



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

I.371

(03/93)

**INTEGRATED SERVICES DIGITAL
NETWORK (ISDN)**

OVERALL NETWORK ASPECTS AND FUNCTIONS

**TRAFFIC CONTROL AND CONGESTION
CONTROL IN B-ISDN**

ITU-T Recommendation I.371

(Previously "CCITT Recommendation")

FOREWORD

The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of the International Telecommunication Union. 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, established the topics for study by the ITU-T Study Groups which, in their turn, produce Recommendations on these topics.

ITU-T Recommendation I.371 was prepared by the ITU-T Study Group XVIII (1988-1993) and was approved by the WTSC (Helsinki, March 1-12, 1993).

NOTES

1 As a consequence of a reform process within the International Telecommunication Union (ITU), the CCITT ceased to exist as of 28 February 1993. In its place, the ITU Telecommunication Standardization Sector (ITU-T) was created as of 1 March 1993. Similarly, in this reform process, the CCIR and the IFRB have been replaced by the Radiocommunication Sector.

In order not to delay publication of this Recommendation, no change has been made in the text to references containing the acronyms "CCITT, CCIR, or IFRB" or their associated entities such as Plenary Assembly, Secretariat, etc. Future editions of this Recommendation will contain the proper terminology related to the new ITU structure.

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

© ITU 1993

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 Introduction	1
1.1 General objectives	1
1.2 Generic functions.....	2
1.3 A Reference configuration for traffic control and congestion control.....	2
1.4 Events, actions, time-scales and response times	2
1.5 Quality of Service, network performance and cell loss priority	2
2 Traffic descriptors and parameters	4
2.1 Definitions	4
2.2 Requirements	5
2.3 User-network traffic contract.....	5
2.4 Traffic parameter specifications	6
3 Functions and procedures for traffic control and congestion control.....	9
3.1 Introduction	9
3.2 Traffic control functions	10
3.3 Congestion control functions	19
Annex A – Peak cell rate cell monitor algorithms accounting for cell delay variation tolerance.....	20
Annex B – List of Abbreviations.....	22
Appendix I – Examples of application of the equivalent terminal for the peak cell rate definition	23

TRAFFIC CONTROL AND CONGESTION CONTROL IN B-ISDN

(Helsinki, 1993)

1 Introduction

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

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

- The main body describes the objectives and mechanisms of traffic control and congestion control.
- Examples of application of the concepts are in annexes.

In B-ISDN, congestion is defined as a state of network elements (e.g. switches, concentrators, cross-connects and transmission links) in which the network is not able to meet the negotiated network performance objectives for the already established connections and/or for the new connection requests.

In general, congestion can be caused by

- unpredictable statistical fluctuations of traffic flows;
- fault conditions within the network.

Congestion is to be distinguished from the state where buffer overflow is causing cell losses, but still meets the negotiated Quality of Service.

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

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

1.1 General objectives

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

The uncertainties of broadband traffic patterns, traffic control and congestion control complexity suggest a step-wise approach for defining traffic parameters and network traffic control and congestion control mechanisms. This Recommendation defines a restricted initial set of traffic control and congestion control capabilities aiming at simple mechanisms and realistic network efficiency.

It may subsequently be appropriate to consider additional sets of such capabilities, for which additional traffic control mechanisms will be used to achieve increased network efficiency.

The objectives of ATM layer traffic control and congestion control for B-ISDN are as follows:

- ATM layer traffic control and congestion control should support a set of ATM layer Quality of Service (QOS) classes sufficient for all foreseeable B-ISDN services; the specification of these QOS classes should be consistent with network performance at present under study.

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

1.2 Generic functions

To meet these objectives, the following functions form a framework for managing and controlling traffic and congestion in ATM networks and may be used in appropriate combinations.

- Network resource management (NRM): Provisioning may be used to allocate network resources in order to separate traffic flows according to service characteristics.
- Connection admission control (CAC) is defined as the set of actions taken by the network during the call set up phase (or during call re-negotiation phase) in order to establish whether a virtual channel/virtual path connection request can be accepted or rejected (or whether a request for re-allocation can be accommodated). Routing is part of connection admission control actions.
- Feedback controls are defined as the set of actions taken by the network and by the users to regulate the traffic submitted on ATM connections according to the state of network elements.
- Usage/network parameter control (UPC/NPC) is defined as the set of actions taken by the network to monitor and control traffic, in terms of traffic offered and validity of the ATM connection, at the user access and the network access, respectively. Their main purpose is to protect network resources from malicious as well as unintentional misbehaviour which can affect the QOS of other already established connections by detecting violations of negotiated parameters and taking appropriate actions.
- Priority control: The user may generate different priority traffic flows by using the cell loss priority bit (see Recommendation I.150). A congested network element may selectively discard cells with low priority if necessary to protect as far as possible the network performance for cells with high priority.
- Other control functions are for further study.

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

1.3 A Reference configuration for traffic control and congestion control

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

1.4 Events, actions, time-scales and response times

Figure 2 illustrates the time-scales over which various traffic control and congestion control functions operate. The response time defines how quickly the controls react. For example, cell discarding can react on the order of the insertion time of a cell. Similarly, feedback controls can react on the time scale of round-trip propagation times. Since traffic control and resource management functions are needed at different time-scales, no single function is likely to be sufficient.

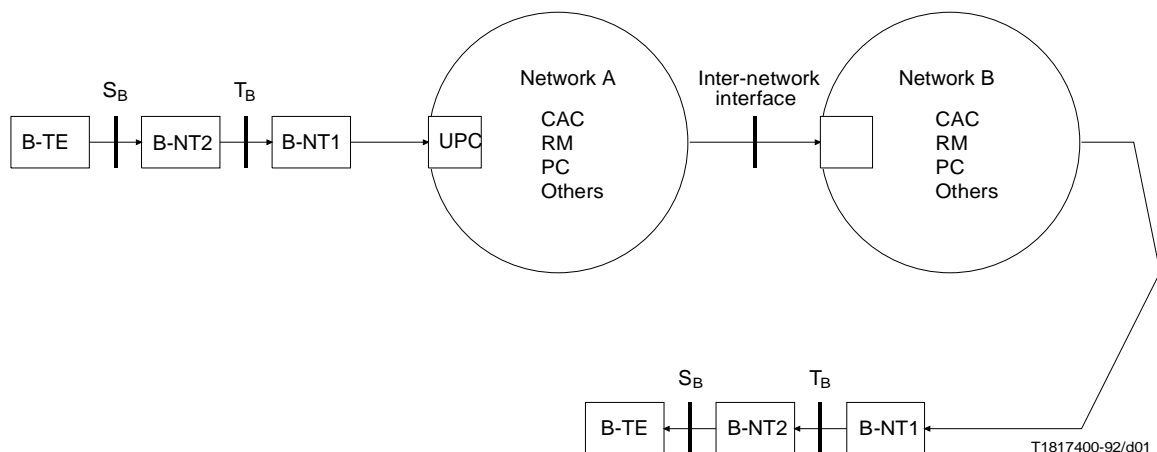
1.5 Quality of Service, network performance and cell loss priority

The ATM layer Quality of Service is defined by a set of parameters such as delay and delay variation sensitivity, cell loss ratio, etc. Other QOS parameters are for further study.

A user requests a specific ATM layer QOS from the QOS classes which a network provides. This is part of the traffic contract at connection establishment (see 2.3.1). It is a commitment for the network to meet the requested Quality of Service as long as the user complies with the traffic contract. If the user violates the traffic contract, the network need not respect the agreed QOS.

A user may request at most two different QOS classes for a single ATM connection, which differ with respect to the cell loss ratio objectives. The cell loss priority bit of the ATM cell header allows for two cell loss ratio objectives for a given ATM connection.

Network performance objectives at the ATM SAP are intended to capture the network ability to meet the requested ATM layer Quality of Service. It is the role of the upper layers, including the AAL, to translate this ATM layer QOS to any specific application requested QOS.

