INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

I.356

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

(11/93)

INTEGRATED SERVICES DIGITAL NETWORK (ISDN)

OVERALL NETWORK ASPECTS AND FUNCTIONS

B-ISDN ATM LAYER CELL TRANSFER PERFORMANCE

ITU-T Recommendation I.356

(Previously “CCITT Recommendation”)

FOREWORD

The ITU-T (Telecommunication Standardization Sector) is a permanent organ of the International Telecommunication Union (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 (Helsinki, March 1-12, 1993).

ITU-T Recommendation I.356 was prepared by ITU-T Study Group 13 (1993-1996) and was approved under the WTSC Resolution No. 1 procedure on the 26th November 1993.

NOTE

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

CONTENTS

Recommendation I.356

B-ISDN ATM LAYER CELL TRANSFER PERFORMANCE

(Geneva, 1993)

1 Introduction

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. A provisional list of generic B-ISDN performance parameters is provided in Recommendation I.350.

NOTES

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. It is intended that one or more ATM cell transfer performance objectives will be specified for each of the defined parameters.

2 The defined parameters apply to cell streams in which all cells conform with a negotiated Recommendation I.371 traffic contract. It is intended that parameter definitions and measurement methods applicable to cell streams in which some cells do not conform with such a contract will be developed. Appendix I contains material relevant to this problem and it is recognized that further work is needed in this area.

3 The defined parameters are intended to characterize ATM connections in the available state. Availability decision parameters and associated availability parameters and their objectives will be the subject of a separate Recommendation.

2 Performance model

Recommendation I.353 defines measurement points (MPs) and associated reference events that provide a basis for ISDN performance description. ATM cell transfer performance is measured by observing the reference events created as ATM cells cross MPs. The MPs should be located at interfaces at which the ATM layer is accessible. For broadband ISDN, the location of the MPI (measurement point international) is for further study.

Figure 1 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 a virtual channel or virtual path is switched by equipment using the ATM technique, and thus 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.

– 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.

In the context of Recommendation I.353 and this Recommendation:

- a cell exit event occurs when the first bit of an ATM cell has completed transmission across an MPT out of a TE, or across an MPI out of an SSN;

- a cell entry event occurs when the last bit of an ATM cell has completed transmission across an MPT into a TE, or across an MPI into an SSN.

Unassigned cells do not create cell transfer reference events.

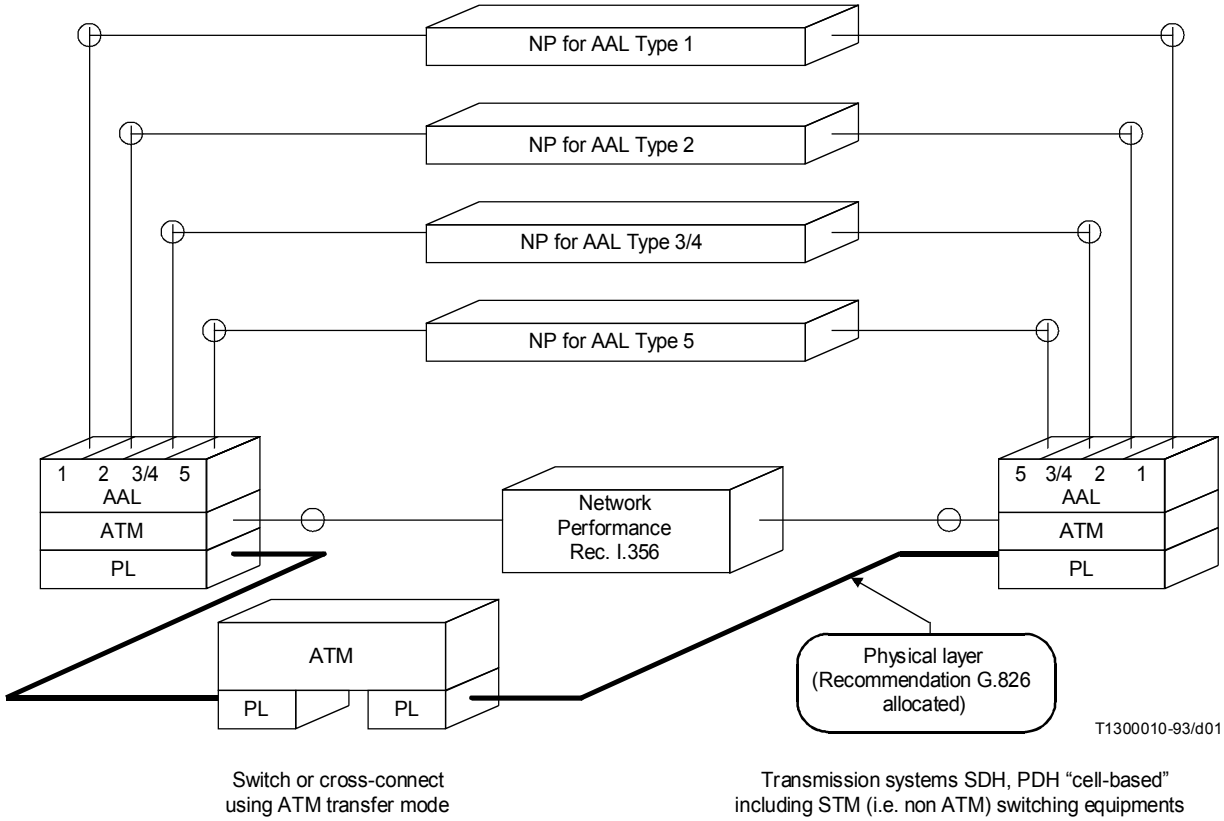

FIGURE 1/I.356

Layered model of performance for B-ISDN

3 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 events are said to be corresponding if they occur on a predefined channel and at a pair of predefined boundaries.

