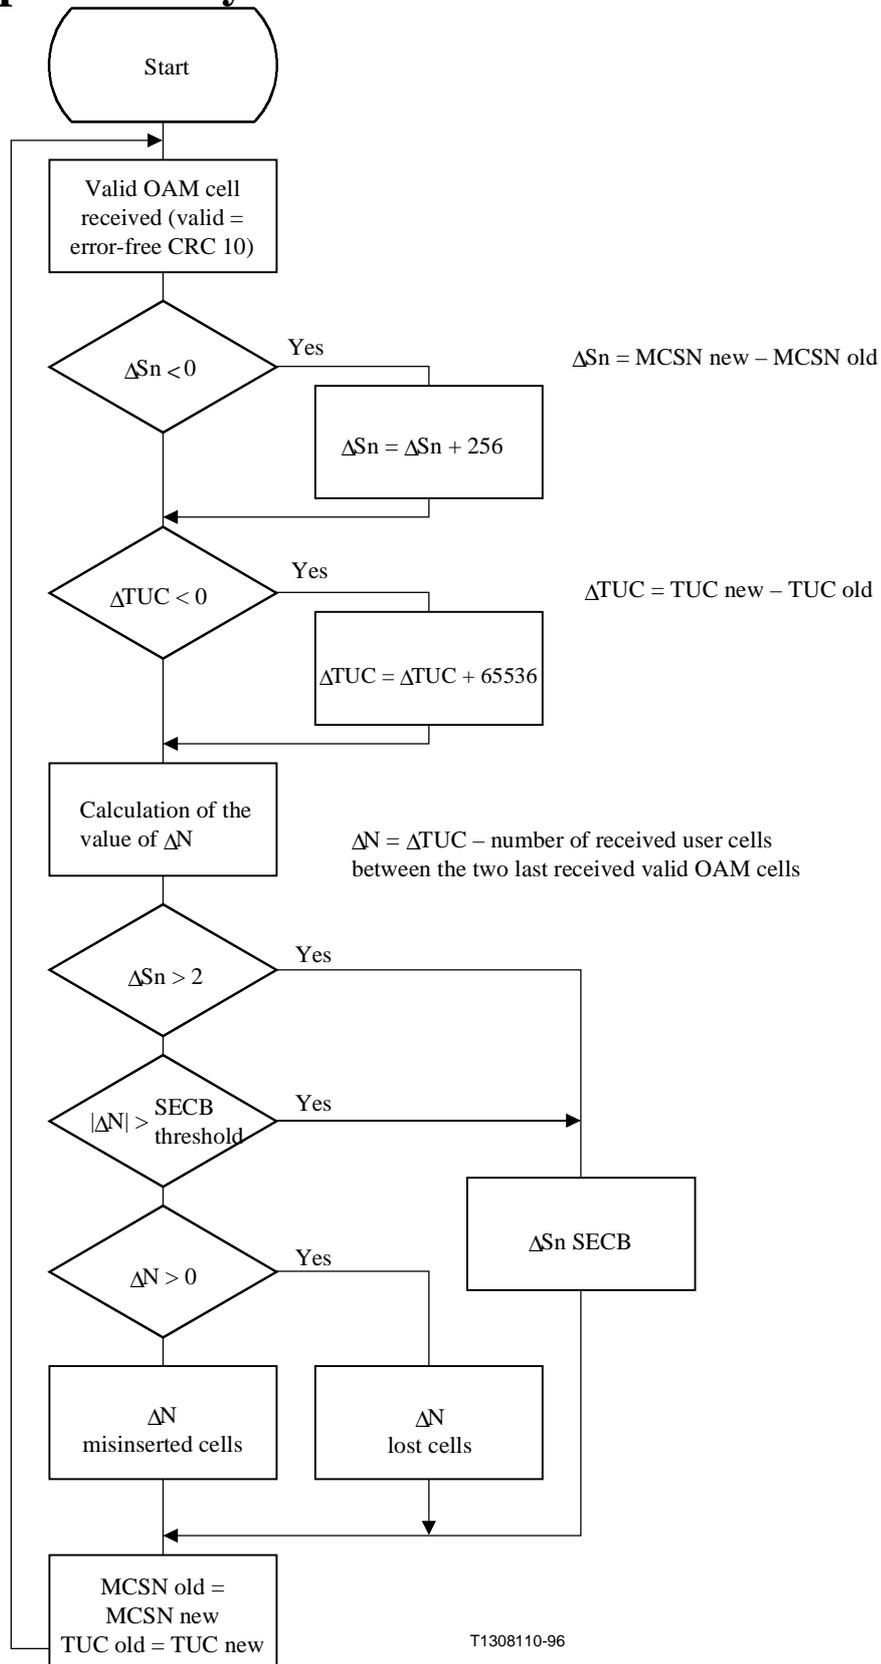
- ΔS_n represents the difference between the previously received MCSN and the present MCSN (MCSN is a field in OAM FM cells defined as the monitoring cell sequence number);
- ΔTUC is the difference between the previously received TUC₀₊₁ and the present TUC₀₊₁. It represents the number (nt) of user monitored cells that have been transmitted between the last two OAM FM cells that have been received. ΔN represents the difference between the numbers of received (nr) and the number of transmitted cells (nt), i.e. ($\Delta N = nt - nr$).