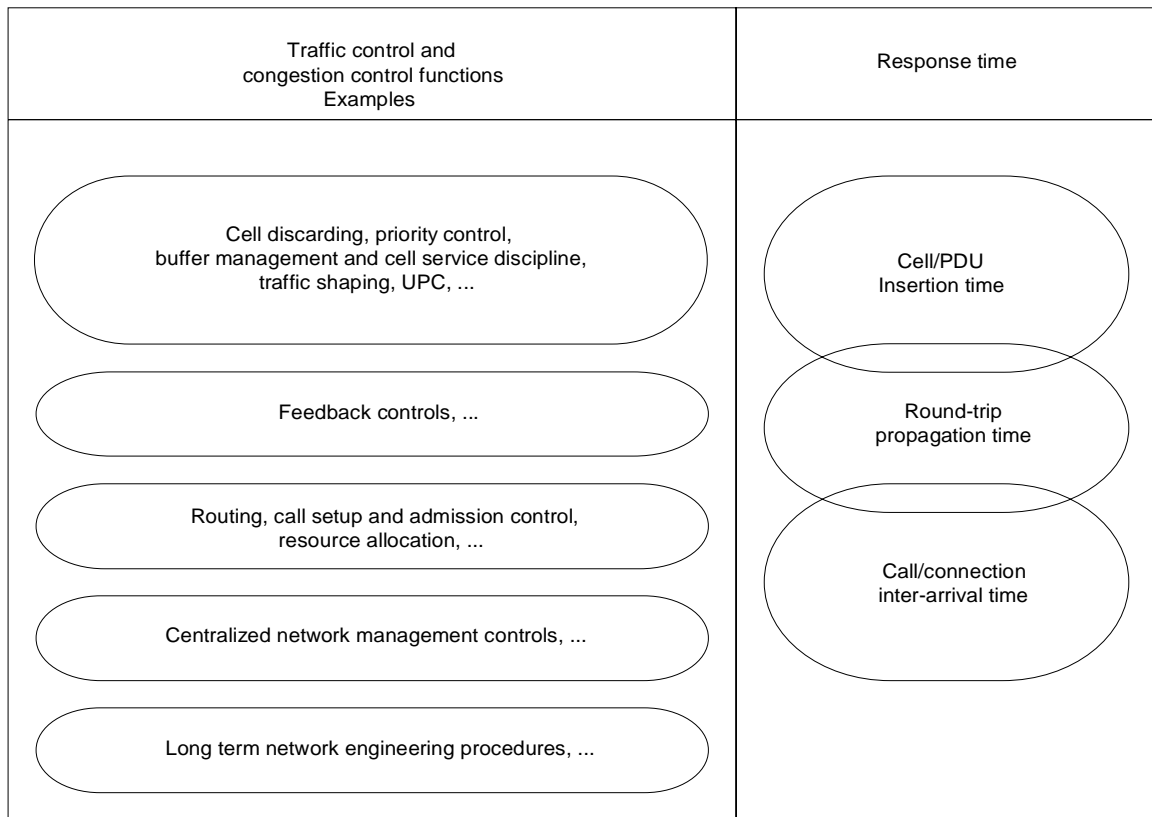


- UPC Usage parameter control
- CAC Connection admission control
- PC Priority control
- NPC Network parameter control
- RM Resource management
- Others For further study

NOTES

- 1 NPC may apply as well at some intra-network NNIs.
- 2 The arrows are indicating the direction of the cell flow.

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



T1817410-92/d02

FIGURE 2/I.371
Control response times

2 Traffic descriptors and parameters

Traffic parameters describe traffic characteristics of an ATM connection. Traffic parameters are grouped into source traffic descriptors for exchanging information between the user and the network.

Connection admission control procedures will use source traffic descriptors to allocate resources and derive parameters for the operation of UPC/NPC.

2.1 Definitions

For the purpose of this Recommendation, the following definitions apply:

2.1.1 traffic parameters: A traffic parameter is a specification of a particular traffic aspect. It may be qualitative or quantitative.

Traffic parameters may for example describe peak cell rate, average cell rate, burstiness, peak duration and source type (e.g. telephone, videophone).

Only peak cell rate is presently defined in this Recommendation.

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

2.1.2 traffic descriptors: The ATM traffic descriptor is the generic list of traffic parameters which can be used to capture the intrinsic traffic characteristics of an ATM connection.

Introduction of additional parameters to enhance the network resource management procedures or to capture traffic characteristics of a new type of connection is left open for further study.

A description of the characteristics of the traffic that any given requested connection may offer has to be provided by the user at the connection set-up phase.

A source traffic descriptor is the set of traffic parameters belonging to the ATM traffic descriptor used during the connection set-up to capture the intrinsic traffic characteristics of the connection requested by the source.

2.2 Requirements

Any traffic parameter to be involved in a source traffic descriptor should:

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

These criteria should be respected since users may have to provide these traffic parameters at connection set-up. In addition, these traffic parameters should be useful to the CAC procedure so that network performance objectives can be maintained once the connection has been accepted. Finally, they should be enforceable by the UPC/NPC to maintain network performance in case of non-compliant usage.

2.3 User-network traffic contract

2.3.1 traffic contract definition: CAC and UPC/NPC procedures require the knowledge of certain parameters to operate efficiently: they should take into account the source traffic descriptor, the requested QOS and the CDV tolerance (see 2.4) in order to decide whether the requested connection can be accepted.

The source traffic descriptor, the requested QOS for any given ATM connection and the maximum CDV tolerance allocated to the CEQ define the traffic contract at the T_B reference point. Source traffic descriptors and QOS are declared by the user at connection set-up by means of signalling or subscription. Whether the maximum allowable cell delay variation tolerance is also negotiated on a subscription or on a per connection basis is for further study.

The connection admission control (CAC) and usage/network parameter control (UPC/NPC) procedures are operator specific. Once the connection has been accepted, the value of the CAC and UPC/NPC parameters are set by the network on the basis of the network operator's policy.

NOTE – All ATM connections handled by network connection related functions (CRF) have to be declared and enforced by the UPC/NPC. ATM layer QOS can only be assured for compliant ATM connections. As an example, individual VCCs inside a user end-to-end VPC are neither declared nor enforced at the UPC and hence no ATM layer QOS can be assured to them.

2.3.2 Source traffic descriptors, Quality of Service and cell loss priority

If a user requests two levels of priority for an ATM connection, as indicated by the CLP bit value, the intrinsic traffic characteristics of both cell flow components have to be characterized in the source traffic descriptor. This is by means of a set of traffic parameters associated with the $CLP = 0$ component and a set of traffic parameters associated with the $CLP = 0 + 1$ component.

As indicated in subclause 1.5, the network provides an ATM layer QOS for each of the components ($CLP = 0$ and $CLP = 0 + 1$) of an ATM connection. The traffic contract specifies the particular QOS choice (from those offered by the network operator) for each of the ATM connection components. There may be a limited offering of QOS specifications for the $CLP = 1$ component.

Cell loss ratio objectives are for further study.

2.3.3 Impact of cell delay variation on UPC/NPC and resource allocation

ATM layer functions (e.g. cell multiplexing) may alter the traffic characteristics of ATM connections by introducing cell delay variation as illustrated in Figure 3. When cells from two or more ATM connections are multiplexed, cells of a given ATM connection may be delayed while cells of another ATM connection are being inserted at the output of the multiplexer. Similarly, some cells may be delayed while physical layer overhead or OAM cells are inserted. Therefore, some randomness affects the time interval between reception of ATM cell data-requests at the end-point of an ATM connection to the time that an ATM cell data-indication is received at the UPC/NPC. Besides, AAL multiplexing may originate CDV.

The UPC/NPC mechanism should not discard or tag cells in an ATM connection if the source conforms to the source traffic descriptor negotiated at connection establishment. However, if the CDV is not bounded at a point where the UPC/NPC function is performed, it is not possible to design a suitable UPC/NPC mechanism and to allocate resources properly. Therefore, it is required that a maximum allowable value of CDV be standardized edge-to-edge, e.g. between the ATM connection end-point and T_B , between T_B and an Inter-Network Interface and between Inter-Network Interfaces (see Figure 1).