By considering two cell transfer reference events, CRE_1 and CRE_2 at MP_1 and MP_2 respectively, a number of possible cell transfer outcomes may be defined. A transmitted cell is either successfully transferred, errored, or lost. A received cell for which no corresponding transmitted cell exists is said to be misinserted. Cell misinsertion can occur as a result of errors in the cell header. Figure 2 illustrates the cell transfer outcome definitions.

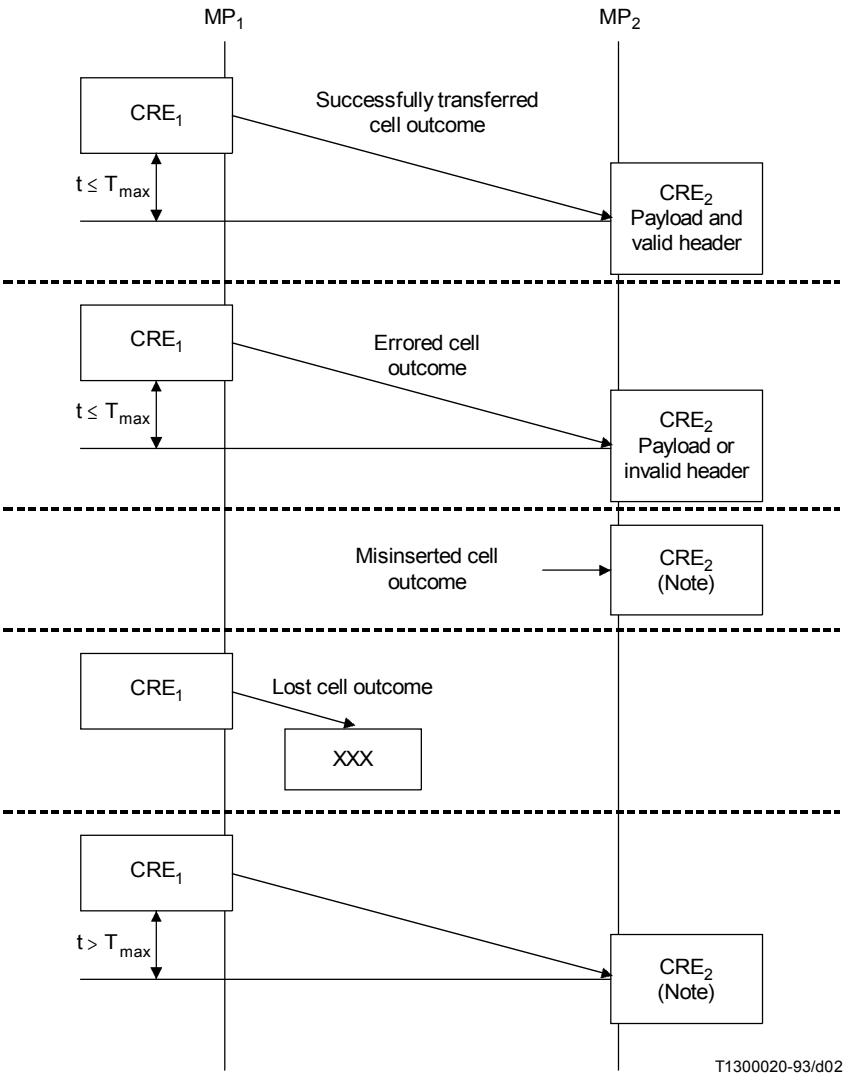

NOTE – Outcome occurs independent of cell content.

FIGURE 2/I.356

Cell transfer outcomes

3.1 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.

3.2 Errored cell outcome

An errored cell outcome occurs when a CRE_2 corresponding to CRE_1 happens within a specified time T_{max} of CRE_1, 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.[1)]

3.3 Lost cell outcome

A lost cell outcome occurs when a CRE_2 fails to happen within time T_{max} of the corresponding CRE_1.

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.

3.4 Misinserted cell outcome

A misinserted cell outcome occurs when a CRE_2 happens without a corresponding CRE_1.

3.5 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.

For practical measurement purposes, a cell block will normally correspond to the number of user information cells transmitted between successive OAM cells. The size of a cell block is to be specified.

4 ATM performance parameters

This clause defines a set of ATM cell transfer performance parameters using the cell transfer outcomes defined in clause 3. All parameters may be estimated on the basis of observations at the MPs. Cell transfer performance measurement methods are described in Annex C.

4.1 Cell error ratio

Cell error ratio (CER) is the ratio of total errored cells to total successfully transferred cells plus errored cells in a population of interest. Successfully transferred cells and errored cells contained in cell blocks counted as severely errored cell blocks should be excluded from the population used in calculating cell error ratio. See 4.4.

4.2 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 cell blocks counted as severely errored cell blocks should be excluded from the population used in calculating cell loss ratio. See 4.4.

---

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 CRE_2 occurs. These cell transfer attempts will be classified as lost cell outcomes.