The principle of the algorithm is to discriminate isolated or small groups of performance impairments (that should lead to increased CLR and CMR) from large groups of performance impairments (that should lead to increase the number of SECBs). This translates as follows:

- if ΔS_n equals 1, no FM OAM cell loss is identified and the ordinary estimation process is used;
- if ΔS_n equals 2, one FM OAM is considered as lost. If the absolute value of ΔN is smaller than the SECB threshold, either the cell loss counter or the misinserted cell counter is incremented by this absolute value (depending on the sign of ΔN), and the SECB counter is not modified. Conversely, if the absolute value of ΔN is larger than M, neither the cell loss counter nor the misinserted cell counter are modified, but the SECB counter is increased by 2;
- if ΔS_n is larger than 2, the SECB counter is increased by ΔS_n .

The flowchart for the algorithm is given in Figure C.4.

Superseded by a more recent version



NOTE 1 – Initialization phase is not considered in this algorithm.

NOTE 2 – MCSN new and TUC new fields are related to the last received valid OAM cell.

NOTE 3 – MCSN old and TUC old fields are related to the previously received valid OAM cell.

Figure C.4/I.356 – Estimation of cell transfer performance parameters in case of lost OAM cells

Superseded by a more recent version

NOTE – In this algorithm the SECB threshold is a "dynamic" value that depends on the actual in-service PM OAM block size. This value is obtained using the following formula:

$$\Delta TUC / (\Delta S_n \times 32)$$

which applies only when $\Delta S_n = 1$ or 2 .

APPENDIX I

Assessing the performance of a UPC/NPC mechanism

This Appendix addresses the assessment of the performance of UPC/NPC mechanisms. This information is provided solely to assist network providers meet the CLR objective allocated to their portion. A network portion is considered compliant with the performance recommendations of I.356 if it meets its allocated performance objectives, regardless of whether its UPC/NPC mechanisms comply with these suggestions.

As specified in Recommendation I.371, assessment of UPC/NPC performance is done by comparing the behaviour of the UPC/NPC with the ideal UPC/NPC mechanism as represented by the cell conformance definition.

Both aspects of the UPC/NPC mechanism performance have to be considered:

- the UPC/NPC mechanism should never discard/tag more cells than the ideal UPC/NPC mechanism;
- when there are non-conforming cells, the UPC/NPC mechanism should be capable of discarding/tagging a number of cells at least equal to some lower bounds derived from the ideal UPC/NPC mechanism.

Only the first point is considered in this Appendix. The second point requires further study.

For any type of traffic contract, a well-behaved UPC/NPC mechanism should always discard/tag less than or equal to the numbers of cells obtained by the maximally constraining tests defined in 7.2.