Standardization of a number of CDV tolerance values less than the maximum allowable value of CDV tolerance to apply to certain interfaces, e.g. on a subscription basis or on a per connection basis, is for further study.

UPC/NPC should accommodate the effect of the maximum CDV allowed on ATM connections within the limit resulting from the accumulated CDV allocated to upstream subnetworks (including CEQ).

Traffic shaping partially compensates for the effects of CDV on the peak cell rate of the ATM connection. Examples of traffic shaping mechanisms are re-spacing cells of individual ATM connections according to their peak cell rate or suitable queue service schemes.

Values of the cell delay variation are network performance issues.

The definition of a source traffic descriptor and the standardization of a maximum allowable CDV may not be sufficient for a network to allocate resources properly. When allocating resources, the network should take into account the worst case traffic passing through UPC/NPC in order to avoid impairments to other ATM connections. This worst case traffic depends on the specific implementation of the UPC/NPC. The trade-offs between UPC/NPC complexity, worst case traffic and optimization of network resources are made at the discretion of network operators. The quantity of available network resources and the network performance to be provided for meeting QOS requirements can influence these trade-offs.

2.4 Traffic parameter specifications

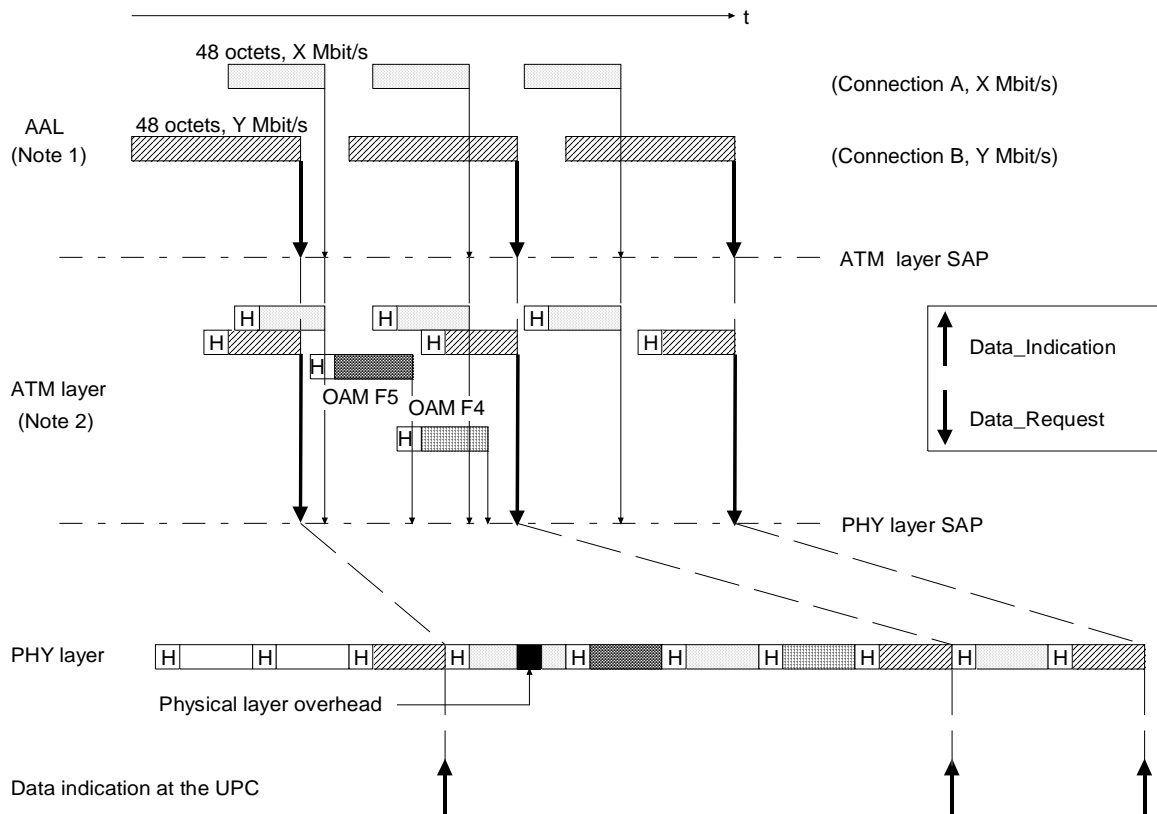
Peak cell rate is a mandatory traffic parameter to be explicitly or implicitly declared in any source traffic descriptor. In addition to the peak cell rate of an ATM connection, it is mandatory for the user to declare either explicitly or implicitly the cell delay variation tolerance τ within the relevant traffic contract.

Additional standardized parameters beyond peak cell rate which may be specified in the future should provide for a significant improvement of network utilization.

2.4.1 Peak cell rate

The following definition applies to ATM connections supporting both CBR and VBR services.

The peak cell rate in the source traffic descriptor specifies an upper bound on the traffic that can be submitted on an ATM connection. Enforcement of this bound by the UPC/NPC allows the network operator to allocate sufficient resources to ensure that the performance objectives (e.g. for cell loss ratio) can be achieved.



T1817420-92/d03

NOTES

- 1 ATM SDUs are accumulated at the upper layer service bit rate. Besides, CDV may also originate in AAL multiplexing.
- 2 GFC delay and delay variation is part of the delay and delay variation introduced by the ATM layer.
- 3 CDV may also be introduced by the network because of random queuing delays which are experienced by each cell in concentrators, switches and cross-connects.

FIGURE 3/I.371
Origins of cell delay variation

2.4.1.1 Peak cell rate definition for a VPC/VCC

Location:

At the physical layer SAP for an equivalent terminal representing the VPC/VCC (this is only a reference configuration; see Figure 4).

Basic event:

Request to send an ATM_PDU in the equivalent terminal.

Definition:

The peak cell rate of the ATM connection is the inverse of the minimum inter-arrival time T between two basic events defined above. T is the peak emission interval of the ATM connection.

The source traffic descriptor of the ATM connection currently reduces to the peak cell rate defined above.

It is noted that conformance control of the peak cell rate by UPC/NPC requires that the CDV tolerance τ allocated to the upstream portion of the ATM connection be specified (see 2.3.1). Whether additional parameters may be useful is for further study.

On a terminal with a single AAL and without ATM layer OAM flows, location and basic event are equivalent to the following ones

Location :

At the ATM layer SAP.

Basic event :

Request to send an ATM_SDU.

In order to properly allocate resources to a VPC/VCC, a peak cell rate as defined above has to be defined for each component of the ATM connection, i.e. the $CLP = 0$ substream (not including the OAM), the aggregate ($CLP = 0 + 1$) substream and the OAM substream. The CDV tolerance accounts for delay variation that will be present in respective cell substreams of the ATM connection. Their values and interpretation are defined by algorithms described in Annex A.

Examples of application of the peak cell rate definition to specific configurations are in Appendix I.

