



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

E.736

(05/97)

SERIES E: OVERALL NETWORK OPERATION,
TELEPHONE SERVICE, SERVICE OPERATION AND
HUMAN FACTORS

Quality of service, network management and traffic
engineering – Traffic engineering – ISDN traffic
engineering

Methods for cell level traffic control in B-ISDN

ITU-T Recommendation E.736

(Previously CCITT Recommendation)

ITU-T E-SERIES RECOMMENDATIONS

OVERALL NETWORK OPERATION, TELEPHONE SERVICE, SERVICE OPERATION AND HUMAN FACTORS

OPERATION, NUMBERING, ROUTING AND MOBILE SERVICES

INTERNATIONAL OPERATION	E.100–E.229
OPERATIONAL PROVISIONS RELATING TO CHARGING AND ACCOUNTING IN THE INTERNATIONAL TELEPHONE SERVICE	E.230–E.299
UTILIZATION OF THE INTERNATIONAL TELEPHONE NETWORK FOR NON-TELEPHONY APPLICATIONS	E.300–E.329
ISDN PROVISIONS CONCERNING USERS	E.330–E.399

QUALITY OF SERVICE, NETWORK MANAGEMENT AND TRAFFIC ENGINEERING

NETWORK MANAGEMENT	E.400–E.489
International service statistics	E.400–E.409
International network management	E.410–E.419
Checking the quality of the international telephone service	E.420–E.489
TRAFFIC ENGINEERING	E.490–E.799
Measurement and recording of traffic	E.490–E.505
Forecasting of traffic	E.506–E.509
Determination of the number of circuits in manual operation	E.510–E.519
Determination of the number of circuits in automatic and semi-automatic operation	E.520–E.539
Grade of service	E.540–E.599
Definitions	E.600–E.699

ISDN traffic engineering **E.700–E.749**

Mobile network traffic engineering	E.750–E.799
------------------------------------	-------------

QUALITY OF TELECOMMUNICATION SERVICES: CONCEPTS, MODELS, OBJECTIVES AND DEPENDABILITY PLANNING E.800–E.899

Terms and definitions related to the quality of telecommunication services	E.800–E.809
Models for telecommunication services	E.810–E.844
Objectives for quality of service and related concepts of telecommunication services	E.845–E.859
Use of quality of service objectives for planning of telecommunication networks	E.860–E.879
Field data collection and evaluation on the performance of equipment, networks and services	E.880–E.899

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

ITU-T RECOMMENDATION E.736

METHODS FOR CELL LEVEL TRAFFIC CONTROL IN B-ISDN

Summary

This Recommendation is concerned with the definition of control procedures which allow cell level GOS objectives to be fulfilled. The primary objective is to define practical CAC procedures allowing the network operator to decide when a new connection can be accepted on individual ATM links or network VPCs. Theoretical background is included where necessary to clarify assumptions and to situate the context of proposed control options. This Recommendation also addresses adaptive resource management techniques, where these are required by defined ATM transfer capabilities, and it identifies procedures for service integration.

Source

ITU-T Recommendation E.736 was prepared by ITU-T Study Group 2 (1997-2000) and was approved under the WTSC Resolution No. 1 procedure on the 26th of May 1997.

FOREWORD

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

The World Telecommunication Standardization Conference (WTSC), which meets every four years, establishes the topics for study by the ITU-T Study Groups which, in their turn, produce Recommendations on these topics.

The approval of Recommendations by the Members of the ITU-T is covered by the procedure laid down in WTSC Resolution No. 1.

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

NOTE

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

INTELLECTUAL PROPERTY RIGHTS

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

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

© 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.

CONTENTS

		<i>Page</i>
1	Scope	1
2	References	1
3	Terms and definitions	1
4	Abbreviations	2
5	Introduction	2
6	Traffic parameters and parameter control	4
	6.1 Source traffic descriptor	4
	6.1.1 Source peak rate and cell delay variation tolerance	4
	6.1.2 Sustainable cell rate parameter set	4
	6.2 Cell traffic variables	5
	6.3 Quality of service requirements	5
	6.4 Cell loss priority bit	6
	6.5 Parameter control	6
	6.6 Traffic shaping	6
7	Performance of an ATM multiplexer	7
	7.1 Multiplexing constant bit-rate streams	7
	7.1.1 Buffer overflow probability	7
	7.1.2 Impact of CDV	9
	7.1.3 Negligible CDV for CBR streams	9
	7.1.4 Nominal multiplexer rate	10
	7.2 Rate envelope multiplexing	10
	7.2.1 Cell loss ratio	11
	7.2.2 Multiplexing efficiency	11
	7.2.3 Cell loss priority	12
	7.2.4 Cell delay priority	12
	7.3 Rate sharing statistical multiplexing	12
	7.3.1 Buffer saturation probability	13
	7.3.2 Cell loss priority	13
	7.3.3 Cell delay priority	13
	7.4 Networks of multiplexer queues	13
	7.4.1 Multiplexing constant rate streams	13
	7.4.2 Rate envelope multiplexing	14
	7.4.3 Rate sharing	14
8	Connection admission control for DBR and SBR transfer capabilities	14
	8.1 CAC for peak rate allocation	15
	8.2 CAC for rate envelope multiplexing	15
	8.2.1 Known cell traffic variables	16
	8.2.2 Worst-case resource allocation	17
	8.2.3 Adaptive CAC	17
	8.3 CAC for rate sharing statistical multiplexing	18
	8.3.1 Known cell traffic variables	18
	8.3.2 Worst-case resource allocation	18
	8.3.3 Adaptive CAC	18
9	Adaptive resource management	18
	9.1 ABT resource management	19
	9.1.1 REM and block transfer	19
	9.1.2 Rate sharing and block transfer	19
	9.2 ABR resource management	19

	<i>Page</i>
10 Service integration	20
10.1 Dedicated resources	20
10.2 Loss priorities	20
10.3 Delay priorities	20
10.4 Scheduling policies	20
11 Bibliography.....	20
12 History.....	21
Appendix I – Example methods for calculation of equivalent cell rate for rate envelope multiplexing	21
I.1 Equivalent cell rate method 1 [RMV96].....	22
I.2 Equivalent cell rate method 2 [RMV96].....	22
I.3 Equivalent cell rate method 3 [ViS 97].....	22

METHODS FOR CELL LEVEL TRAFFIC CONTROL IN B-ISDN

(Geneva, 1997)

1 Scope

This Recommendation describes performance evaluation methods and traffic control methods enabling a network operator to meet objectives for cell level network performance. This Recommendation clarifies the traffic engineering consequences of traffic control and congestion control mechanisms and procedures defined in Recommendation I.371. Complementary Recommendations are E.735, which outlines the B-ISDN resource allocation framework, and E.737 which provides dimensioning guidelines enabling the network operator to meet call level performance objectives.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; all users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

- ITU-T Recommendation E.177 (1996), *B-ISDN routing*.
- ITU-T Recommendation E.716 (1996), *User demand modelling in B-ISDN*.
- ITU-T Recommendation E.737 (1997), *Dimensioning methods for B-ISDN*.
- ITU-T Recommendation E.735 (1997), *Framework for traffic control and dimensioning in B-ISDN*.
- ITU-T Recommendation I.356 (1996), *B-ISDN ATM layer cell transfer performance*.
- ITU-T Recommendation I.371 (1996), *Traffic control and congestion control in B-ISDN*.

3 Terms and definitions

This Recommendation defines the following terms:

3.1 equivalent cell rate: A cell rate attributed to a connection such that cell level GOS objectives are satisfied on an ATM link or network VPC as long as the sum of equivalent cell rates is not greater than the rate of the ATM link or VPC.

3.2 rate envelope multiplexing: A statistical multiplexing scheme where CAC aims to make negligible the probability that the combined arrival rate of multiplexed connections exceeds multiplexer capacity; buffering is employed solely to account for the deviation of the cell arrival process from a fluid ideal where no buffer would be required to meet cell level GOS objectives.

3.3 rate sharing: A statistical multiplexing scheme where a buffer is used to absorb excess cells when the arrival rate is greater than the multiplexer output rate for significant periods of time; a buffer would be necessary to meet cell level GOS objectives even in the fluid ideal.

4 Abbreviations

This Recommendation uses the following abbreviations.

ABR	Available Bit Rate
ABT	ATM Block Transfer
ATM	Asynchronous Transfer Mode
BECN	Backward Explicit Congestion Notification
CAC	Connection Admission Control
CDV	Cell Delay Variation
CLR	Cell Loss Ratio
DBR	Deterministic Bit Rate
ECR	Equivalent Cell Rate
FECN	Forward Explicit Congestion Notification
GCRA	Generic Cell Rate Algorithm
GOS	Grade of Service
IBT	Intrinsic Burst Tolerance
INI	Inter-Network Interface
MBS	Maximum Burst Size
NPC	Network Parameter Control
PCR	Peak Cell Rate
QOS	Quality of Service
REM	Rate Envelope Multiplexing
RM	Resource Management
SBR	Statistical Bit Rate
SCR	Sustainable Cell Rate
STD	Source Traffic Descriptor
UNI	User-Network Interface
UPC	Usage Parameter Control
VCC	Virtual Channel Connection
VPC	Virtual Path Connection

5 Introduction

Recommendation I.371 defines the scope of ATM layer traffic control and congestion control identifying a variety of functions ranging from network resource management to priority controls, acting over a wide range of time scales. This Recommendation is concerned with the traffic engineering implications of the different ATM transfer capabilities standardized in Recommendation I.371 including the definition of Connection Admission Control (CAC) and resource allocation procedures. Other control actions such as usage parameter control are briefly considered in so far as it is necessary to achieve a high level of consistency between the various control capabilities.