4.3 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 duration (or equivalently, the number of misinserted cells per connection second)[2]. Misinserted cells and time intervals associated with cell blocks counted as severely errored cell blocks should be excluded from the population used in calculating cell misinsertion rate. See 4.4.

4.4 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 preventing bursts of cell transfer failures from inappropriately influencing the observed values for cell error ratio, cell loss ratio, cell misinsertion rate, and the associated availability parameters.

4.5 Cell transfer delay

Cell transfer delay (CTD) is the time, $t_2 - t_1$, between the occurrence of two corresponding successful cell transfer events, CRE_1 at time t_1 and CRE_2 at time t_2, where $t_2 > t_1$ and $t_2 - t_1$ ### T_{max}. The value of T_{max} is for further study.

4.5.1 Mean cell transfer delay

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

4.5.2 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 on the basis of observation of a sequence of consecutive cell arrivals at a single MP. The second parameter, 2-point cell delay variation, is defined on the basis of 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 is relatable to cell conformance at the MP, and to network queues. It is also related to the buffering procedures used in AAL 1 of the receiving side 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 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. Cell misinsertion on a particular connection is most often caused by an undetected error in the header of a cell being transmitted on a different connection. Since the mechanism that most often causes misinserted cells has 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.

4.5.2.1 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 3a)]: $y_k = c_k - a_k$. The reference arrival time pattern (c_k) is defined as follows:

$$c_0 = a_0 = 0,$$

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

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)]

Annex B illustrates one measurement method that calculates, for a cell stream received at an MP, the number of cells that do not conform with a specified peak cell rate at a specified CDV tolerance. It is anticipated that one or more values for maximum CDV tolerance (τ) will be specified.

4.5.2.2 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 the same two MPs [see Figure 3b)]: $v_k = x_k - d_{1,2}$.

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.

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 CDV is identical to the distribution of absolute cell transfer delay for any specified population of transferred cells. It is anticipated that the specification of 2-point CDV objectives will be in terms of upper and lower quantiles. The specified upper and lower quantile values may depend on the negotiated peak cell rate.

Annex C illustrates one method of estimating the range of the 2-point CDV distribution for a succession of transferred cells on the basis of observations of 1-point CDV values (y_k) for connections providing CBR services. Annex B relates an upper quantile of the probability distribution for 2-point cell delay variation to the cell loss ratio.

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

4.6 Cell flow related parameters

The need for network performance parameters describing the actual flow of cells in an ATM connection is for further study. Such parameters will be needed if flow control mechanisms are implemented in ATM networks. One useful parameter could be the (positive) difference between the negotiated peak cell rate and the actual cell transfer rate. The difference between the requested peak cell transfer rate and the negotiated peak cell transfer rate could also be useful. Measures of specific flow control mechanisms may also be of value.

---

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) It is defined for all corresponding cell transfer reference event pairs (CRE_1, CRE_2); cell transfer delay as defined in 4.5 applies only to successful cell transfer outcomes. Variables a_{2k} and a_{1k} are measured with reference to the same reference time in calculating absolute cell transfer delay.

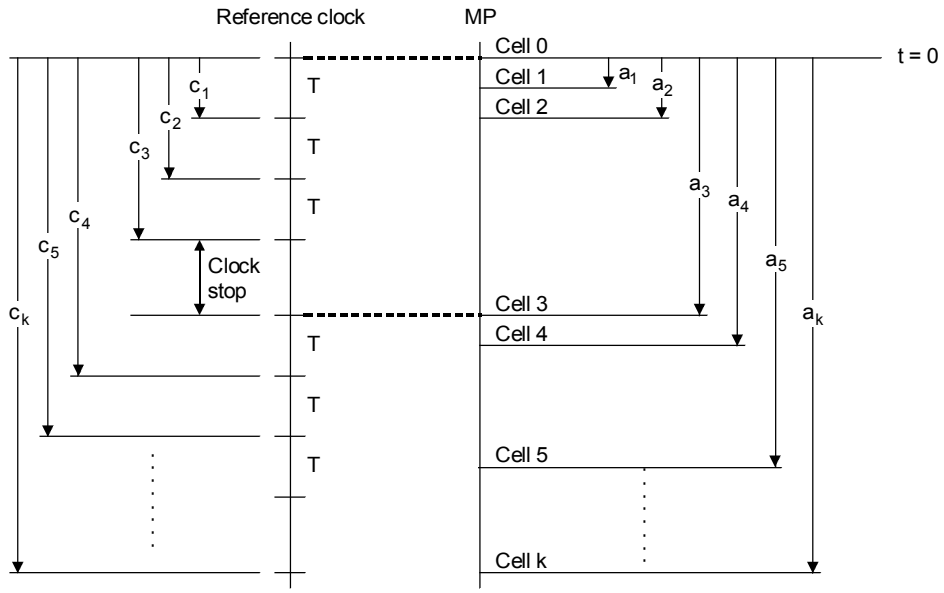

Variables:

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) Call delay variation – 1-point definition