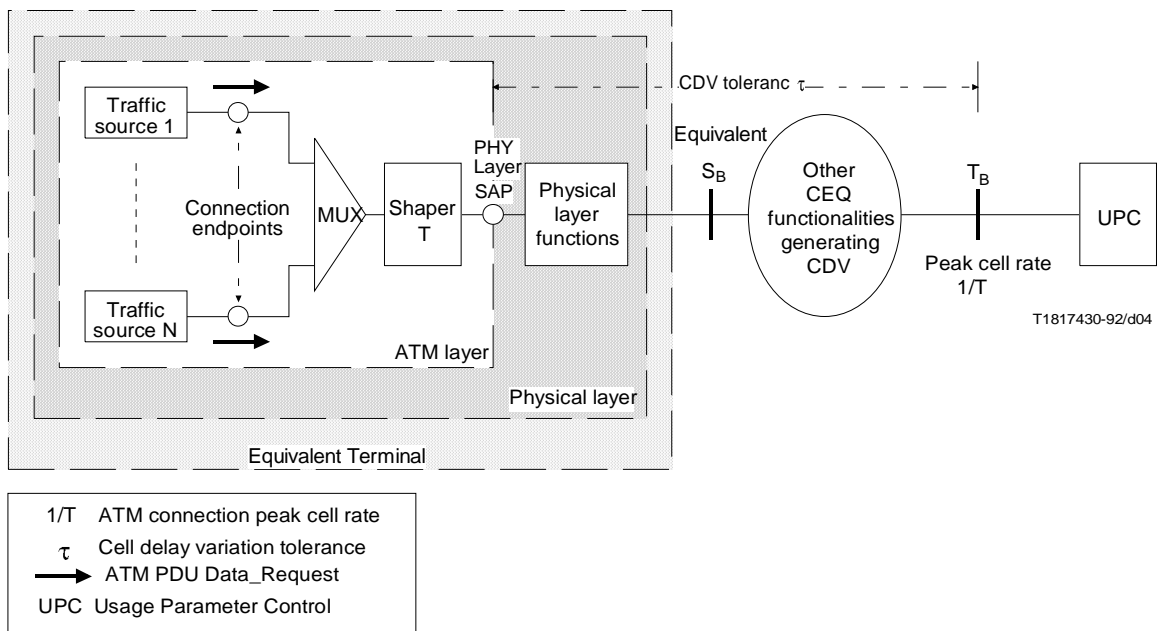


FIGURE 4/I.371

Reference configuration and equivalent terminal for the definition of the peak cell rate of an ATM connection

2.4.1.2 Peak cell rate granularity specification

Network functions such as UPC/NPC cannot be requested to handle every specific peak cell rate value but only a restricted, discrete and finite set of values. The ordered list of these values is referred to as the ATM peak cell rate granularity.

As the peak cell rate definition, the peak cell rate granularity specification should also be based on the peak emission interval.

The ATM peak cell rate granularity and its coding are for further study.

2.4.2 Other traffic parameters

For further study.

3 Functions and procedures for traffic control and congestion control

3.1 Introduction

Generic traffic control and congestion control functions are defined as the set of actions respectively taken by the network in all the relevant network elements to avoid congestion conditions or to minimize congestion effects and to avoid the congestion state spreading once congestion has occurred.

Under normal operation, i.e. when no network failures occur, functions referred to as traffic control functions in this Recommendation are intended to avoid network congestion.

However, congestion may occur, e.g. because of malfunctioning of traffic control functions caused by unpredictable statistical fluctuations of traffic flows or of network failures. Therefore, additionally, functions referred to as congestion control functions in this Recommendation are intended to react to network congestion in order to minimize its intensity, spread and duration.

3.1.1 Traffic control and congestion control functions

A range of traffic and congestion control functions will be used in the B-ISDN to maintain the QOS of ATM connections.

The following functions are described in this Recommendation.

- a) *Traffic control functions:*
 - i) Network resource management (3.2.1).
 - ii) Connection admission control (3.2.2).
 - iii) Usage/network parameter control (3.2.3).
 - iv) Priority control and selective cell discarding (3.2.4).
 - v) Traffic shaping (3.2.5).
 - vi) Fast resource management (3.2.6).
- b) *Congestion control functions:*
 - i) Selective cell discarding (3.3.1).
 - ii) Explicit forward congestion indication (3.3.2).

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

- Connection admission control that reacts to and takes account of the measured load on the network.
- Variation of usage monitored parameters by the network. For example, reduction of the peak rate available to the user.
- Other traffic control functions (e.g. re-routing, connection release, OAM functions) are for further study.

The impact on standardization of the use of these additional functions (e.g. the impact on ATM layer management, user-network signalling and control plane) requires further study.

Different levels of network performance may be provided on ATM connections by proper routing, traffic shaping, priority control and resource allocation to meet the required ATM layer QOS for these connections.

3.2 Traffic control functions

3.2.1 Network resource management

Use of virtual paths is described below. Other networking techniques are for further study.

3.2.1.1 Use of virtual paths

Virtual paths are an important component of traffic control and resource management in the B-ISDN. With relation to traffic control, VPCs can be used to:

- simplify CAC;
- implement a form of priority control by segregating traffic types requiring different QOS;
- efficiently distribute messages for the operation of traffic control schemes (for example to indicate congestion in the network by distributing a single message for all VCCs comprising a VPC);
- aggregate user-to-user services such that the UPC/NPC can be applied to the traffic aggregate;
- aggregate network capabilities such that the NPC can be applied to the traffic aggregate.

VPCs also play a key role in network resource management. By reserving capacity on VPCs, the processing required to establish individual VCCs is reduced. Individual VCCs can be established by making simple connection admission decisions at nodes where VPCs are terminated. Strategies for the reservation of capacity on VPCs will be determined by the trade-off between increased capacity costs and reduced control costs. These strategies are left to operators' decision.

The peer-to-peer network performance on a given VCC depends on the performances of the consecutive VPCs used by this VCC and on how it is handled in CRF(VC)s (see Figure 5).


If handled similarly by CRF(VC)s, different VCCs routed through the same sequence of VPCs experience similar expected network performance – e.g. in terms of cell loss ratio, cell transfer delay and cell delay variation – along this route.

Conversely, when VCCs within a VPC require a range of QOS, the VPC performance objective should be set suitably for the most demanding VCC carried. The impact on resource allocation is for further study.

Combining common routing and priority control by CLP may be used by call admission control for services requiring a number of VCCs with low differential delays and different cell loss ratios (e.g. multimedia services).

On the basis of the applications of VPCs contained in 2.3.2/I.311, namely:

- A) User-user application: the VPC extends between a pair of T_B reference points;
- B) User-network application: the VPC extends between a T_B reference point and a network node;
- C) Network-network application: the VPC extends between network nodes.



NOTES

- 1 VCCs 1 and 2 experience a network performance which depends on network performance on VPCs b and c and on how these VCCs are handled by CRF(VC)s. It may differ from network performance experienced by VCCs 3, 4 and 5, at least due to different network performances provided by VPCs.
- 2 VCCs 3, 4 and 5 experience similar network performances in terms of cell delay and cell delay variation if handled similarly by CRF(VC)s, whilst providing for two different cell loss ratios by using the CLP bit.
- 3 On a user-to-user VPC, the QOS experienced by individual VCCs depends on CEQ traffic handling capabilities.

FIGURE 5/I.371

Mapping cell loss ratios for virtual channel connections and virtual path connections

The above implies:

In case A: because the network has no knowledge of the QOS of the VCCs within the VPC, it is the user's responsibility to determine in accordance with the network capabilities the necessary QOS for the VPC.

In cases B and C: the network is aware of the QOS of the VCCs carried within the VPC and has to accommodate them.

Statistical multiplexing of VC links within a VPC where the aggregate peak of all VC links may exceed the virtual path connection capacity, is only possible when all virtual channel links within the virtual path connection can tolerate the QOS that results from this statistical multiplexing. The way this is managed is for further study.

As a consequence, when statistical multiplexing of virtual channel links is applied by the network operator, virtual path connections may be used in order to separate traffic thereby preventing statistical multiplexing with other types of traffic. This requirement for separation implies that more than one virtual path connection may be necessary between network origination/destination pairs to carry a full range of QOS between them. Implications of this are for further study.

3.2.1.2 Other networking techniques

For further study.

3.2.2 Connection admission control

3.2.2.1 General

Connection admission control is defined as the set of actions taken by the network at the call set up phase (or during call re-negotiation phase) in order to establish whether a virtual channel connection or a virtual path connection can be accepted or rejected.

On the basis of connection admission control in an ATM based network, a connection request for a given call is accepted only when sufficient resources are available to establish the call through the whole network at its required Quality of Service (QOS) and to maintain the agreed QOS of existing calls. This applies as well to re-negotiation of connection parameters within a given call.