Traffic controls may be distinguished according to whether their function is to enable quality of service guarantees at cell level (e.g. cell loss ratio) or at call level (e.g. call blocking probability). This Recommendation is concerned with cell level controls. Consideration is restricted to the CAC procedure applied to a single ATM link or network-to-network VPC which determines simply if that ATM link or VPC is capable or not of handling the requested connection. The issue of traffic routing (i.e. determining a network path from among those possible) is dealt with in Recommendation E.177.

B-ISDN is a connection-oriented network. Each connection is defined by a set of traffic parameters and QOS requirements. When the establishment of a new connection is requested, the network must decide if it has sufficient resources to accept it without infringing cell level GOS requirements for all established connections as well as the new connection; this is the role of CAC. Given that a connection is accepted, the network must ensure that the user in fact emits traffic in conformity with the declared traffic parameters; this is the role of Usage Parameter Control (UPC). When more than one network is involved in a connection, it is also incumbent on each network to verify that the traffic it receives from the neighbouring network conforms; this is Network Parameter Control (NPC). Standard traffic parameters and the algorithms by which the conformity of connections can be checked in UPC/NPC mechanisms are defined in Recommendation I.371. This Recommendation is concerned with the definition of CAC procedures which allow QOS requirements to be fulfilled taking account of the information available about connection traffic and the accuracy with which it can be controlled. This Recommendation also addresses adaptive resource management techniques, where these are required by defined ATM transfer capabilities, and it identifies procedures for service integration.

Depending on network architecture, cell level traffic controls may be applied to different transmission entities. Recommendation E.735 defines the physical and logical network entities which constitute the framework for cell level traffic control. In this Recommendation, it is generally assumed that a connection is offered to an ATM link or a network-to-network VPC defined by a DBR traffic descriptor or by traffic variables as considered in Recommendation E.735 (i.e. an uncontrolled constant rate VPC or a variable rate VPC).

Effective traffic control of ATM connections, particularly when it is an objective to perform statistical multiplexing of variable bit rate connections, relies on a sound understanding of how the performance of multiplexing stages depends on source traffic characteristics. Much of this understanding will only be gained with experience in operating the B-ISDN. Even though considerable progress has been made over recent years, it is necessary to recognize that knowledge of both traffic characteristics and their impact on network performance remains limited. It is also true that there is still considerable divergence in the scientific community about the effectiveness of different modelling approaches and their applicability to the range of connection types. For these reasons, this Recommendation is not restricted to a simple list of traffic control recipes. It has been the intention to also include some theoretical background in order to clarify assumptions and to situate the context of proposed control options. Where possible, unambiguously defined control rules and algorithms are clearly presented in the text of this Recommendation and can be applied without knowledge of the background material.

The following ATM transfer service capabilities are defined in Recommendation I.371:

- Deterministic Bit Rate (DBR);
- Statistical Bit Rate (SBR);
- Available Bit Rate (ABR);
- ATM Block Transfer (ABT).

This Recommendation considers the traffic engineering implications of implementing these different service categories.

Clause 6 discusses the question of defining traffic parameters and UPC/NPC algorithms since these strongly influence the choice of possible CAC procedures and their efficiency. Clause 7 then introduces a number of traffic modelling considerations which underly the relationship between traffic, capacity and performance. This relationship is the basis for resource allocation in connection admission control. Clause 8 presents a number of CAC possibilities for DBR and SBR transfer capabilities. Clause 9 is devoted to the use of resource management procedures to adapt resource allocation

to changing traffic conditions during the lifetime of a connection as allowed in ABR and ABT transfer service capabilities. Finally, in Clause 10, this Recommendation discusses cell level traffic controls enabling network resources to be shared between connections with different characteristics and QOS requirements established using different ATM transfer capabilities.

6 Traffic parameters and parameter control

A connection request is specified by a traffic descriptor, Cell Delay Variation (CDV) tolerance and QOS requirements. Based on current traffic conditions, the network must decide whether or not it is possible to accept the connection request (Connection Admission Control). If the connection is accepted, there is implicitly defined a traffic contract whereby the network operator provides the requested quality of service on condition that the user emits traffic in conformity with the declared traffic descriptor. Note that CAC can be performed on the basis of traffic parameters included in the traffic descriptor or, alternatively, on the basis of the cell traffic variables defined in Recommendation E.716 when these are known or can be deduced.

6.1 Source traffic descriptor

The source traffic descriptor is a list of traffic parameters each of which should (see Recommendation I.371):

- be understandable for the user or his terminal; conformance should be possible;
- participate in resource allocation schemes meeting network performance requirements;
- be enforceable by the UPC and NPC.

The traffic parameters may relate explicitly to connection traffic characteristics such as the peak cell rate or implicitly define these characteristics by reference to a service type.

6.1.1 Source peak rate and cell delay variation tolerance

Peak Cell Rate (PCR) is defined in Recommendation I.371 as the inverse of the minimal cell inter-arrival interval observed at a certain equivalent terminal. For the network to verify conformity with a declared value, however, it is necessary to account for cell delay variation occurring between the equivalent terminal and the observation point. Conformity to the declared peak cell rate is determined by the so-called Generic Cell Rate Algorithm (GCRA) defined in Recommendation I.371 based on the cell delay variation tolerance τ .

When PCR is less than the rate of the link on which a connection is carried, the CDV tolerance allows a certain variability in the connection bit rate. Of some interest is the maximum length of a burst at link rate compatible with the parameters PCR and τ . A succession of such bursts, separated by silent periods, may be considered as a worst-case traffic for traffic engineering purposes. Let the link rate be LR . The definition of τ is such that the Maximum Burst Size (MBS) is:

$$MBS = \lfloor 1 + \tau / (1/PCR - 1/LR) \rfloor \quad (6-1)$$

where $\lfloor x \rfloor$ denotes the integer part of x .

6.1.2 Sustainable cell rate parameter set

Traffic parameters Sustainable Cell Rate (SCR) and Intrinsic Burst Tolerance (IBT) are defined in Recommendation I.371 with respect to the generic cell rate algorithm.

The *SCR/IBT* parameter set applies to a wide range of traffic streams with different cell traffic variables. For traffic engineering purposes, it is useful to characterize a worst-case traffic compatible with a given STD. It is assumed in this Recommendation that the worst-case traffic compatible with given *PCR*, *SCR* and *IBT* is a stream of maximal length bursts at rate *PCR*. The Maximum Burst Size (*MBS*), measured in cells, is:

$$MBS = \lfloor 1 + IBT/(1/SCR - 1/PCR) \rfloor \quad (6-2)$$

For non-zero CDV tolerance, the "peak rate" bursts can themselves be a succession of link rate bursts as discussed in 6.1.1.

6.2 Cell traffic variables

The source traffic descriptor traffic parameters are necessarily defined with respect to a rule to enable enforcement at the UPC and NPC. In some cases, notably when a connection is described by a service type or when traffic characteristics are determined by network controlled operations (e.g. the formation of a network-to-network VPC), it may be possible to more closely characterize a connection using cell traffic variables. Cell traffic variables refer directly to the statistical properties of the connection traffic. Examples of cell traffic variables useful for traffic engineering are given in Recommendation E.716.

Of particular relevance for the CAC procedures considered in this Recommendation for rate envelope multiplexing (see 7.2) are cell traffic variables describing the probability distribution of the source rate at an arbitrary instant t denoted Λ_t . For example, cell traffic variables for an on/off source might describe the rate of the source when on (its peak rate) and the probability the source is on.

To predict performance, and therefore to perform CAC in the case of rate sharing (see 7.3), further cell traffic variables describing the transient nature of rate variations (i.e. not just the stationary probability distribution of the instantaneous rate) are necessary. Such traffic variables relate to the number of cells arriving in certain time intervals or the number of arrivals exceeding a certain rate, as discussed in Recommendation E.716.

The relationship between cell traffic variables and source traffic descriptor parameters is discussed in Recommendation E.716.

6.3 Quality of service requirements

End-to-end cell level QOS criteria use the following performance parameters defined in Recommendation I.356:

- cell transfer delay;
- cell delay variation;
- cell loss ratio.

End-to-end performance objectives relevant to traffic engineering are identified in Recommendation E.735 as the following:

- maximum end-to-end queueing delay, defined as a remote quantile (10^{-8} , say) of the delay distribution;
- mean queueing delay;
- cell loss ratio.

These performance objectives must be apportioned to the various network elements contributing to the performance degradation of a given connection so that the end-to-end QOS criteria are satisfied.