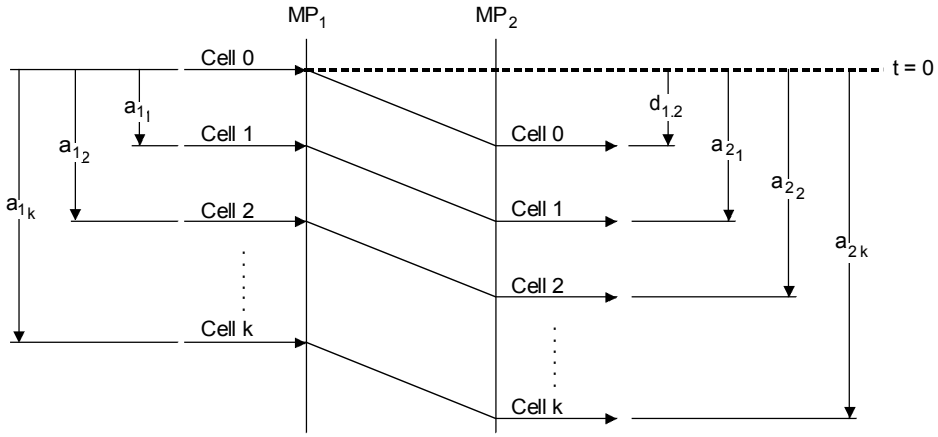

Variables:

a_{1k} Cell k actual arrival time at MP_1
a_{2k} Cell k actual arrival time at MP_2
$d_{1.2}$ Absolute cell O transfer delay between MP_1 and MP_2
x_k Absolute cell k transfer delay between MP_1 and MP_2
v_k 2-point CDV value between MP_1 and MP_2

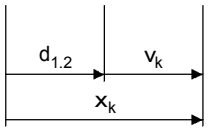

$$x_k = a_{2k} - a_{1k}$$

$$v_k = x_k - d_{1.2}$$

T1300030-93/d03

b) Cell delay variation – 2 point definition

FIGURE 3/I.356

Cell delay variation parameter definitions

Annex A

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

(This annex forms an integral part of this Recommendation)

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 of such dummy AAL SDUs (e.g. all "1", all "0", repeat the previous cell, etc.) require further study (see Recomendation I.363).

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.

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 BER of 0.5) if dummy AAL SDUs are inserted;

– an errored cell 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

(This annex forms an integral part of this Recommendation)

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 (queuing, switching and routing, etc.) delay (see Notes 4 and 5).

NOTES

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

2 The amount of buffering delay consumed in AAL handling equipment will depend on the amount of cell delay variation for which the ATM network is responsible.

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.

4 ATM nodes may perform both virtual channel (VC) and virtual path (VP) switching.

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 4.5.2.1 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 B.1 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]. This number could be divided by the number of cells (k_0) arriving at the MP during an observation period to calculate a cell non-conformance ratio (n/k_0).

Note that the method modifies the values computed for $\{c_k\}$ and $\{y_k\}$ if non-conforming cells are observed. The modified set of variables c'_k and y'_k represent a revised theoretical arrival time and the 1-point CDV, respectively, of the k-th cell for a specific value of CDV tolerance (τ). These variables are obtained as follows:

$$c'_0 = a_0$$

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

$$y'_k = c'_k - a_k$$

Note that $c_k = c'_k$ and $y_k = y'_k$ if only conforming cells are observed (up to cell k).

The method of Figure B.1 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 Annex A/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 two equivalent algorithms is summarized in Table B.1.

TABLE B.1/I.356

Mapping between the variables defined in this Recommendation and those of virtual scheduling and continuous state leaky bucket algorithms defined in Annex A/I.371

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

5) Other methods of calculating non-conforming cell are possible.