In a B-ISDN environment a call can require more than one connection (e.g. for multimedia or multiparty services such as videotelephony or videoconferencing). In this case, connection admission control procedures should be performed for each virtual channel connection or virtual path connection.

Priority control using the CLP bit allows at most for two cell loss ratio objectives for ATM connections. Delay sensitivity is part of the required QOS.

In the case of an on-demand service, the connection establishment procedures will enable CAC to derive at least the following information :

- source traffic descriptors;
- required QOS class.

In the case of permanent or reserved service (e.g. using a permanent virtual path connection or a permanent virtual channel connection), this information is indicated with an appropriate OAM procedure, either on-line (e.g. signalling) or off-line (e.g. service order) basis.

Connection admission control makes use of this information to determine:

- whether the connection can be accepted or not;
- traffic parameters needed by usage parameter control;
- routing and allocation of network resources.

The role of priority control in connection admission control is for further study. Further information on priority control can be found in 3.2.4.

3.2.2.2 Parameters for connection admission control

3.2.2.2.1 Required QOS class

For a single ATM connection, a user indicates at most two QOS classes from the QOS classes which the network provides, only differing by the cell loss ratio. Specific QOS classes are the subject for further study.

3.2.2.2.2 Negotiation of traffic characteristics

The user will negotiate the traffic characteristics of the ATM connections with the network at connection establishment. These characteristics may be re-negotiated during the lifetime of the connection at the request of the user. The network may limit the frequency of these re-negotiations.

The re-negotiation procedure and the impact on network element complexity require further study.

3.2.2.3 Resource allocation

In order to ensure network performance and to protect the network, both $CLP = 0$ and $CLP = 1$ traffic flows have to be allocated resources.

Different strategies of network resource allocation may be applied for $CLP = 0$ and $CLP = 1$ traffic flows. In addition, information such as the measured network load may be used when performing CAC. This may allow a network operator to achieve higher network utilization while still meeting the performance objectives.

Resource allocation schemes are for further study. They may be left to network operators' decision.

3.2.3 Usage parameter control and network parameter control

Usage parameter control (UPC) and network parameter control (NPC) perform similar functionalities at different interfaces: the UPC function is performed at the user-network interface, whereas the NPC function is performed at the Inter-Network Interfaces.

The use of a UPC function is recommended, and the use of an NPC function is a network option. Whether or not the operator chooses to use the NPC function, the network-edge-to-network-edge and user-to-user performance objectives still need to be met.

3.2.3.1 UPC/NPC functions

Usage/network parameter control is defined as the set of actions taken by the network to monitor and control traffic in terms of traffic offered and validity of the ATM connection, at the user access and the network access, respectively. Their main purpose is to protect network resources from malicious as well as unintentional misbehaviour which can affect the QOS of other already established connections by detecting violations of negotiated parameters and taking appropriate actions.

Connection monitoring encompasses all connections crossing the UNI or inter-network interface usage parameter control and network parameter control apply to both user VCCs/VPCs and signalling virtual channels. Methods for monitoring meta-signalling channels and OAM flows are for further study.

The monitoring task for usage parameter control and network parameter control is performed for VCCs and VPCs, respectively by the following two actions :

- 1) checking the validity of VPI and VCI (i.e. whether or not VPI/VCI values are assigned), and monitoring the traffic entering the network from active VCCs in order to ensure that parameters agreed upon are not violated;
- 2) checking the validity of VPI (i.e. whether or not VPI values are assigned), and monitoring the traffic entering the network from active VPCs in order to ensure that parameters agreed upon are not violated.

3.2.3.2 UPC/NPC requirements

The need for and the definition of a standardized UPC/NPC algorithm requires further study. A number of desirable features of the UPC/NPC algorithm can be identified as follows:

- capability of detecting any illegal traffic situation;
- selectivity over the range of checked parameters (i.e. the algorithm could determine whether the user behaviour is within an acceptance region);
- rapid response time to parameter violations;
- simplicity of implementation.

There are two sets of requirements relating to the UPC/NPC:

- those which relate to the Quality of Service impairments the UPC/NPC might directly cause to the user cell flow;
- those which relate to the resource the operator should allocate to a given path/channel and the way the network intends to protect those resources against misbehaviour from the user side or another network (due to fault conditions or maliciousness).

There is a practical uncertainty in determining the values of the controlled parameters. Hence, in order to have adequate control performance, tolerances of controlling performance parameters need to be defined. The definition of these tolerances is for further study.

Two performance parameters have been identified. They have to be considered when assessing the performances of UPC/NPC mechanisms. Methods for evaluating UPC/NPC performance and the need to standardize these methods are for further study.

- Response time: the time to detect a given non-compliant situation on a VPC/VCC under given reference conditions.
- Transparency: for the same set of reference conditions, the accuracy with which the UPC/NPC initiates appropriate control actions on a non-compliant connection and avoids inappropriate control actions on a compliant connection.

Additional UPC/NPC performance parameters are for further study.

A specific UPC/NPC mechanism may commit errors by taking policing actions on a compliant connection, i.e. declaring a cell as non-compliant although the connection is actually compliant. It can also fail to take the appropriate policing actions on a non-compliant connection.

Unappropriate actions of the UPC/NPC on a compliant connection are part of the overall network performance degradation. Safety margins may be provisioned depending upon the UPC/NPC algorithm to limit the degradation introduced by the UPC/NPC.

Policing actions performed on the excess traffic in case of traffic contract violation are not to be included in the network performance degradation allocated to the UPC/NPC.

Impact of UPC/NPC on cell delay should also be considered. Cell delay and cell delay variation introduced by UPC/NPC is also part of the delay and delay variation allocated to the network.

3.2.3.2.1 Performance of peak cell rate UPC/NPC

A method to determine whether a traffic flow is conforming to a negotiated peak cell rate at a given interface is currently considered for network performance purposes. Non-conformance is measurable by a 1 point-measurement process in terms of the ratio γ_M between the number of cells exceeding the traffic contract and the total number of submitted cells.

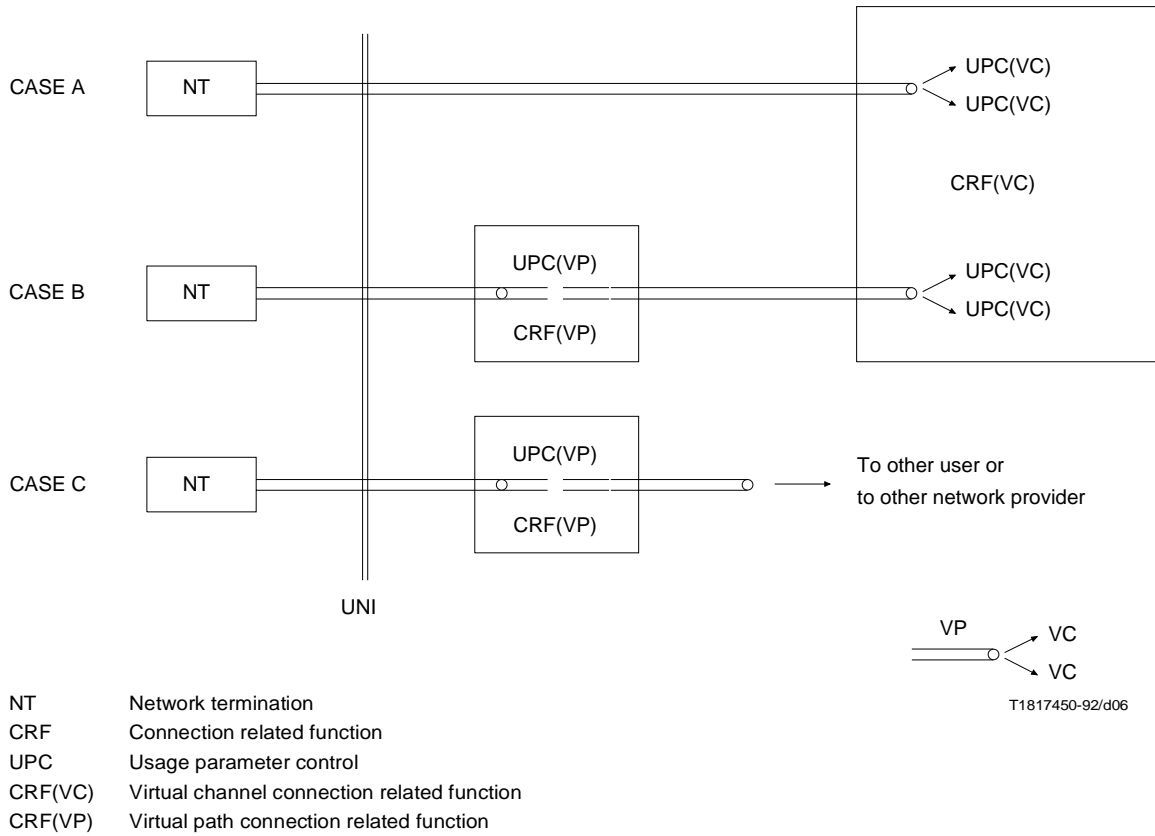
An ideal UPC/NPC implementing the 1 point-measurement process would just take policing actions on a number of cells according to this ratio. Although the process allows for a cell-based decision, it is not possible to predict which particular cells of a non-compliant connection will suffer from the policing action (this is because of measurement phasing).