For the purposes of this Recommendation, it is assumed that each link has assigned target values for the three performance objectives. The defined target values must be mutually compatible (e.g. the mean delay must be less than the maximum) but all three do not necessarily intervene in the CAC procedure.

6.4 Cell loss priority bit

The use of the CLP bit in the cell header is to assign high (CLP = 0) and low (CLP = 1) loss priorities to the cells offered to the network. The assignment is made by the user but, as an option, the UPC/NPC can tag non-compliant cells by setting their CLP bit to 1. It is understood that in the event of congestion a network might discard CLP = 1 cells in preference to CLP = 0 cells. Traffic parameters and quality of service parameters must then be declared distinguishing the two types of cell. Current Recommendations (Recommendation I.371, for example) are that parameters be declared for the CLP = 0 stream, on the one hand, and the combined CLP = 0 + 1 stream, on the other hand.

6.5 Parameter control

One of the three requirements on traffic parameters included in the STD is that they be enforceable by the UPC and NPC. This has led to a definition of traffic parameters peak cell rate, sustainable cell rate and intrinsic burst tolerance allowing user conformance to be determined by reference to a rule or algorithm, namely the generic cell rate algorithm.

The GCRA is standardized in Recommendation I.371. The algorithm actually used for parameter control is not standardized, however. The implemented algorithm should be transparent to a conforming cell stream and take appropriate action when any declared parameter value is exceeded in order to protect the quality of service of other connections.

The network can implement a traffic contract with declared traffic parameters for CLP = 0 cells alone and declared traffic parameters for CLP = 0 + 1 cells. Non-conforming CLP = 0 cells may be "tagged" and admitted to the network as CLP = 1 cells. The traffic engineering implications of these possibilities are for further study.

6.6 Traffic shaping

Users or networks may introduce supplementary cell delays to shape the characteristics of a given stream. By smoothing cell rate variations, shaping generally allows an increase in the utilization of network resources leading to greater multiplexing gains. On the other hand, shaping may introduce non-negligible delays and a part of the end-to-end GOS objective must be allocated to the shaper.

Shaping may be performed by the user to ensure compliance with declared traffic parameters and CDV tolerance. The network operator may employ shaping at the network entrance, within the network or at the network egress (to meet constraints on output traffic characteristics). Shaping is an option for users and networks.

A particular example of shaping is the reduction of CDV by means of cell spacing. The spacer tries to produce a cell stream with a time between consecutive cells at least equal to the peak cell emission interval (the inverse of the PCR) by imposing a variable delay on each cell. Figure 6-1 depicts a spacer acting on an access line of rate LR assuming a total of m virtual connections are to be spaced to their individual peak rate parameter PCR_i with $\sum PCR_i \leq LR$. Cells for connection i are dispatched to a FIFO queue served at rate PCR_i before being re-integrated with the cells of other connections in a FIFO queue served at rate LR . The residual CDV at the output of this spacer is equivalent to that generated by a single FIFO multiplexing stage receiving periodic cell streams. Spacer realizations typically rely on scheduling algorithms and do not involve a physical queue for each connection. Alternative shaper designs using more sophisticated scheduling algorithms can produce connections with CDV less than that of the above example.




Figure 6-1/E.736 – Conceptual spacer for m peak rate allocated connections ($\sum PCR_i \leq LR$)

Another example of shaping is the reduction of the PCR for variable rate connections.

7 Performance of an ATM multiplexer

The generic network element determining network performance is the ATM multiplexer. This is defined for present purposes to be a device allowing several streams of cells to exclusively share a transmission capacity of rate c (e.g. an ATM link or DBR VPC) with a buffer of capacity B allowing cells to be stored temporarily while awaiting transmission. The multiplexer may implement different queue disciplines or scheduling policies. In this clause, FIFO service is generally assumed. Service disciplines discriminating cells according to the CLP bit are however discussed in 7.2.3 and 7.3.2. Head of line priority and scheduling algorithms are discussed in Clause 10 on service integration. Note that the buffer capacity B may correspond to the physical capacity of the equipment or be determined by the maximum delay performance constraint: if the maximum delay is W_{\max} we assume the ATM multiplexer has a buffer of size B given by:

$$B = \min \{ cW_{\max}, \text{physical capacity} \} \quad (7-1)$$

This clause discusses models which may be used to predict the performance of an ATM multiplexer. These performance models may be useful in dimensioning the multiplexer buffers or in determining the amount of traffic which can be handled by a multiplexer of given dimensions. It is assumed here that source traffic characteristics are independent of multiplexer performance. The considered models do not therefore apply to the available bit rate transfer capability. The performance of ABR multiplexing is considered in clause 9.

It is useful to distinguish three different operating principles: multiplexing of constant bit rate streams (7.1), statistical multiplexing of variable bit rate streams with "Rate Envelope Multiplexing" (7.2), statistical multiplexing of variable bit rate streams with "Rate Sharing" (7.3). The impact on performance of a succession of multiplexer stages is considered in 7.4.

7.1 Multiplexing constant bit-rate streams

It is assumed here that the sum of bit rates of the multiplexed CBR streams is less than the multiplexer bit rate. Streams are first assumed to be perfectly periodic at the multiplexer input (i.e. no cell delay variation). Let the multiplex rate be c cells/sec and consider N streams, all with the same cell inter-arrival interval of T seconds, emitting cells independently in the sense that, in any interval of length T , the arrival epoch of each of the N cells is uniformly and independently distributed in the interval. Let $D = Tc$ be the normalized cell inter-arrival interval when the cell transmission time is taken as time unit.

7.1.1 Buffer overflow probability

For any fixed set of streams, the buffer occupancy state is a periodic process of period T . If the buffer overflows during this period and cells are lost, some streams will lose all their cells while the others will lose none. To render the probability of this occurrence less than a target level (10^{-9} , say), the multiplexer buffer may be dimensioned so that, for a randomly chosen set of stream phases, the probability $Q(B)$ of the queue in an unlimited buffer exceeding the capacity

B cells at an arbitrary instant is less than the target value. For the small probabilities generally considered, $Q(B)$ can be assimilated to the saturation probability of a buffer of capacity B . This probability is given by the following expression [RMV96]:

$$Q(B) = \sum_{B < n \leq N} \binom{N}{n} \left(\frac{n - B}{D} \right) \left(1 - \frac{n - B}{D} \right)^{N-n} \frac{D - N + B}{D - n + B} \quad (7-2)$$

Note that the equivalent probability evaluated at an arrival instant is given by the same formula with N replaced by $N - 1$. A simpler approximate formula which gives good order of magnitude estimates at load (N/D) greater than 0.8 is [RMV96]:

$$Q(B) \approx \exp \left\{ -2B \left(\frac{B}{N} + \frac{D - N}{N} \right) \right\} \quad (7-3)$$

For the very small probabilities of interest, $Q(B)$ constitutes a tight upper bound on the cell loss ratio.

In general, streams do not have the same rate. The mixture of rates extends the period of the queue length process to the lowest common multiple of the stream cell inter-arrival intervals and attenuates the possible concentration of cell loss on particular streams. Let the N multiplexed streams have normalized inter-arrival intervals D_i for $i = 1, \dots, N$. The multiplexer load is $\sum 1/D_i$. The calculation of the buffer overflow probability $Q(B)$ for such a mixture of bit rates proves complicated. However, empirical results suggest that an upper bound on buffer requirements can be obtained on assuming N identical streams of period $D = N \times (\sum 1/D_i)^{-1}$ and using the Formula 7-2.

Numerical evaluations of Formula 7-2 provide the results of Table 7-1 giving the buffer size B cells such that $Q(B) < 10^{-9}$ for various numbers of sources and multiplexer loads.

Table 7-1/E.736

N	50	50	500	500	5000	5000	∞	∞
N/D	.80	.95	.80	.95	.80	.95	.80	.95
B	19	22	37	61	47	135	48	204

The last two columns correspond to results for the M/D/1 queue, i.e. assuming Poisson arrivals. This traffic model may be used as a tool for worst-case dimensioning in the absence of an upper limit on the number of multiplexed connections. As seen in the results of Table 7-1, the M/D/1 results provide conservative estimates of buffer requirements and constitute a good approximation when the number of sources is high and the multiplexer load is not too close to 1. An accurate approximation for the M/D/1 queue length distribution is [RMV96]:

$$Q(B) \approx C e^{-rB} \quad (7-4)$$

where $C = (1 - \rho) / (\rho e^r - 1)$ and r is the solution of the equation $\rho(e^r - 1) - r = 0$.

Note that the assumption of Poisson arrivals corresponds to a worst-case traffic model for any superposition of periodic streams (homogeneous or heterogeneous) having the same overall average arrival rate in the sense that all quantiles of the delay distribution are greater. In particular, the Cell Loss Ratio (CLR) estimated by $Q(B)$ is greatest for Poisson arrivals.