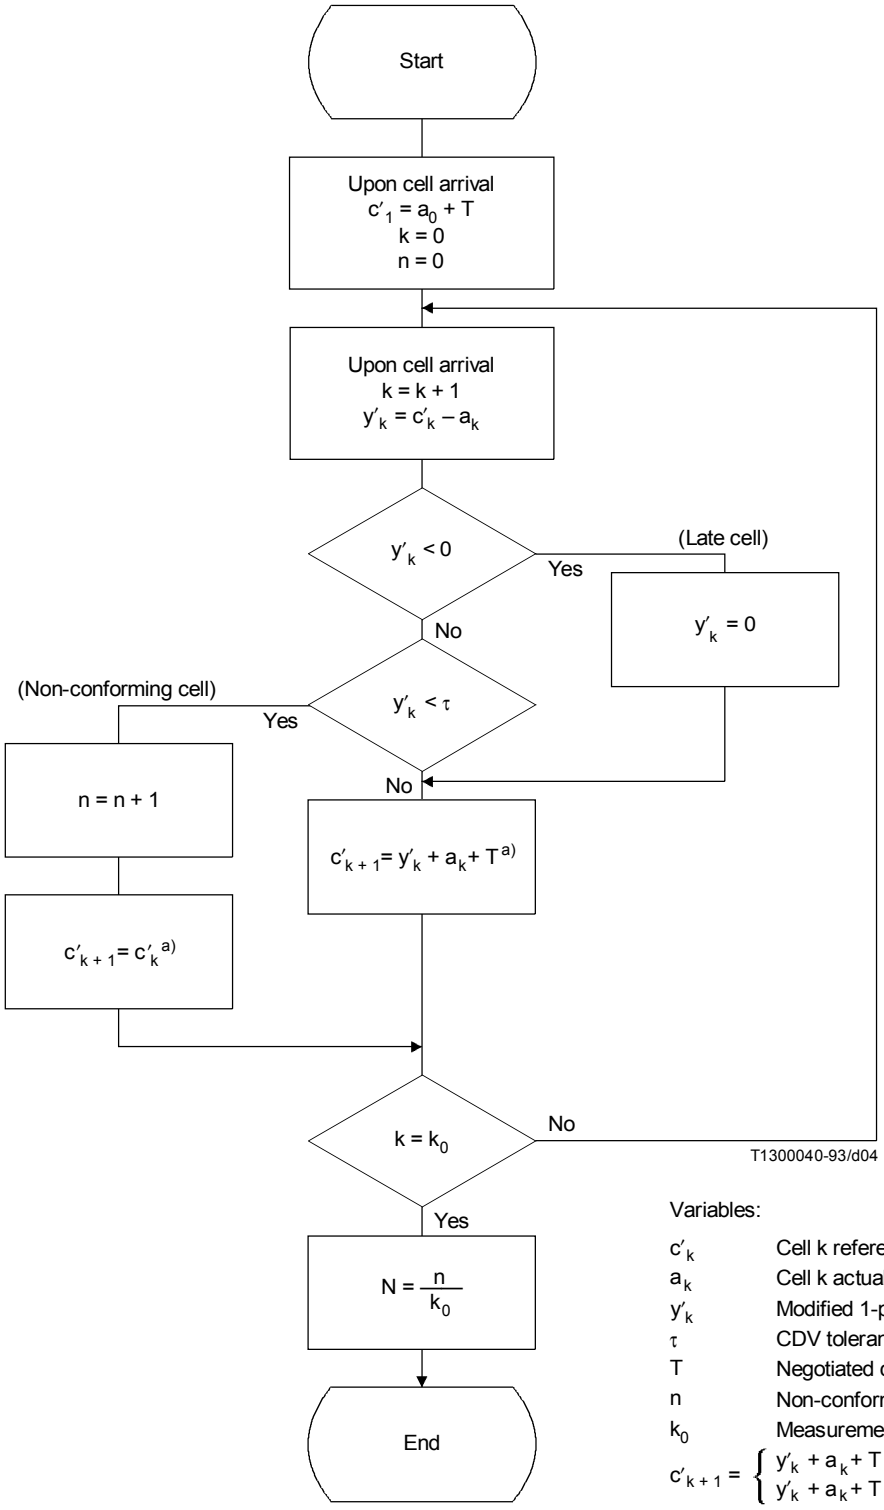

Variables:

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 ($c'_k - a_k$)
τ	CDV tolerance at MP
T	Negotiated cell interarrival time
n	Non-conforming cell count
k_0	Measurement size

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

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

$$LCT_{k+1} = \begin{cases} a_k & \text{if } k^{th} \text{ cell is conforming,} \\ LCT_k & \text{otherwise.} \end{cases}$$

FIGURE B.1/I.356

One method of calculating non-conforming cell total for a given CDV tolerance and peak cell rate

B.4 Relationship between 2-point CDV and cell loss in a 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 $\rho_{offered}$. Any cell arriving at this buffer experiences a random waiting time W before it reaches the link and is transmitted. Figure B.2 illustrates this situation, together with some representative probability density functions for W.

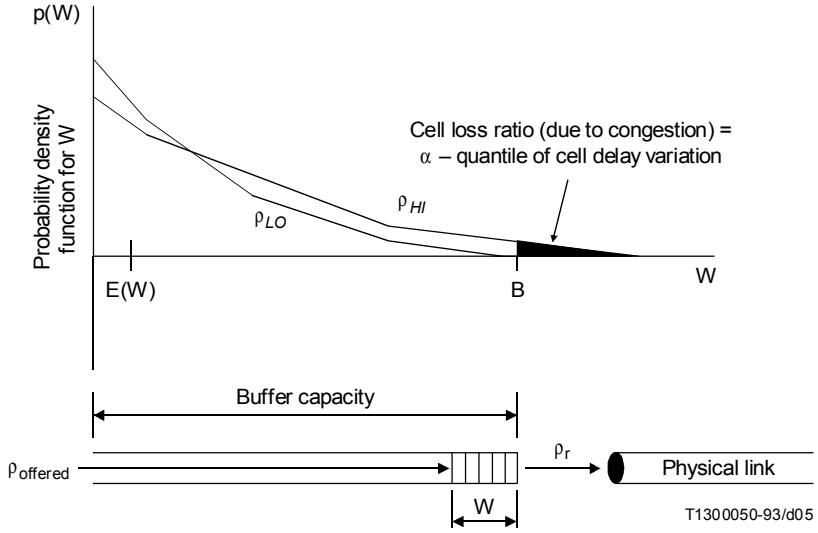

FIGURE B.2/I.356

Illustration of random waiting time (W)

With a sufficiently high value of offered load, characterized in Figure B.2 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 times[6]. The area under this curve can be interpreted as the cell loss ratio (due to congestion) and also as a quantile of cell delay variation.

With a lower value of offered load, characterized in Figure B.2 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}, and in the specification of 2-point CDV and cell loss ratio values.

---

6) 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.

B.5 Allocation of 2-point CDV values

The maximum 2-point CDV $V_{(p, q)}$ between two non-adjacent MPs (p, q) is related to the 2-point CDVs of the portions between those MPs by the inequality

$$V_{(p, q)} \leq \sum_{i=p}^{q-1} V(i, i + 1)$$

This inequality could be useful in allocating end-to-end values for 2-point CDV among connection portions.