According to the definition of the conformance of a traffic flow to a peak cell rate, the transparency of a UPC/NPC mechanism can be defined by the accuracy with which this mechanism approaches the ideal mechanism, i.e. the difference between the reference policing ratio γ_M and the actual policing ratio γ_p . A positive difference means that the UPC/NPC is taking less policing action than a measurement process would do. A negative difference means that policing actions are unduly taken by the UPC/NPC.

The exact way of measuring the transparency of a given mechanism for peak cell rate UPC/NPC and its dependence on time requires further study.

3.2.3.3 UPC location

Usage parameter control is performed on VCCs or VPCs at the point where the first VP or VC links are terminated within the network. Three possibilities can be identified as shown in Figure 6:



NOTES

- 1 In case A, the VPI value does not identify a negotiated VPC.
- 2 Provision of UPC at other locations is for further study.

FIGURE 6/I.371

Location of the usage parameter control functions

NOTE – In the following cases, CRF(VC) stands for virtual channel connection related function, and CRF(VP) stands for virtual path connection related function. A CRF(VC) or a CRF(VP) may respectively be a VC or VP concentrator.

Case A (Figure 6): User connected directly to CRF(VC):

Usage parameter control is performed within the CRF(VC) on VCCs before the switching function is executed (action 1, 3.2.3.1).

Case B (Figure 6): User connected to CRF(VC) via CRF(VP):

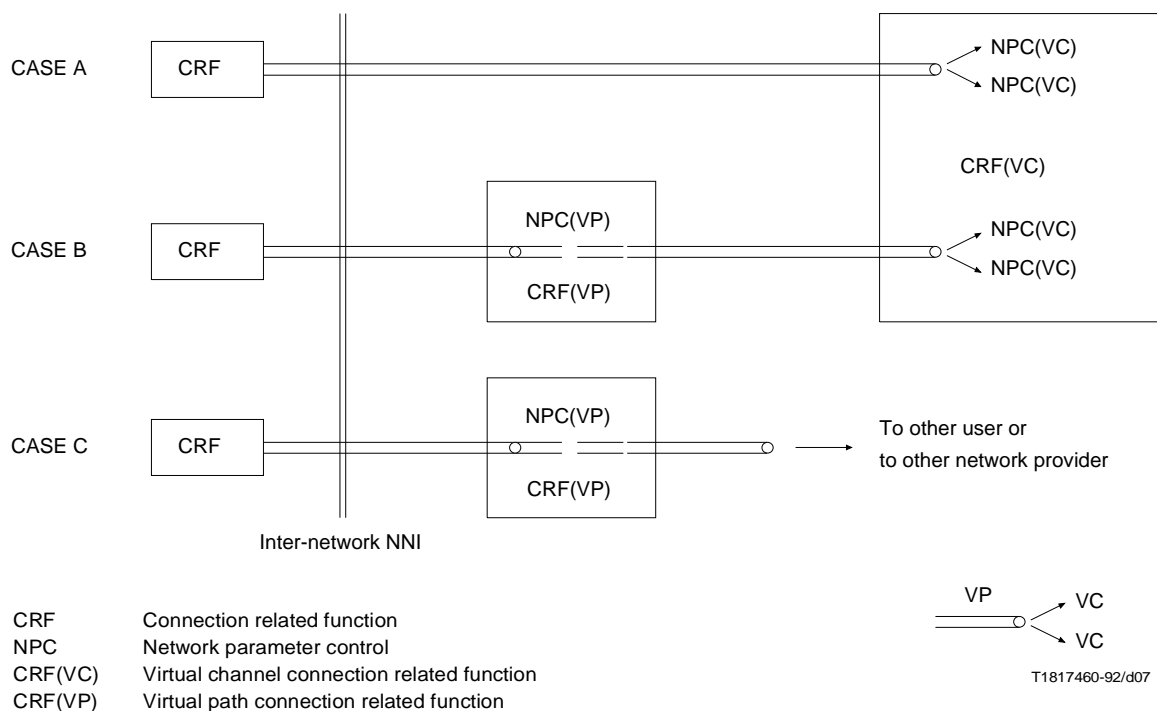
Usage parameter control is performed within the CRF(VP) on VPCs only (action 2, 3.2.3.1) and within the CRF(VC) on VCCs only (action 1, 3.2.3.1).

Case C (Figure 6): User connected to user or to another network provider via CRF(VP):

Usage parameter control is performed within the CRF(VP) on VPCs only (action 2, 3.2.3.1). VCC usage parameter control will be done by another network provider when CRF(VC) is present.

3.2.3.4 NPC location

Network parameter control is performed on VCCs or VPCs at the point where they are first terminated within the network. Three possibilities can be identified as shown in Figure 7. This requires further study.



NOTES – In case A and B the VPI value does not identify a negotiated VPC.

FIGURE 7/I.371

Location of the network parameter control functions

NOTE – In the following cases, CRF(VC) (resp. CRF(VP)) stands for virtual channel connection (resp. virtual path connection) related functions.

Case A (Figure 7): Originating network connected directly to CRF(VC) :

NPC is performed within the CRF(VC) before the switching function is executed (action 1, 3.2.3.1).

Case B (Figure 7): Originating network connected to the CRF(VC) via the CRF(VP):

NPC is performed within the CRF(VP) on VPCs (action 2, 3.2.3.1) only before the VP switching functions is executed and within the CRF(VC) on VCCs only (action 1, 3.2.3.1) before the switching function is executed.

Case C (Figure 7): Originating network connected to user or another network provider via CRF(VP):

NPC is performed within the CRF(VP) on VPCs only (action 2, 3.2.3.1). VCC network parameter control is performed by another network provider when CRF(VC) is present.

3.2.3.5 Traffic parameters subject to control at the UPC/NPC

Traffic parameters which may be subject to control are those included in the source traffic descriptor (see 2). Whether all these parameters or a subset are subject to control depends upon CAC and UPC/NPC mechanism. The peak cell rate has to be controlled for all types of connections.

3.2.3.6 UPC/NPC actions

The UPC/NPC is intended to control the traffic offered by an ATM connection to ensure compliance with the negotiated traffic contract. The objective is that a user will never be able to exceed the traffic contract.

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

- a) cell passing;
- b) cell re-scheduling (when traffic shaping and usage parameter control are combined, optional);
- c) cell tagging (network operator optional); cell tagging operates on CLP = 0 cells only, by overwriting the CLP bit to 1;
- d) cell discarding.

Cell passing and cell re-scheduling are performed on cells which are identified by a UPC/NPC as compliant. Cell tagging and cell discarding are performed on cells which are identified by a UPC/NPC as non-compliant.

The specific monitoring actions to be taken depend on the access network configuration.