It is sometimes convenient to define batch arrival processes where exactly k cells arrive at each arrival instant. The above formulae can be used to estimate the buffer saturation probability for corresponding batch arrival systems on replacing B by B/k . For example, the buffer saturation probability when cells arrive in batches of k according to a Poisson process (the $M^{(k)}/D/1$ queue) may be estimated by $Q(B/k)$ where $Q(\cdot)$ is given by Formula 7-4.

7.1.2 Impact of CDV

Cell delay variation, due to a variety of reasons as discussed in Recommendation I.371, alters the precise periodicity of a CBR stream. In particular, CDV is acquired by a CBR stream as it passes each multiplexing stage on its path through a network due to the different queueing delays affecting successive cells. For traffic engineering purposes, and notably for CAC, it is important to understand how this CDV affects multiplexer performance.

In fact, it is often sufficient to characterize a stream subject to CDV by whether or not it leads to better performance in comparison with a given reference stream in the sense that relevant performance parameters [e.g. CLR or the buffer saturation probability $Q(B)$] would be worse if the considered stream were replaced by the reference stream. In this sense, as indicated in the previous subclause, any superposition of periodic CBR sources is better than a Poisson arrival process of the same average rate.

In this Recommendation, the impact of CDV on CBR streams is distinguished as being "negligible" or "non-negligible". The notion of negligible CDV is defined precisely below and allows for fairly simple CAC procedures. The impact on performance and traffic engineering of non-negligible CDV is for further study.

7.1.3 Negligible CDV for CBR streams

The notion of negligible CDV is defined with respect to a reference arrival process and a given performance parameter. A stream is said to have negligible CDV if the realized value of the considered performance parameter would not be better if the stream were replaced by the reference process having the same rate. An appropriate performance parameter for CBR streams is the saturation probability which takes into account both loss and delay constraints. The reference process defines the traffic assumed for traffic engineering purposes. Note that by this definition, the reference process itself has negligible CDV and any superposition of independent streams with negligible CDV also has negligible CDV.

In this Recommendation, the considered reference process is a batch arrival Poisson process with constant sized batches of k cells for some value of $k \geq 1$. This process will be referred to as a k -batch Poisson process (or just Poisson process if $k = 1$). This choice is motivated by the fact that the superposition of streams satisfying $\tau \cdot PCR \leq k - 1$ has negligible CDV with respect to a k -batch Poisson process.

If all connections handled by a multiplexer have rates defined with negligible CDV with respect to a k -batch Poisson process (with a common value of k), then performance is better than that of the corresponding $M^{(k)}/D/1$ queue and the saturation probability can be conservatively estimated by $Q(B/k)$ as given by Formula 7-4.

To establish whether a given stream has negligible CDV, the following guidelines are proposed:

- if a stream of rate PCR has been shaped in a cell spacer as depicted in Figure 6-1, it has negligible CDV with respect to a Poisson process of rate PCR ;
- if a stream is characterized by the traffic descriptor PCR and an associated CDV tolerance τ , it has negligible CDV with respect to a k -batch Poisson process with $k \geq \tau \cdot PCR + 1$;
- according to the conjecture formulated in 7.4.1 below, a stream which has negligible CDV and which is multiplexed in stable queues with other streams with negligible CDV (with respect to the same reference process) retains the property of negligible CDV on output.

Based on the above guidelines, a possible operator policy based on the notion of negligible CDV might be the following:

- fix a certain reference process (i.e. a certain value of k) and use this to perform traffic engineering as detailed in this Recommendation;
- shape connections for which the CDV tolerance τ satisfies $\tau \cdot PCR > k - 1$.

7.1.4 Nominal multiplexer rate

For a multiplexer of given output rate c and buffer size B , it is useful to define a multiplexer nominal capacity c' determined such that the CLR objective is satisfied when the sum of connection PCR parameters is not greater than c' . The choice of c' must take account of the CDV tolerance of the connections to be multiplexed without being specific to a given mix of connections (e.g. a worst-case calculation compatible with maximum allowed CDV tolerance). A new connection is accepted if its PCR added to the sum of PCR values of existing connections is less than or equal to c' .

This notion of nominal rate may be related to that of negligible CDV discussed in 7.1.3. Suppose the value of c' is determined such that a given reference traffic (for example, a Poisson cell stream) would result in cell loss less than ϵ as long as the multiplexer load is less than c'/c . The simple admission criterion of comparing the sum of PCR values with c' then applies to traffic streams which have negligible CDV with respect to the given reference arrival process.

7.2 Rate envelope multiplexing

For many service types it is natural to identify source states during which the cell emission rate is approximately constant (e.g. on/off sources or sources whose rate changes between different levels). For present purposes, it is assumed that the notion of instantaneous cell arrival rate is well-defined for a connection or group of connections. For instance, the instantaneous arrival rate of a group of on/off sources would be the sum of the rates of currently active sources.

Statistical multiplexing of variable rate streams can be performed by assuring that the combined instantaneous input rate, as discussed above, is not greater than the multiplexer service rate. This can be achieved by restricting the offered cell stream by appropriate CAC and UPC/NPC mechanisms and/or by modifying the multiplexer service rate by dynamic resource management, for example. The objective being to keep the cell arrival rate within the limit defined by the service rate, this multiplexing scheme is referred to as Rate Envelope Multiplexing (REM).

In a fluid analogy where a cell arrival stream of combined rate λ is viewed as a fluid flow of the same rate, REM is clearly distinguished from other multiplexing schemes by the fact that it does not rely on buffering: whenever the input rate λ is greater than the service rate c then cells are lost at rate $(\lambda - c)$. In this fluid analogy, traffic control procedures are considerably simpler than those necessary for multiplexers which rely on buffers to absorb momentary rate overloads. This simplification motivates the present definition of REM which adapts control procedures applicable in the fluid analogy to the cell-by-cell traffic streams of the real network.

With REM, a small multiplexer buffer (i.e. about 100 cell places) is required just to handle the queue arising due to asynchronous arrivals from streams whose combined rate is less than the multiplexer rate. However, REM traffic controls do not exploit any gain in efficiency arising from the absorption of rate overloads by this buffer. The size of the buffer is determined as for multiplexing CBR streams so that the time to serve a cell entering the last place is compatible with connection delay GOS objectives.

Reasons for using REM as a resource sharing strategy include the following:

- it is possible to provide performance guarantees without knowing statistical details of the burst structure (only the stationary rate distribution is relevant);
- only small buffers are required simplifying multiplexer design;
- cell transfer delay is very small and meets strict performance objectives.

Conversely, REM may be considered necessary when traffic characteristics describing the burst structure are unknown, when multiplexers are not equipped with large buffers or when cell transfer delay requirements are strict. A further advantage of REM is that CAC procedures are simplified.

7.2.1 Cell loss ratio

It is convenient in the case of REM to decompose the cell loss ratio into a "burst-scale" component CLR_{bs} corresponding to losses due to rates greater than multiplexer capacity calculated using the fluid analogy and a "cell scale" component CLR_{cs} corresponding to a correction term necessary to account for the deviation of the real cell arrival process from the fluid ideal.

Let Λ_t be the combined bit rate of all streams at time t . The burst-scale component of CLR is then:

$$CLR_{bs} = E\left\{\left(\Lambda_t - c\right)^+\right\} / E\left\{\Lambda_t\right\} \quad (7-5)$$

It proves difficult to estimate exactly the cell scale component. However, when the rate of multiplexed traffic streams is defined with negligible CDV with respect to a k -batch Poisson process, the following approximation is satisfactory for traffic engineering purposes:

$$CLR_{cs} \approx Q(B/k) \quad (7-6)$$

where $Q'(B/k)$ is given by Formula 7-4 evaluated for an arrival rate equal to the mean arrival rate of the multiplexed streams. Some guidelines for determining when rates are defined with negligible CDV are given in 7.1.3.

7.2.2 Multiplexing efficiency

A constraint on CLR_{bs} defines an implicit relation between the characteristics of the offered traffic and the achievable multiplexer occupancy. In particular, consider the following example which illustrates the role of the connection peak rate.

N identical on/off sources of peak rate p and mean rate m are multiplexed on a link of capacity c . CLR_{bs} is then estimated by:

$$CLR_{bs} \approx \sum_{ip > c'} (ip - c) \binom{N}{i} \left(\frac{m}{p}\right)^i \left(1 - \frac{m}{p}\right)^{N-i} \times \frac{1}{Nm} \quad (7-7)$$

A constraint on CLR_{bs} (e.g. $CLR_{bs} < 10^{-9}$) imposes a limit on achievable multiplex utilization Nm/c . Assume a target CLR_{bs} of 10^{-9} . The achievable load compatible with the limiting overload probability can then be calculated as a function of the source peak rate. This function is plotted in Figure 7-1 for several values of N (including the limiting case where the number of sources tends to infinity) and several values of the ratio p/m . The figure illustrates that a high-link utilization is only possible here when the peak rate is a small fraction of the multiplex link rate unless the sources have a low peak to mean rate ratio. For illustration purposes, consider a link of capacity 100 Mbit/s; to accommodate bursts of peak rate 20 Mbit/s ($c/p = 5$) with 10^{-9} CLR_{bs} would require mean utilization to be limited to around 2%; to achieve 50% utilization with the same objective requires either the sources to be very slightly bursty (e.g. $p/m = 2$) when N is small or of very low peak rate ($p \ll c$) when N is large.