Annex C

Cell transfer performance measurement methods

(This annex forms an integral part of this Recommendation)

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 correlation between user and OAM cells. 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 – The accuracy of measurement of these events may only be about plus or minus 200 microseconds at SDH interfaces if the event times for cells embedded in SDH frames are approximated by the frame event times.

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 copied or extracted at any similar point downstream. The corresponding approach for out-of-service monitoring is to establish a virtual path or connection at an appropriate measurement point, introduce a cell stream of known content and timing at that point, and then observe the cell stream at a remote measurement point.

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. 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.

C.1 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.

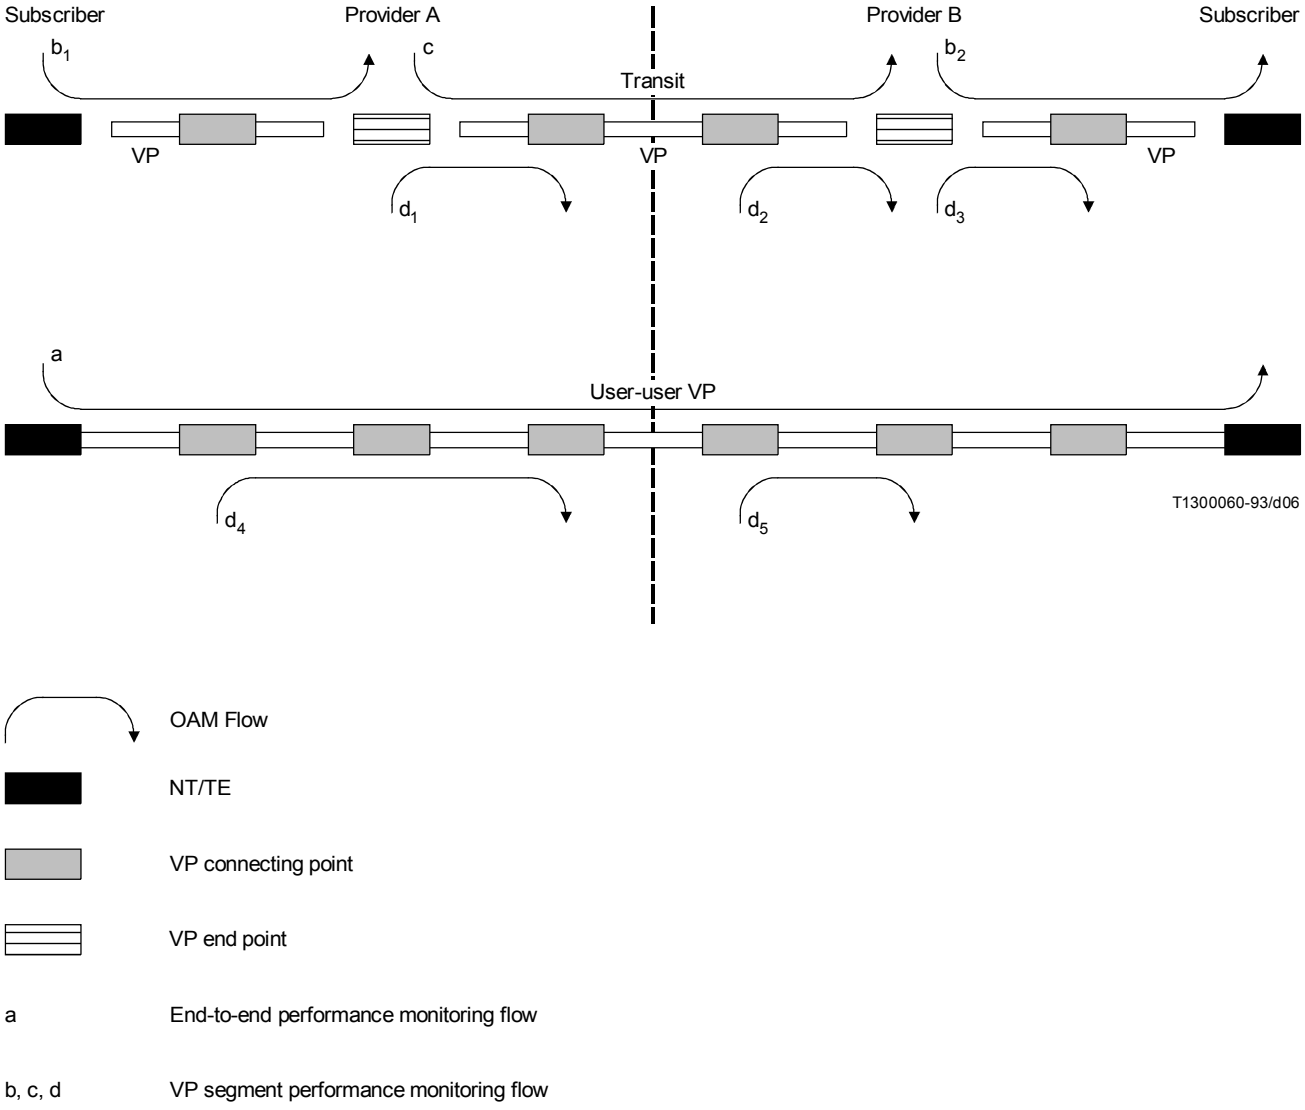

FIGURE C.1/I.356

OAM cell flow for VP performance monitoring

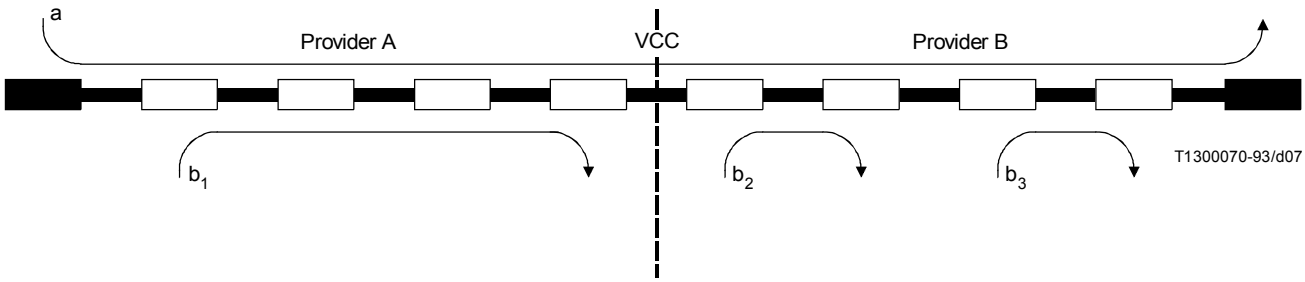

(An end-to-end performance monitoring flow and a network maintenance flow can be provided at any VC cross section).

OAM flow

NT/TE

VC connecting point

a End-to-end performance monitoring flow

b VC segment performance monitoring flow

FIGURE C.2/I.356

OAM cell flow for VC performance monitoring

The method assumes that the number of cells within a block is not too large (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.2 Cell loss ratio