APPENDIX II

Hypothetical reference connections for validating the ATM performance objectives

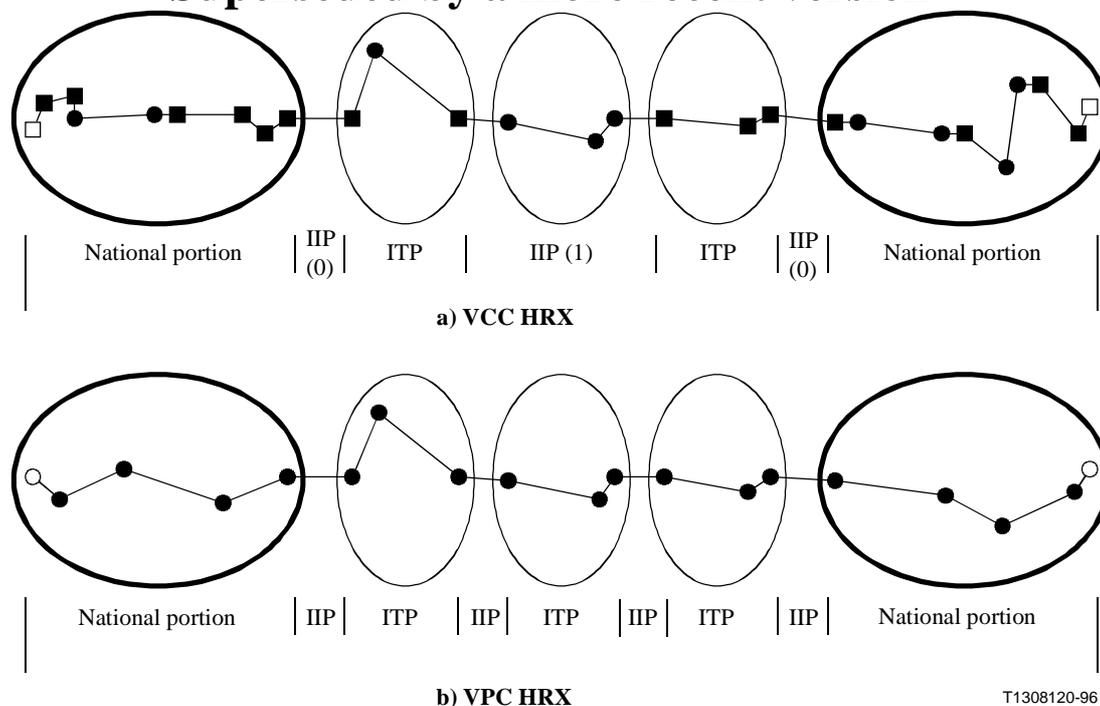
This Appendix presents the hypothetical reference connections considered in validating the feasibility of the end-to-end performance objectives presented in clause 8. These hypothetical reference connections are examples only. The material in this Appendix is not normative and does not recommend or advocate any particular connection architectures. Any connection that satisfies the performance objectives of clause 8 or connection portion that satisfies the allocation rules of clause 9 can be considered fully compliant with the normative recommendations of I.356.

The following complex, but realistic hypothetical reference connections (HRXs) were considered in evaluating whether the end-to-end objectives of clause 8 could be supported in the year 2001.

- a) VCC.
- b) VPC.

These HRXs are illustrated in Figure II.1.

Superseded by a more recent version



- VC source ■ VC switch
- VP source ● VP switch

Figure II.1/I.356 – 27 500 km hypothetical reference connections

Both international HRXs (VCC and VPC) include two national portions and one international portion (IP). The international portions each include ATM (VP and/or VC) switching or cross-connect functions in each of three transit countries. The HRXs are 27 500 km long, as in Recommendation G.826.

After analysis, it was agreed that the end-to-end objectives and QOS classes proposed in Table 1 could be achieved on the HRXs.

II.1 Number ATM nodes in the HRXs

By definition, an international interoperator portion (IIP) for a VP connection does not contain any ATM nodes. An IIP for a VC connection may contain several ATM nodes that access the VP layer. Let IIP(i) denote an IIP that spans i countries all accessing the VP layer.

Table II.1 indicates the number of ATM (VP and VC) nodes that are crossed in the standardized connection portions of the two HRXs.

Table II.1/I.356 – Number of ATM nodes (VP or VC nodes) in each portion of 2 HRXs

	National portion	IIP(0)	IIP(1)	IIP(2)	IIP(3)	ITP
a) VCC	8	0	3	6	9	3
b) VPC	4	0	Not applicable	Not applicable	Not applicable	3

Superseded by a more recent version

Based on Table II.1, Table II.2 calculates the total number of ATM nodes crossed by both of the HRXs.

Table II.2/I.356 – Total number of nodes in each HRX

	a) VCC	b) VPC
Number of nodes	$25 = 8 + (3 \times 3) + 8$	$17 = 4 + (3 \times 3) + 4$

II.2 Switching speeds in the HRXs

Two types of nodes were considered in the HRXs:

- nodes with output links at a rate equal to 34 or 45 Mbit/s;
- nodes with output links at least equal to STM1 (155 Mbit/s).

Table II.3 lists the number of ATM nodes assumed to operate at 34/45 Mbit/s for each standardized connection portion of the 4 HRXs. The remaining links were assumed to operate at 150 Mbit/s or more.

Table II.3/I.356 – Maximum number of ATM nodes at 34/45 Mbit/s in each portion

	National portion	IIP(0)	IIP(1)	IIP(2)	IIP(3)	ITP
a) VCC	3	0	1	2	3	1
b) VPC	2	0	Not applicable	Not applicable	Not applicable	1

NOTE – In the near future, it is likely that many ATM connections will have access links at rates lower than 34/45 Mbit/s. Two particular cases were considered:

- the ingress access link rate is below 34/45 Mbit/s but output rates are at least at 34/45 Mbit/s: in this case, no supplementary CDV degradation is to be expected;
- the egress link rate is below 34/45 Mbit/s: in this case, a supplementary CDV performance degradation may be expected beyond the end-to-end objective presented in Table 2. For the considered HRXs, the egress link rate was assumed to be at or above 34/45 Mbit/s.