In general, while the achievable link load depends on the precise traffic mix, it may be stated that REM can be efficient for bursty sources with relatively low peak rates but can require a rather low network link utilization if bursty streams with peak rates comparable to the link rate are to be carried. In all cases, however, REM is never less efficient than peak rate allocation.




Figure 7-1/E.736 – Achievable load against link to peak ratio (c/p)

7.2.3 Cell loss priority

The CLP bit allows a multiplexer to implement discriminatory cell discard and thus to satisfy two cell loss ratio constraints, one for CLP = 0 cells and one for CLP = 1 (or CLP = 0 + 1). It is also possible to discriminate between different connections by rejecting cells in some priority order according to the value of the virtual connection identifier. In the case of the CLP bit, two discard strategies have been identified:

- "threshold" or partial buffer sharing, where CLP = 1 cells are discarded by the multiplexer when its queue length exceeds a given threshold; and
- "pushout", where CLP = 0 cells arriving to find a full queue can displace a queued CLP = 1 cell.

These queue mechanisms are referred to as "loss priority" mechanisms in opposition to traditional head of line disciplines which offer delay priority. Loss priority mechanisms should preserve the order of cells in a given connection whether these be CLP = 0 or CLP = 1 cells. The performance of such queueing disciplines has been studied in [RMV96] and [LuP90].

As a first approximation, it can be assumed that, in case of overload, no CLP = 0 cell is discarded unless the combined arrival rate of CLP = 0 cells is greater than the multiplexer nominal rate, the models described in 7.2 can be used to predict the performance of CLP = 0 cells alone (using the CLP = 0 arrival rate as Λ_r) and of all cells (using the overall arrival rate as Λ_t).

When discriminating between different connections, it is not necessary as in the case of CLP to preserve cell sequence integrity. However, the use of the above loss priority queueing disciplines ensures that delays of transmitted cells are limited.

7.2.4 Cell delay priority

Head of line queueing disciplines can be used to provide delay priorities to different connections. However, this type of operation negates the advantages of using REM as detailed in 7.2. Cell delay priority is a more appropriate option in the case of rate sharing. Delay priority may also be employed to allow REM to be used for a group of higher priority connections while the remainder use buffering to share the remaining bandwidth.

7.3 Rate sharing statistical multiplexing

Statistical multiplexing of VBR streams can be performed with higher link utilization than that achievable with REM if multiplexers are equipped with a larger buffer to absorb the excess traffic arriving when the combined arrival rate is momentarily greater than the link transmission rate. This statistical multiplexing scheme is referred to as rate sharing to

distinguish it from REM where it is as if streams have a dedicated bit rate whenever they require it. In general, the larger the buffer, the higher the achievable link utilization for a given cell loss ratio. However, larger buffers also imply potentially longer delays and it is necessary to verify that these are compatible with GOS objectives.

7.3.1 Buffer saturation probability

A number of models of a rate sharing ATM multiplexer have been proposed in the literature. These models all rely on some particular representation of the offered traffic and, as a rule, are more complex as the characteristics of multiplexed streams are more general. Further study is required to recommend any one model as providing a dimensioning tool enabling the definition of acceptable traffic mixes for given buffer size.

Simpler results valid in an asymptotic regime may however provide practical traffic engineering rules for certain traffic classes. It has been shown that, for a large class of arrival processes, the survivor function of the queue length in a multiplexer equipped with an unlimited buffer would be asymptotically exponential for large x , i.e. $Pr \{ \text{queue length} > x \} \approx \alpha e^{-\gamma x}$ (see, for example, [RMV96]). Note that this asymptotic limit may only be useful for estimating very small probabilities [CLW 94]. Furthermore, the exponential limit is not valid for certain types of traffic exhibiting long term dependence [RMV96].

Queueing models have been used to investigate the impact on performance of different traffic characteristics. It has been shown for traffic streams of the on/off type (see Recommendation E.716) in particular, that performance depends significantly on the first two moments of the burst and silence length distributions. Correlations in the burst generation process (bursts of bursts, etc.) also significantly affect the multiplexer queue length distribution. It follows that to predict the values of performance parameters such as the cell loss ratio and maximum or mean queueing delays requires the knowledge of such complex traffic characteristics.

Of particular interest is a model for a superposition of periodic on/off sources constituting a "worst-case traffic" compatible with *SCR* and *IBT* traffic parameters. Such models have been studied in the literature (e.g. see [RMV96]) but the deduction of practically useful traffic engineering procedures remains for further study.

7.3.2 Cell loss priority

Strategies for discriminating between $CLP = 0$ and $CLP = 1$ cells described in 7.2.3 are also applicable here. It is also possible to distinguish loss priorities between all the cells of different connections. However, prediction of the performance of the different priority streams suffers from the same problems outlined in 7.3.1 above.

7.3.3 Cell delay priority

Head of line delay priority queueing disciplines can be used to provide different qualities of service to specific groups of connections. The overall performance of the highest priority group can be evaluated approximately by considering only the arrival process of this group. The overall mean delay of the i highest groups can also be evaluated by considering the overall arrival process of these i groups.

7.4 Networks of multiplexer queues

Characteristics of connection traffic streams are altered as the cells progress through the multiplexing stages of the network path. It is necessary to be able to account for these alterations, notably, in performing CAC. The impact on the cell stream is different depending on the type of multiplexing employed.

7.4.1 Multiplexing constant rate streams

CBR streams are defined by the parameter *PCR* and its associated CDV tolerance. The following statement is supported by evidence from analytical and simulation studies on network performance but has not been formally proved:

If the CDV of all streams is negligible with respect to a k -batch Poisson reference process at the network input (i.e. at the UNI or INI) and multiplexing is performed subject to the condition that the sum of *PCR* values is less than the service rate at each multiplexing stage, then CDV remains negligible with respect to the same reference process throughout the network.

In particular, streams which were initially exactly periodic or which have been spaced at the network ingress to their nominal *PCR* do not acquire non-negligible CDV with respect to a Poisson process, no matter how many multiplexing stages they cross. If all streams in the network have CDV tolerance less than $(k - 1) T$ peak emission intervals ($\tau \leq (k - 1) T$) they have and retain negligible CDV with respect to a k -batch Poisson process.

7.4.2 Rate envelope multiplexing

With REM, the sum of rates of active sources can exceed the multiplexer rate c and cells can be lost. Cell loss changes source characteristics. However, if stream rates on input are defined with negligible CDV compared to a Poisson or k -batch Poisson reference process, they retain negligible CDV on output compared to the same reference process (i.e. a Poisson or k -batch Poisson stream with the rate of the input process, not the modified rate accounting for cell loss). This property allows CAC to be performed on all multiplexers within the network assuming the rate distribution is the same as that observable at the network ingress.

7.4.3 Rate sharing

The impact of rate sharing multiplexing on connection traffic characteristics is for further study.

8 Connection admission control for DBR and SBR transfer capabilities

When a user requests the setting up of a new connection, it is necessary for the network to decide if the connection can be admitted while satisfying the quality of service requirements of both new and existing connections. This decision can sometimes be made by allocating resources to specific connections or groups of connections and refusing new requests when insufficient resources are available. Note that the allocation is generally logical: no particular physical resources are attributed to a specific connection. The resources in question are typically bandwidth and buffer space. It is assumed in the following that resources are allocated independently for each ATM link or VPC of a path with a separate decision made for each transmission direction of a virtual connection. A connection will be established only if resources are available on every link of its path, in both directions.

The following discussion relates to a single ATM link or VPC, as defined in Recommendation E.735. A shaped DBR VPC is considered like an ATM link and both are assumed to be completely characterized by an output bit rate c cells/s and a buffering capacity B cells. For a shaped DBR VPC, the buffer size is determined by the shaping algorithm as depicted in Recommendation E.735. Uncontrolled constant rate and variable rate VPCs are characterized by traffic variables. In all cases, it is assumed that the output bit rate is "fully accessible" in the sense that the only access restriction derives from the total amount of allocated bit rate. A single set of performance objectives (cell loss ratio, maximum and mean delay) is considered corresponding to the most stringent requirements of all multiplexed connections. Extension to more general resource sharing schemes including priority controls is considered in clause 10.

Resources may be allocated once and for all at the start of a call or, following renegotiation, at some time in the course of the call. Resource renegotiation may be performed using Resource Management (RM) cells, in the case of ABR and ABT transfer capabilities, or using the out-of-band signalling system. This clause is confined to CAC procedures for DBR and SBR transfer capabilities. It is assumed that these procedures are applicable both for initial resource allocation at call set-up and subsequently in case of renegotiation conducted by signalling. The use of adaptive resource management procedures in ABR and ABT transfer capabilities is discussed in clause 9.

The way resource allocations can be related to connection characteristics is considered below in the three multiplexing schemes considered in clause 7 above.