Cell loss ratio can be estimated in-service as follows. The transmitter inserts OAM cells into a transmitted user information cell stream at suitable intervals. Each OAM cell contains a count of the number of user information cells transmitted since the last OAM cell. The receiver keeps a running count of the number of user information cells transmitted (Nt) and received (Nr), excluding cells in severely errored cell blocks. Cell loss ratio can then be calculated by dividing the positive difference (Nt – Nr) by Nt. The method will undercount cell loss events if cell misinsertion occurs during the measurement period.

C.3 Cell misinsertion rate

Cell misinsertion rate can be estimated in-service using a method similar to that described in C.2. Running counts Nt and Nr are obtained during a timed measurement period Tm (excluding cells in severely erorred cell blocks), and the cell misinsertion rate is calculated by dividing the positive difference (Nr – Nt) by Tm. The method will 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.4 Severely errored cell block ratio

Severely errored cell block ratio can be estimated in-service for a set of S consecutive or non-consecutive cell blocks by computing the number of lost cell or misinserted cell outcomes in each cell block (as described in C.2 and C.3); identifying cell blocks with more than M lost cell or misinserted cell outcomes as severely errored cell blocks; and dividing the total number of such severely errored cell blocks 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.5 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.6 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.

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

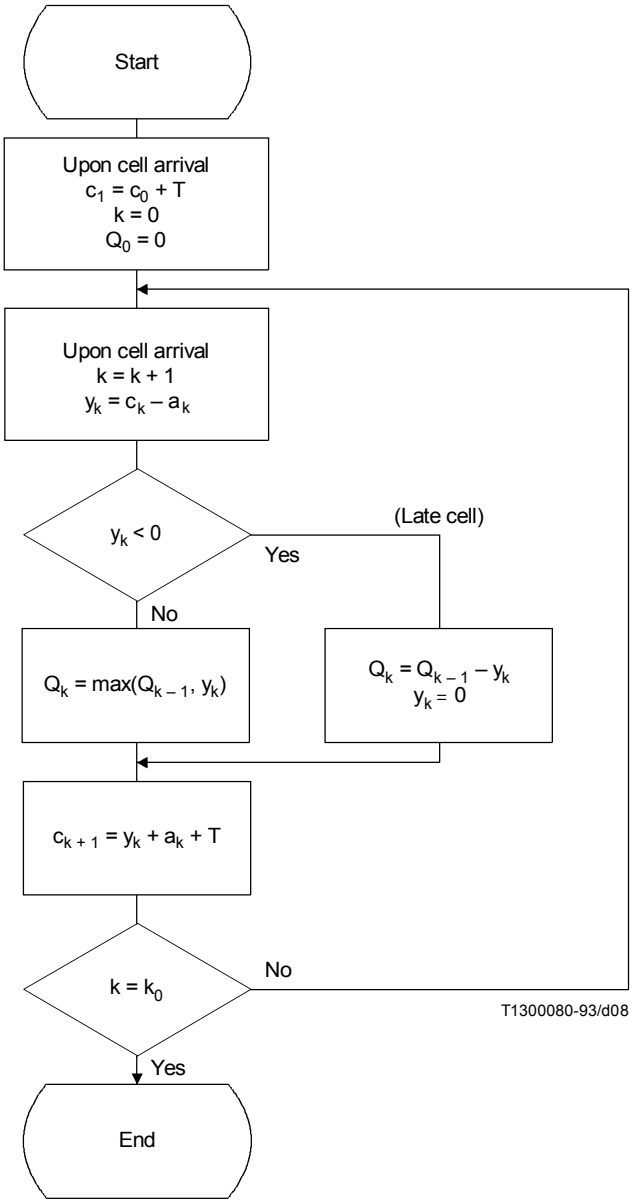

Variables:

c_k Reference arrival time for cell k at MP
a_k Actual arrival time for cell 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