Besides the above actions at the cell level, as an option, one other action performed at the connection level may be initiated by the UPC/NPC:

- releasing the connection.

3.2.3.7 Relationship between UPC/NPC cell loss priority and network performance

When an ATM connection utilizes the CLP capability on user request, network resources are allocated to CLP = 0 and CLP = 1 traffic flows as described in 3.2.2.3. By controlling CLP = 0 and CLP = 0 + 1 traffic flows (see Figure 8), allocating adequate resources and suitably routing, a network operator may provide the two requested QOS classes for CLP = 0 and CLP = 0 + 1 cell flows.

If the tagging option is used by a network operator, CLP = 0 cells identified by the UPC/NPC function performed on CLP = 0 flow as non-compliant are converted to CLP = 1 cells and merged with the user-submitted CLP = 1 traffic flow before the CLP = 0 + 1 traffic flow enters the UPC/NPC mechanism.

A cell identified as non-compliant by the UPC/NPC function performed on the aggregate CLP = 0 + 1 flow is discarded.

When no additional network resource has been allocated for CLP = 1 traffic flow (either on user request or due to network provisioning), CLP = 0 cells identified by the UPC/NPC as not compliant are discarded. In this case, tagging is not applicable.

Since cell sequence integrity is maintained on any ATM connection, the UPC/NPC including its optional tagging action must operate as a single server using First In First Out (FIFO) service discipline for each ATM connection.

Subclause 3.2.3.2 addresses undue UPC/NPC actions on compliant ATM connections. This is part of the network performance degradation allocated to the UPC/NPC and should remain of a very low probability.

When the CLP capability is used by an ATM connection and the CLP = 0 + 1 aggregate flow is not complying to the traffic contract, the UPC/NPC function performed on the aggregate flow may discard CLP = 0 cells which were not considered in excess by the UPC/NPC function performed on the CLP = 0 cell stream.

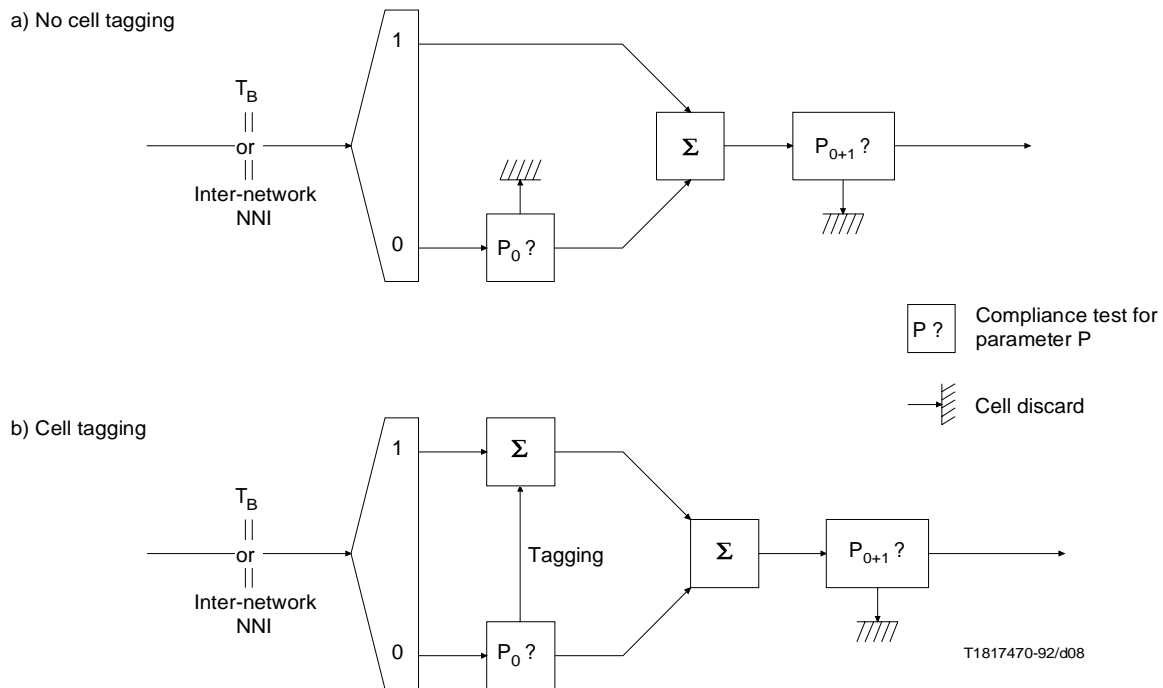


FIGURE 8/L.371
Possible actions of the UPC/NPC

3.2.3.8 Relationship between UPC/NPC, OAM and network management

OAM alarm indications may be provided by the UPC/NPC to the user and to the network management when enforcement actions occur on non-compliant VCCs/VPCs (e.g. cell discard). These alarm indications may initiate other enforcement actions (e.g. connection release). This is for further study.

As the output flow of any UPC should conform with the contract negotiated at connection set-up, alarms due to misbehaving users should not propagate through the network. This is for further study.

OAM information inserted at the ATM layer or above is part of the corresponding ATM connection. Therefore, it will be subject to enforcement at the UPC/NPC and needs properly allocated resources.

The use of OAM cells for traffic control and resource management purposes (e.g. to estimate delay and delay variation) is for further study.

3.2.4 Priority control and selective cell discard

Network elements may selectively discard cells of the lower priority flow while still meeting network performance objectives on both flows.

3.2.5 Traffic shaping

Traffic shaping is a mechanism that alters the traffic characteristics of a stream of cells on a VCC or a VPC to achieve a desired modification of those traffic characteristics. Traffic shaping must maintain cell sequence integrity on an ATM connection.

Examples of traffic shaping are peak cell rate reduction, burst length limiting, reduction of CDV by suitably spacing cells in time, queue service schemes.

It is an option that the traffic shaping be used in conjunction with suitable UPC functions, provided the additional delay remains within the acceptable QOS negotiated at call set-up.

The options available to the network operator/service provider are the following:

- Re-shape the traffic at the entrance of the network and allocate resources in order to respect both the CDV and the propagation delay allocated to the network.
- Dimension the network in order to accommodate the input CDV and provide for a shaper at the output.
- Dimension the network in order both to accommodate the input CDV and comply with the output CDV without any shaping function.

Traffic shaping may also be used within the customer equipment or the terminal to ensure that the traffic generated by the source or at the UNI is conforming to the traffic contract.

Traffic shaping is an option for network operators and users.

3.2.6 Fast resource management

Fast resource management functions operate on the time scale of the round-trip propagation delay of the ATM connection. Potential fast resource management functions are for further study.

One possible fast resource management function that has been identified is as follows:

- In response to a user request to send a burst, the network may allocate capacity (e.g. bandwidth, buffer space) for the duration of the burst. When a source requests an increase of its peak cell rate, it has to wait until resources have been reserved in all network elements along the ATM connection before the new peak cell rate can be used. UPC/NPC parameters would be adjusted accordingly.

3.3 Congestion control functions

For low priority traffic, some adaptive rate control facilities at the ATM layer or above may be used. Such cell-based reactive techniques are for further study.

The following congestion control functions have been identified. Other congestion control functions are for further study.

3.3.1 Selective cell discard

A congested network element may selectively discard cells explicitly identified as belonging to a non-compliant ATM connection and/or those cells with CLP = 1 cell loss priority. This is to primarily protect, as long as possible, high priority CLP = 0 flows.

3.3.2 Explicit forward congestion indication

The EFCI is a congestion notification mechanism which may be used to assist the network in avoidance of and recovery from a congested state. Since the use of this mechanism by the CEQ is optional, the network operator should not rely on this mechanism to control congestion.

A network element in a congested state may set an explicit forward congestion indication in the cell header so that this indication may be examined by the destination CEQ. For example, the end user's CEQ may use this indication to implement protocols that adaptively lower the cell rate of the connection during congestion. A network element that is not in a congested state will not modify the value of this indication.