8.1 CAC for peak rate allocation

When multiplexing constant bit rate streams, an obvious resource allocation scheme consists in allocating to each connection a bandwidth on each link equal to its declared bit rate (with due allowance for CDV). The same resource allocation procedure can be applied to variable rate connections if a bandwidth equal to the connection peak bit rate is reserved on every link.

Connections are characterized by their *PCR* and CDV tolerance parameter τ or by their *PCR* and the fact that CDV is negligible with respect to a given *k*-batch Poisson reference process. CAC can be performed by comparing *PCR* values with respect to a nominal multiplexer rate c' or by comparing an equivalent rate to the actual multiplexer capacity c . The nominal multiplexer rate is defined in 7.1.4. The equivalent rate in the present case is defined as follows:

Equivalent cell rate

To each connection is attributed an Equivalent Cell Rate (ECR) such that if the sum of ECR values of all multiplexed connections is less than the multiplexer output rate c , then the CLR objective is satisfied. The definition of ECR in general is for further study. However, in the case of negligible CDV, the ECR may be seen to be equal to the *PCR* multiplied by the ratio c/c' where c' is the nominal output rate defined in 7.1.4 above.

Example CAC procedure for an ATM link or shaped DBR VPC of rate c and buffer size B .

- It is assumed that connection peak rates are defined with negligible CDV compared to a *k*-batch Poisson reference process (see 7.1.3).
- Estimate the nominal rate c' such that the overflow probability of an $M^{(k)}/D/1/B$ queue would be less than the target value ϵ . $Q(B/k)$ given by Formula 7-4 can be used to estimate the overflow probability provided ϵ is small.
- Admit connections of rates pcr_i while $\sum pcr_i \leq c'$.
- Equivalently, calculate the equivalent cell rate for connection i as $ecr_i = pcr_i \times c/c'$ and admit connections while $\sum ecr_i \leq c$.

Note that mean and maximum queueing delay requirements are assumed to be satisfied for all admissible traffic mixes through the choice of buffer size B and nominal capacity c' .

Example CAC procedure for an uncontrolled constant rate VPC of rate r_{VPC} .

- The VPC and multiplexed VCCs are assumed to have negligible CDV with respect to a common *k*-batch Poisson process.
- Admit connections of rates pcr_i while $\sum pcr_i \leq r_{VPC}$.

Note that there is no need to apply the factor c/c' to evaluate an equivalent rate since this is already taken into account in the CAC of the link applied to the VPC. In this case $ecr_i = pcr_i$.

A variable rate VPC is equivalent to an uncontrolled constant rate VPC in the present case of peak rate allocation and therefore has the same CAC procedure.

8.2 CAC for rate envelope multiplexing

Multiplexing variable bit rate streams using peak rate allocation can lead to inefficient link utilisation. It can be possible to use resources more efficiently while still satisfying GOS objectives by overbooking link bandwidth in statistical multiplexing schemes using REM as discussed in 7.2 above.

The GOS objective $CLR \leq \epsilon$ must be decomposed into two parts, one for the cell scale, $CLR_{cs} \leq \epsilon_{cs}$, and one for the burst scale, $CLR_{bs} \leq \epsilon_{bs}$ with $\epsilon = \epsilon_{cs} + \epsilon_{bs}$. To ensure that $CLR_{cs} \leq \epsilon_{cs}$ when rates are defined with negligible CDV with respect to a *k*-batch Poisson process, it is proposed to compare mean arrival rates to a nominal capacity c' determined such that $Q(B/k) \leq \epsilon_{cs}$ where Q is given by Formula 7-4 with $\rho = c'/c$. The control of the burst-scale component CLR_{bs} relies on being able to estimate the stationary probability distribution of the instantaneous bit rate of multiplexed streams (or, at least, its first moments), either collectively for all existing connections or individually for each connection.

8.2.1 Known cell traffic variables

If the statistical cell traffic variables relating to burst structure, as defined in Recommendation E.716, are known (e.g. for a given source type), these variables can be used in CAC. For example, a speech connection coded according to a given algorithm with silence elimination (i.e. only cells transporting a significant signal are transmitted) can be accurately characterized as an on/off source of given peak rate p and mean rate m . Rate envelope multiplexing of such sources can be performed on an ATM link or a VPC and the burst-scale component CLR_{bs} can be estimated using Formula 7-7.

A mixture of sources of known but different rate distributions can be handled in the same way by calculating CLR_{bs} by Formula 7-5 with expectations calculated using the distribution of Λ_t derived by convolution of the individual distributions. CAC must ensure that a new connection is admitted only if the resulting CLR_{bs} would be less than the target value for all connections.

Equivalent cell rate

CAC is greatly simplified by exploiting known properties of the convolution of rate distributions. In particular, it is possible to attribute to each connection an Equivalent Cell Rate (ECR) such that GOS objectives are met if the sum of ECR values is less than the multiplexer rate c . In other words, a connection i is assigned an ECR of ecr_i such that $CLR_{bs} < \epsilon_{bs}$ while $\sum ecr_i \leq c$. The ECR can be calculated using a rule which depends only on the connection traffic characteristics and static parameters describing the multiplexer and its expected traffic mix. It may alternatively also depend on the traffic characteristics of the other connections using the multiplexer and consequently change dynamically as connections are set up and released. Possible methods for calculating ecr_i are described in Appendix I.

Example CAC procedure for an ATM link or shaped DBR VPC of rate c and buffer size B .

- It is assumed in this example that rates are given with negligible CDV with respect to a k -batch Poisson process, as defined in 7.1.3, and that the mean cell rate of each connection is known; let mcr_i be the mean rate of connection i .
- Targets ϵ_{cs} and ϵ_{bs} are assigned to CLR_{cs} and CLR_{bs} , respectively, such that $\epsilon_{cs} + \epsilon_{bs}$ is less than the CLR GOS objective.
- The multiplexer nominal capacity c' is determined from the link rate c and buffer capacity B such that $Q(B/k) \leq \epsilon_{cs}$ where $Q(B/k)$ is the buffer saturation probability for an $M^{(k)}/D/1$ queue estimated by Formula 7-4 for a load of $\rho = c'/c$.
- Calculate the equivalent cell rate ecr_i of connection i according to one of the methods described in Appendix I; depending on the definition of equivalent cell rate it may be necessary to calculate ecr_i once and for all as the connection is requested or to re-evaluate it as the occupancy state changes.
- Admit connections while $\sum mcr_i \leq c'$ and $\sum ecr_i \leq c$.

Example CAC procedure for an uncontrolled constant rate VPC of rate r_{VPC} .

- The VPC and multiplexed connections are assumed to have negligible CDV with respect to a common k -batch Poisson process.
- Evaluate an equivalent cell rate ecr_i for each connection i (e.g. by using one of the methods of Appendix I) with, however, r_{VPC} in place of c .
- Admit connections while $\sum ecr_i \leq r_{VPC}$.

The condition on the sum of VCC mean rates occurring in the previous example is not applicable here since cell scale congestion is accounted for in the CAC of the link applied to the VPC. Note that in this case CLR_{bs} is not strictly a cell loss ratio but the fraction of cells which violate the rate r_{VPC} declared for the VPC. Thus, the target ϵ_{bs} for CLR_{bs} must here be negligible compared to the target CLR of the ATM links of the VPC path to avoid deteriorating the performance of other connections sharing these links.

Example CAC procedure for a variable rate VPC.

Variable rate VPCs provide for more efficient multiplexing at the cost of more complex CAC taking account individually of the ATM links over which the VPC is routed. CAC is facilitated when the variable rate VPC is characterized by a set of equivalent cell rates as discussed in Recommendation E.735.

- All multiplexed connections on the considered links are assumed to have rates defined with negligible CDV with respect to a common k -batch Poisson process.
- Each ATM link j of the VPC is characterized in the originating VC-node by the following set of parameters:
 - ECR^j , an equivalent cell rate determined off-line by a network dimensioning procedure taking account of the required capacity of the VPC and the rate and expected traffic mix of link j , as discussed in Recommendation E.735;
 - the link rate and expected traffic mix parameters necessary to allow computation of equivalent cell rates of the VCCs to be multiplexed in the VPC (the same parameters used to compute the equivalent cell rate of all VCCs and VPCs handled by the ATM link).
- The VPC is also characterized by a Mean Cell Rate (MCR). Let mcr_i be the mean cell rate of VCC $_i$.
- For every VCC $_i$ to be handled within the VPC, evaluate an equivalent cell rate ecr_i^j for each link j using the link characteristics specified above.
- Accept a new connection while $\sum_i ecr_i^j \leq ECR^j$ for every link j and $\sum_i mcr_i \leq MCR$.

In the present case of a variable rate VPC, it is not envisaged that the equivalent cell rates should depend on the actual occupancy states of the links. For the equivalent cell rate methods detailed in Appendix I, method 1 does not use a traffic mix parameter while methods 2 and 3 use a single parameter denoted α . Parameters describing the second and subsequent links of the VPC must be communicated to the originating VC-node via the management plane when the VPC is established or modified.