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

FIGURE C.3/I.356

Estimation of the range of 2-point CDV from 1-point CDV for connections providing CBR service

Appendix I

Cell loss performance in the case of non-conformance

(This appendix does not form an integral part of this Recommendation)

I.1 Introduction

The parameters defined in this Recommendation apply to cell streams in which all cells conform with a negotiated Recommendation I.371 traffic contract. This appendix extends the definition of cell loss ratio to cell streams in which some cells do not conform with a negotiated Recommendation I.371 traffic contract (consisting only of a peak cell rate and CDV tolerance) and no separately specified OAM cell flow exists at the MPT or MPI. This appendix assumes that the CLP = 0 traffic descriptor is equal to the CLP = 0 + 1 traffic descriptor. The appendix addresses the cell loss ratio for aggregate cell streams regardless of cell tagging. This appendix discusses the cell loss ratio over the number of conforming cells.

Once a traffic contract is negotiated between a user and network, resources are allocated by the network to the connection on the basis of the negotiated traffic and the characteristics of the UPC/NPC. These resources should allow the network to deliver the negotiated cell loss ratio to the user.

A peak cell rate monitor algorithm, which is defined in Recommendation I.371, classifies the cells which are sent by the user as conforming or non-conforming. When non-conforming cells are sent by the user, the network is allowed to discard a number of cells equal to the number of non-conforming cells. Such discarded cells should not be counted as lost cells in assessing the network's cell loss performance; however, all discarded cells may have an impact from a user perspective.

This appendix does not imply any change in Recommendation I.371. Indeed, allowable network actions in the presence of non-conforming user cells (such as reduced traffic admission or disconnection) may render portions of this material irrelevant.

I.2 Network performance perspective

I.2.1 Modification of the severely errored cell block outcome

As defined in 3.5, 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 excess traffic, a cell block could wrongly be considered as severely errored if the effect of excess traffic is not excluded when assessing whether more than M cells have been lost in the block.

In order to address this specific issue, the following modification to the I.356 SECB definition has been proposed:

– If some cells in the block N consecutive cells are non-conforming with the negotiated traffic contract, the lost cell outcome total should include only cells which are lost in excess of the total number of non-conforming cells in the cell block because only these cells reflect the network behaviour.

The total number of non-conforming cells in the cell block may be computed using e.g. the method which is described in B.2 or estimated by other network functions.

Since traffic control functions may not be synchronized with the process which identifies non-conforming cells, the cells which are considered as non-conforming in a cell block by the measurement process might not be discarded by the traffic control functions. However, an equivalent amount of cells could be discarded, which could belong to a different cell block.

NOTE – It is anticipated that the UPC will take action on the cell stream as soon as the non-conformance is detected by the UPC function.

This situation should be carefully studied in order to evaluate its impact on the assessment of this redefined SECB outcome.

I.2.2 Computing CLR_{NP}

In order to make sure that CLR_{NP} relates to the resources which are allocated by the network to the connection, the set of cells for which CLR_{NP} is assessed should be at most as large as the set of cells which are negotiated in the traffic contract.

It is therefore proposed, in population of interest:

- to compute the total number N_T of cells transmitted in non-SECBs;
- to compute the number N_I of lost cells observed in non-SECBs;
- to compute the number N_{nc} of non-conforming cells observed in non-SECBs (using e.g. the method described in B.2, or estimation methods based on other network functions);
- to derive the quantity CLR_{NP}:

$$CLR_{NP} = \max(0, N_I - N_{nc})/N_T - N_{nc}$$

The value for CLR_{NP} relates directly to the way the network meets the negotiated loss performance: indeed, the quantity $(N_T - N_{nc})$ is at most equal to the amount of negotiated traffic which could be observed in the population of interest and $\max(0, N_I - N_{nc})$ represents the amount of cells which are lost by the network in excess of the non-conforming cells i.e. cell losses which, according to Recommendation I.371, have to be considered as part of the overall network degradation.

Note that the definition for CLR_{NP} reduces to the definition in 4.2 for CLR if all cells are conforming to the traffic contract.

I.2.3 Measurement of CLR_{NP}

To measure CLR_{NP}, information should be collected both at the input and at the output of the network portion under consideration.

The amount of cells transmitted in a cell block, the number of non-conforming cells in this block are to be measured at the input of the network portion, whereas the number of successfully transferred cells is to be measured at the output of the network portion.

When, for each cell block, the numbers of transmitted, successfully transferred, lost and non-conforming cells are available, it is possible:

- for each observed cell block, to assess whether the cell block is severely errored or not;
- to compute the value for CLR_{NP}.

This procedure would allow a network operator or a user to assess CLR_{NP}. The procedure could be implemented in hardware or software as part of stand alone equipment that could be installed either permanently or temporarily, by either the network provider or the customer, at or near the appropriate pair of measurement points. Practical ways of implementing the methods identified in this appendix are for further study.

I.3 User QOS perspective

From a user perspective, the cell loss ratio parameter defined in 4.2 can be applied directly to cell streams in which some cells do not conform with a negotiated traffic contract. This parameter will characterize one aspect of the user-perceived quality of service.

Printed in Switzerland

Geneva, 1994