II.3 Loading within the HRXs

The fraction of each transmission link occupied by active cells was assumed to be 0.85 for both the VCC and the VPC HRXs. It is unlikely that a network will operate continuously at such high loads, especially on access links. For the CDV analyses the load on each link was assumed to vary between 0 and 0.85.

II.4 Geostationary satellites within the HRXs

The use of geostationary satellites was considered during the study of the HRXs. A single geostationary satellite can be used within the HRXs and still achieve the end-to-end objectives of clause 6 on the assumption that it replaces significant terrestrial distance, multiple ATM nodes, and/or transit country portions.

The use of low- and medium- Earth orbit satellites was not considered in connection with these HRXs.

Superseded by a more recent version

When choosing the allocations to be given to portions containing geostationary satellite hops, it was assumed that the satellite would replace significant terrestrial distance and eliminate the need for a number of ATM nodes. Table II.4 presents the HRXs that were used in preparing the allocations for SECBR, CER, and CLR. These HRX designs are not normative.

Table II.4/I.356 – Hypothetical portions with a geostationary satellite

Portion type	Geostationary satellites	Terrestrial distance	ATM nodes (VC or VP)
NP	1	500 km	2 or 3
IIP(0)	1	< 100 km	0
ITP	1	< 100 km	2
IIP(1)	1	1000 km	1 or 2
IIP(2)	1	2500 km	2 or 3
IIP(3)	1	5000 km	4 to 6

II.5 Other aspects of the HRXs

- Each of the HRXs has a ratio of route-to-air distance based on Recommendation G.826.
- The error performance of all transmission facilities is consistent with Recommendation G.826.

NOTE – The need for additional transmission performance parameters and objectives is under study.

- CTD due to terrestrial transmission and physical layer processing is 6.25 microseconds per km.
- Each ATM node (VC or VP) creates a worst-case average of 300 microseconds of queuing delay for QOS class 1.
- Private networks and CEQ are not included.

APPENDIX III

Example applications of the allocation rules of 9.5, 9.6 and 9.7

The following examples illustrate the use of the allocation rules in 9.5, 9.6 and 9.7.

Example 1

An international connection consisting of:

- one NP with air-route distance between its MPT and MPI equal to 1000 km; the calculated route length is 1500 km;
- one IIP(0) with air-route distance between its MPIs equal to 500 km; the calculated route length is 750 km;
- one NP with air-route distance between its MPI and MPT equal to 1000 km; the calculated route length is 1500 km.

The SECBR and CER objective is thus $2 \times (17.5 + 3) + (1 + 2) = 44\%$ of the end-to-end SECBR and CER objectives.

Superseded by a more recent version

The class 1 CLR objective is thus $2 \times (23 + 2) + (1 + 1) = 52\%$ of the end-to-end class 1 CLR objective.

The class 2 CLR₀₊₁ and the class 3 CLR₀ objectives are thus $2 \times 34.5 + 1 = 70\%$ of the end-to-end class 2 CLR₀₊₁ and the class 3 CLR₀ objectives.

Example 2

An international connection consisting of:

- one NP with air-route distance between its MPT and MPI equal to 1000 km; the calculated route length is 1500 km;
- one IIP(3) with air-route distance between its MPIs equal to 5000 km; the calculated route length is 6250 km;
- one NP with air-route distance between its MPI and MPT equal to 1000 km; the calculated route length is 1500 km.

The SECBR and CER objective is thus $2 \times (17.5 + 3) + (10 + 13) = 64\%$ of the end-to-end SECBR and CER objectives.

The class 1 CLR objective is thus $2 \times (23 + 2) + (25 + 7) = 82\%$ of the end-to-end class 1 CLR objective.

The class 2 CLR₀₊₁ and the class 3 CLR₀ objectives are thus $2 \times 34.5 + 31 = 100\%$ of the end-to-end class 2 CLR₀₊₁ and the class 3 CLR₀ objectives.

Example 3

An international connection consisting of:

- one NP with air-route distance between its MPT and MPI equal to 2000 km; the calculated route length is 2500 km;
- one IIP(0) with a geostationary satellite hop;
- one NP with air-route distance between its MPI and MPT equal to 500 km; the calculated route length is 750 km.

The SECBR and CER objective is thus $(17.5 + 5) + 35 + (17.5 + 2) = 77\%$ of the end-to-end SECBR and CER objectives.

The class 1 CLR objective is thus $(23 + 3) + 25 + (23 + 1) = 75\%$ of the end-to-end class 1 CLR objective.

The class 2 CLR₀₊₁ and the class 3 CLR₀ objectives are thus $2 \times 34.5 + 1 = 70\%$ of the end-to-end class 2 CLR₀₊₁ and the class 2 CLR₀ objectives.

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