8.2.2 Worst-case resource allocation

Performance objectives can be met by allocating resources assuming "worst-case"¹ traffic characteristics compatible with declared values of traffic parameters. If only PCR is declared, CAC is then equivalent to peak rate allocation (see 8.1 above).

If both PCR and SCR are known, a worst-case allocation can be derived assuming the source is of on/off type with peak rate PCR and mean rate SCR . The procedures outlined in 8.2.1 can then be followed for CAC. Note that the IBT parameter does not affect CAC when REM is employed.

8.2.3 Adaptive CAC

The CAC decision depends on the traffic characteristics of the connection requesting admission and the combined characteristics of all existing connections. Greater efficiency than worst-case resource allocation can be obtained if the latter combined characteristics can be accurately estimated by real-time traffic measurements. The observed traffic characteristics should allow an estimation of the resulting CLR if the new connection with worst-case traffic characteristics were added to the existing connections: the connection would not be admitted if the CLR estimate were greater than that required. Some work in this area has been reported in the literature [SaS 91], [DMM 94], [KRK 94], [RMV96] but further work is required before recommending any particular method.

¹ By worst case, we imply those traffic characteristics compatible with the declared traffic descriptor requiring the greatest resource allocation to meet QOS requirements.

8.3 CAC for rate sharing statistical multiplexing

As discussed in 7.3, to achieve high-link utilization when multiplexing connections whose peak bit rate is not a small fraction of the multiplexer bit rate requires a large buffer to absorb the cells arriving during momentary overload periods. This type of multiplexing has been termed Rate Sharing (RS). RS may also be used for connections with low peak rates; the essential difference with REM is the reliance on a large buffer to absorb input rate overloads which occur with non-negligible probability. With RS, it may be necessary to perform CAC by allocating a quantity of both bandwidth and buffer space to each connection.

As for REM, three possibilities for estimating the traffic characteristics necessary for predicting multiplexer performance can be distinguished.

8.3.1 Known cell traffic variables

If all necessary traffic variables² of multiplexed connections can be deduced from the fact that each belongs to a given source type, CAC can be performed with reference to a mathematical model for predicting multiplexer performance (i.e. cell loss ratio and mean and maximum queueing delay). The definition of such a model is for further study.

Equivalent cell rate

As for REM, CAC may be simplified in certain cases when it is possible to attribute to each connection an Equivalent Cell Rate (ECR) depending on the multiplexer rate, buffering capacity, and the connection's own intrinsic properties. The CAC procedure then consists in accepting connections until the sum of ECR values would be greater than the multiplexer rate. The precise definition of ECR and its field of application in the present context is for further study.

8.3.2 Worst-case resource allocation

As for REM, GOS objectives can be guaranteed if resources are allocated for the "worst-case" traffic corresponding to the declared traffic descriptors. The worst-case traffic when only the *PCR* is declared is a CBR stream and CAC is equivalent to peak rate allocation (see 8.1).

If *PCR*, *SCR* and *IBT* parameters are declared, a candidate worst-case traffic is an on/off source with maximal length burst and silence periods (see 6.1.2). The definition of CAC rules based on such a worst-case model is for further study.

8.3.3 Adaptive CAC

Traffic measurements (e.g. observations of buffer occupancy) may be used to derive an estimate of the capacity to accept new connections without infringing performance objectives. The definition of such measurements and the corresponding CAC procedures is for further study.

9 Adaptive resource management

Resource sharing efficiency can be improved by employing dynamically adaptive resource management, especially when it is not possible to declare all connection traffic characteristics during the set-up. ABR and ABT transfer service capabilities both rely on the use of Resource Management (RM) cells to adjust resource allocation during the lifetime of a connection. In ABR, it is the network which determines the bit rate which is available to a given connection and informs the user. The ABT service, on the other hand, is based on the user initiating requests for resource allocation changes.

² By necessary traffic variables is meant all the parameters of the cell arrival process which have a significant influence on multiplexer performance (see Recommendation E.716).

9.1 ABT resource management

In this subclause, we consider the possibility of a user re-negotiating resource allocation during the lifetime of a connection to more closely correspond to the user's current traffic demand.

9.1.1 REM and block transfer

When multiplexers are equipped with small buffers dimensioned as discussed in 7.2 for REM, any prolonged rate overload (combined instantaneous bit rate greater than link nominal rate) leads to cell loss for all connections. This can have serious consequences if, for example, cell loss leads to the retransmission of affected Protocol Data Units (PDUs). This possibility is avoided with peak rate allocation as discussed in 8.1. One way of preserving this desirable property while realizing the advantages of statistical multiplexing is to require connections to dynamically reserve sufficient bandwidth for their current instantaneous requirement. For example, a user intermittently emitting bursts of cells at some given bit rate would reserve bandwidth at that rate at the start of a burst and relinquish it at the end of the burst.

CAC would be performed largely as described in 8.2 above although the connection admission criteria could be different bearing in mind the less severe consequences of blocking a requested bandwidth increase (affecting one user only) compared to the generalized cell loss in a system without bandwidth reservation at this level. It would be appropriate to specify QOS criteria explicitly related to the probability of RM request failure.

9.1.2 Rate sharing and block transfer

User initiated RM procedures can be used with rate sharing statistical multiplexing to more closely match the amount of resources allocated to a connection to its current level of activity. Both bandwidth and buffer space could be allocated dynamically.

For connections with a peak bit rate close to the multiplexer rate, it seems inappropriate to seek to reserve bandwidth equal to this peak rate for each burst of activity. To attempt to do this with a suitably low probability of burst blocking would require the multiplexer to operate with a very low mean utilization. This follows from considerations similar to those presented in 7.2.2. Buffer allocation can also be adapted to the current requirements of an established connection. It is only necessary to reserve buffer space when the connection is actually transmitting data.

9.2 ABR resource management

The ABR transfer capability is designed primarily for rate sharing statistical multiplexing. Under rate sharing, queue lengths in buffers can be large and have a sensitive dependence on attributes of arrival processes, which does not occur with REM. Without any feedback mechanism to the sources, it is difficult to engineer buffer capacity and bandwidth to satisfy a cell loss criterion. With ABR, sources receive ongoing feedback of the rate that can currently be supported. When buffers are congested, (selected) sources can be told to decrease their current allowed cell rate.

The appropriate dimensioning and CAC for ABR is dependent on the nature of the service the network operator intends to offer with ABR. For example, based on business decisions that are beyond the scope of this Recommendation, an operator might determine a given amount of bandwidth to dedicate for the ABR-based service; the operator might choose to admit all connection requests, or might set an upper limit. As another example, the operator might offer an ABR-based service with a commitment that when a connection is active it will receive at least a given cell rate with a given probability.

Users of the ABR transfer capability have the option of requesting a Minimum Cell Rate (MCR). If granted, the network commits to provide at least this rate to the source for the duration of the connection. Thus, the portion of the CAC that handles the minimum cell rate of the ABR connections can be similar to that of the DBR transfer capability, see clause 8.

Traffic engineering implications of the ABR transfer service capability are for further study.

10 Service integration

If all cells in multiplexer queues are served in First In First Out (FIFO) order, the most severe cell transfer delay and cell delay variation requirements of the connection types to be multiplexed determine the maximum buffer size. In particular, if services with real-time response requirements like interactive speech are to be handled, it does not appear possible to perform statistical multiplexing with large buffers as discussed in 7.3 above unless some more sophisticated service discipline is employed.

10.1 Dedicated resources

To satisfy the different GOS requirements of different service classes, specific resources may be dedicated to groups of services having similar requirements. In particular, distinct ATM links can be used or distinct shaped DBR virtual path connections can be created, their bandwidth being adapted to the expected demand for the given group of services. Buffer space reserved for a given VPC would also be chosen appropriately (e.g. small buffers for services with real-time response requirements, large buffers for high rate delay tolerant services). Note that unshaped VPCs cannot be used for this purpose; to provide different GOS to different connections using uncontrolled VPCs requires the implementation of priority or scheduling mechanisms as discussed below.

10.2 Loss priorities

The loss priority mechanisms discussed in 7.2.3 and 7.3.2 can be used to differentiate the *CLR* offered to cells of a given connection according to the value of the CLP bit or to offer different cell loss ratios to different connections. The loss priority mechanisms can be combined with delay priority mechanisms.

The definition of traffic engineering rules to provide precise QOS guarantees is for further study.

10.3 Delay priorities

Head of line priority can be given to certain traffic streams, notably to reduce the waiting time of their cells in multiplexers equipped with large buffers. Several priority levels might be defined, the level to which a given stream belongs being identified by the VPI/VCI field of the cell header. Head of line and loss priority using the CLP bit (see 7.2.3 and 7.3.2) can be mixed.

10.4 Scheduling policies

Head of line priority can be used to ensure minimal delay for services with real-time response time constraints. More general service discrimination can be achieved at the price of more sophisticated queue scheduling disciplines such as weighted fair queueing [RMV96]. The impact of such disciplines on multiplexer performance and traffic control is for further study.

11 Bibliography

- [CLW94] CHOUDURY (G.L.), LUCANTONI (D.M.), WHITT (W.): On the effectiveness of effective bandwidth for admission control in ATM networks, *Proceedings of ITC 14*, June 1994.
- [DMM 94] DZIONG (Z.), MONTANUY (O.), MASON (L.G.): Adaptive traffic admission in ATM networks – Optimal estimation framework, *Proceedings of ITC 14*, Juan-les-Pins, June 1994.
- [KRK 94] KRÖNER (H.), RENGER (T.), KNOBLING (R.): Performance modelling of an adaptive CAC strategy for ATM networks, *Proceedings of ITC 14*, Juan-les-Pins, June 1994.

- [LuP90] LUCANTONI (D.M.), PAREKH (S.): Selective cell discarding mechanisms for a B-ISDN congestion control architecture, *ITC Seminar*, Morristown, 1990.
- [RMV96] ROBERTS (J.), MOCCI (U.), VIRTAMO (J.), (Eds): Broadband Network Teletraffic, *Lecture Notes in Computer Science*, Vol. 1155, Springer-Verlag, 1996.
- [SaS 91] SAITO (H.), SHIOMOTO (K.): Dynamic Call Admission Control in ATM Networks, *IEEE JSAC*, Vol. 9, No. 2, September 1991.
- [ViS 97] VILLÉN-ALTAMIRANO (M.), SÁNCHEZ-CAÑABATE (MF.): Effective bandwidth dependent of the actual traffic mix: an approach for bufferless CAC, *ITC 15*, June 1997.

12 History

This is a new Recommendation.

Appendix I

Example methods for calculation of equivalent cell rate for rate envelope multiplexing

The following three methods for the calculation of the Equivalent Cell Rate (ECR) of a connection, ecr_i , are designed for the case of rate envelope multiplexing. All three methods model the cell flow as a fluid and use as input the rate of the ATM link (or VPC) and the parameter ϵ of the performance criterion $CLR_{bs} < \epsilon$.

As discussed in Recommendation E.735, ECR depends in general on the traffic characteristics of the other connections multiplexed on the same ATM link or VPC. In some cases this dependence can be accounted for through a single parameter, denoted α in Recommendation E.735. This parameter may be evaluated for the actual traffic or, to simplify CAC, for a representative set of connections \mathbf{R} . The set \mathbf{R} is such that the numbers of connections of different types are in proportion to those of an expected traffic mix and the CLR_{bs} objective ϵ is attained (i.e. the representative set of connections is situated on the frontier of the admissible region). The use of a parameter α is a characteristic of methods 2 and 3 below. For both methods, although α is defined differently, it is known that the value of ECR does not depend critically on the precise set of connections considered and that the estimation is always conservative, i.e. the calculated ECR is greater than that determined for the actual set of connections.

The first method assumes on/off sources whose mean and peak rates are chosen to match the mean and variance of the instantaneous rate of the source. With this method, the computation of the effective bandwidth of a connection does not depend on the traffic variables of any other connection. This method is the simplest of the three but is the least accurate and does not ensure that the results are conservative (i.e. that the performance criterion is indeed satisfied when $\sum ecr_i \leq c$).

The second method directly uses the distribution of the instantaneous rate of the cell flow at an arbitrary time. Based on the Chernoff bound, this method computes a parameter α which takes account of the traffic characteristics of the actual or a representative set of connections. In the latter case, α can be computed off-line, so that the method can be used in real-time for CAC. Computation of α in real-time for the actual traffic does not seem feasible in view of the complexity of the calculations involved.

The third method assumes on/off sources and relies on the computation of a parameter α that depends on characteristics of the actual or a representative set of connections. The algorithm to determine α for the actual traffic is simpler than that of method 2 and can plausibly be applied in real-time.

Methods 2 and 3 were initially proposed to evaluate an equivalent cell rate to be used in a CAC algorithm ensuring that the probability of the arrival rate Λ_t exceeding the service rate c is less than a target value. The methods are adapted here for CAC algorithms based on *CLR* using the simple order of magnitude relation $CLR \approx Pr\{\Lambda_t > c\} \times 100$ (see [RMV96] page 446).

I.1 Equivalent cell rate method 1 [RMV96]

- For a source of given mean rate m_i and variance σ_i^2 , consider an equivalent on/off source of peak rate h_i such that $\sigma_i^2 = m_i(h_i - m_i)$
- Derive the "effective bandwidth" eb_i as follows:

$$eb_i = \begin{cases} am_i(1 + 3z_i(1 - m_i/h_i)), & \text{for } 3z_i \leq \min(3, h_i/m_i) \\ am_i(1 + 3z_i^2(1 - m_i/h_i)), & \text{for } 3 < 3z_i^2 \leq h_i/m_i \\ ah_i, & \text{otherwise} \end{cases} \quad (\text{I-1})$$

where $a = 1 - \frac{\log_{10}\epsilon}{50}$ and $z_i = \frac{-2 \log_{10}\epsilon}{c/h}$.

- The equivalent cell rate for $CLR_{bs} \leq \epsilon$ is $ecr_i = eb_i$.

I.2 Equivalent cell rate method 2 [RMV96]

- Let $\lambda_i(t)$ be the rate of source i at time t and derive the log moment generating function of its distribution:

$$M_i(s) = \log_e E \left\{ e^{s\lambda_i(t)} \right\}.$$

- Determine α such that the function $\sum_{i \in S} M_i(s) - sc$ is minimized at $s = \alpha$.

α can be evaluated for two definitions of the set of connections S considered in the summations:

- the actual set of connections currently handled by the link or VPC (method 2a);
- the representative set of connections R described above (method 2b).

- Calculate the "effective bandwidth" of source i as $eb_i = M_i(\alpha) / \alpha$.
- Admitting connections while $\sum_i eb_i \leq c - \gamma / \alpha$ ensures $Pr\{\sum \lambda_i(t) > c\} \leq e^{-\gamma}$.
- Use the order of magnitude approximation $CLR_{bs} \approx Pr\{\sum \lambda_i(t) > c\} / 100$ to derive the equivalent cell rate for $CLR_{bs} \leq \epsilon$:

$$ecr_i = \frac{eb_i}{c + \log_e(100\epsilon) / \alpha} \times c \quad (\text{I-2})$$

I.3 Equivalent cell rate method 3 [ViS 97]

- For a source of given mean rate m_i and variance σ_i^2 , consider an equivalent on/off source of peak rate h_i such that $\sigma_i^2 = m_i(h_i - m_i)$.

$$\gamma_i = -\log_e \left(\frac{m_i}{h_i} \right).$$

- Let $e^{-\gamma}$ be a target overload probability and define

- Derive α and M as follows:

$$\alpha = \frac{g - \sum_{i \in \mathbf{P}} g_i}{c - \sum_{i \in \mathbf{P}} h_i - \sum_{i \notin \mathbf{P}} m_i}; \quad M = c - \frac{\gamma}{\alpha} \quad (\text{I-3})$$

where \mathbf{P} is the set of sources satisfying: $\gamma_i/(h_i - m_i) < \alpha$.

- As for method 2, α can be evaluated for two definitions of the set of connections considered in the summations:
 - the actual set of connections currently handled by the link or VPC (method 3a);
 - the representative set of connections \mathbf{R} described above (method 3b).
- Calculate the effective bandwidth, eb_i , and the modified mean rate, m'_i ; of source i as:

$$eb_i = \begin{cases} h_i & \text{if } i \in \mathbf{P} \\ h_i \sum_{k=0}^3 a_k^{(p_i)} (\alpha h_i)^k & \text{if } i \notin \mathbf{P} \end{cases}$$

$$m'_i = \begin{cases} h_i - \frac{\gamma_i}{\alpha} & \text{if } i \in \mathbf{P} \\ m_i & \text{if } i \notin \mathbf{P} \end{cases} \quad (\text{I-4})$$

where:

$$p_i = m_i / h_i$$

$$a_0^{(p_i)} = p_i$$

$$a_1^{(p_i)} = 1,7 p_i (1 - p_i)$$

$$a_2^{(p_i)} = \begin{cases} p_i (1 - p_i) (1,72 - 4,06 p_i + 1,55 p_i^2) & \text{if } p_i \geq 0,04 \\ p_i (1 - p_i) (0,9 + 16,6 p_i) & \text{if } p_i < 0,04 \end{cases}$$

$$a_3^{(p_i)} = \frac{(1 - p_i)^4 - a_1^{(p_i)} (1 - p_i)^2 \left| \log p_i \right| - a_2^{(p_i)} (1 - p_i) \left| \log p_i \right|^2}{\left| \log p_i \right|^3} \quad (\text{I-5})$$

- Admitting connections while:

$$\sum_I (eb_i + m'_i) \leq c + M \quad (\text{I-6})$$

ensures a probability of overload less than $e^{-\gamma}$. Note that with method 3a, $\sum_I m'_i = M$, and thus the above condition could be written as $\sum_I eb_i \leq c$.

- The ECR for a $CLR_{bs} \leq \varepsilon \approx e^{-\gamma}/100$ is:

$$ecr_i = \frac{c}{c + M} (eb_i + m'_i) \quad (\text{I-7})$$

ITU-T RECOMMENDATIONS SERIES

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