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

The impact of explicit forward congestion indication on the traffic control and congestion control functions requires further study.

3.3.3 Reaction to UPC/NPC failures

Due to equipment faults (e.g. in usage parameter control devices and/or other network elements) the controlled traffic characteristics at the UPC/NPC could be different from the values agreed during the call set-up phase. To cope with these situations, specific procedures of the management plane should be designed (e.g. in order to isolate the faulty link). The impact of these malfunctioning situations on the usage parameter control needs further study.

Annex A

(to Recommendation I.371)

Peak cell rate cell monitor algorithms accounting for cell delay variation tolerance

(This annex forms an integral part of this Recommendation)

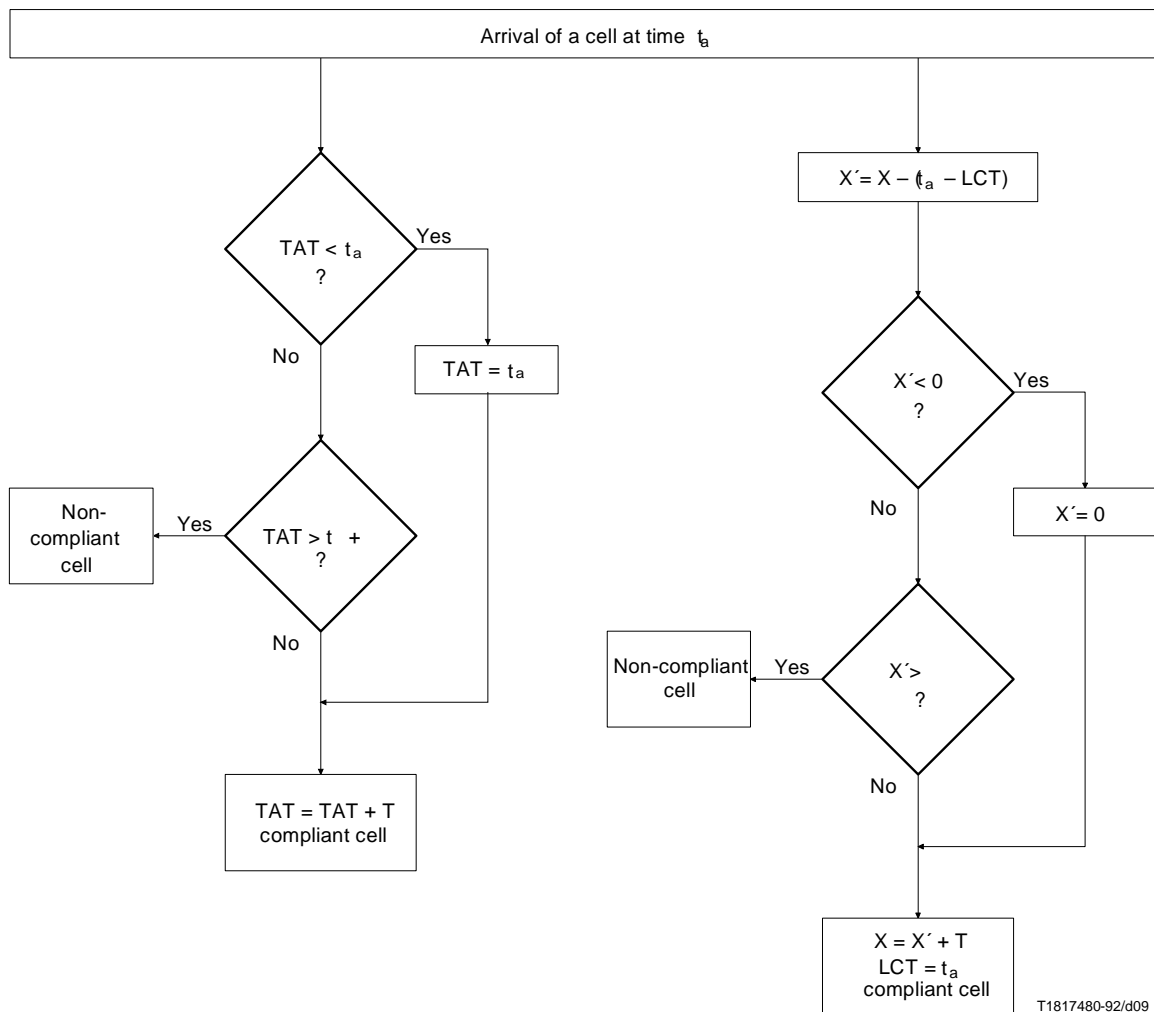
This annex provides two examples of algorithms that may be useful in monitoring the peak cell rate $1/T$ of an ATM connection while taking into account a certain CDV tolerance τ . No recommendation is made that either of these algorithms be used for actual implementation.

A virtual monitor algorithm (which might be located within the equivalent terminal) determines whether or not ATM_PDU requests are compliant with the negotiated values of the peak cell rate descriptor. It mirrors the measurement process of connection compliance currently considered for network performance purposes.

Let T be the peak emission interval and τ the cell delay variation tolerance. τ corresponds to the amount of "distortion" introduced by, for example, either customer equipment multiplexing before the T_B reference point or the mapping of ATM_PDU requests onto cell time slots. τ may be set to the difference between the maximum and minimum cell transfer delay throughout the CEQ.

T and τ are the only parameters needed to define the virtual monitor algorithm.

The virtual monitor algorithm is described in Figure/A.1. Two equivalent versions of this algorithm are shown: the virtual scheduling algorithm and the continuous state leaky bucket algorithm. In the first version, τ is expressed in units of time. In the second one, the leaky bucket capacity measured in units of time is equal to $L_B = T + \tau$.



Virtual scheduling algorithm

TAT Theoretical arrival time
 t_a Time of arrival of a cell

At the time of arrival t_a of the first cell of the connection, TAT = t_a

Continuous-state leaky bucket algorithm

X' Value of the leaky bucket counter
 X Auxiliary variable
 LCT Last compliance time

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

FIGURE A.1/I.371
Equivalent versions of the peak cell rate monitor algorithm

Annex B
(to Recommendation I.371)

List of Abbreviations

(This annex should form an integral part of this Recommendation)

AAL	ATM adaptation layer
ATM	Asynchronous transfer mode
CAC	Connection admission control
CBR	Constant bit rate
CDV	Cell delay variation
CEQ	Customer equipment
CRF(VC)	Virtual channel connection related functions
CRF(VP)	Virtual path connection related functions
CLP	Cell loss priority (bit)
CLR	Cell loss ratio
EFCI	Explicit forward congestion indication
FIFO	First In First Out
FRM	Fast resource management
GFC	Generic flow control
NPC	Network parameter control
OAM	Operation and maintenance
PDU	Protocol data unit
PTI	Payload type indicator
QOS	Quality of Service
SAP	Service access point
SDU	Service data unit
UNI	User network interface
UPC	Usage parameter control
VBR	Variable bit rate
VCC	Virtual channel connection
VCI	Virtual channel identifier
VPC	Virtual path connection
VPI	Virtual path identifier

Appendix I
(to Recommendation I.371)

**Examples of application of the equivalent terminal
for the peak cell rate definition**

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

An equivalent terminal has been used in 2.4.1 to define the peak cell rate of an ATM connection. The following two examples intend to clarify the concepts of peak emission interval T and the cell delay variation tolerance τ at T_B .


For sake of simplicity, the transmission rate at T_B is approximated by 150 Mbit/s. Δ is the cell cycle time at the interface at T_B .

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

Configuration 1

This configuration (Figure I.1) consists of a single terminal connected to T_B by a point-to-point single VCC.

ATM_PDU Data_requests are generated every $T = 1.25 \Delta$. This corresponds to a peak bit rate of 120 Mbit/s.



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

FIGURE I.1/I.371
Traffic configuration 1

Figure I.2 shows the basic events on a time scale and gives the needed CDV tolerance τ at T_B of configuration 1.

For sake of simplicity, the propagation delay between the terminal and T_B is assumed to be zero.

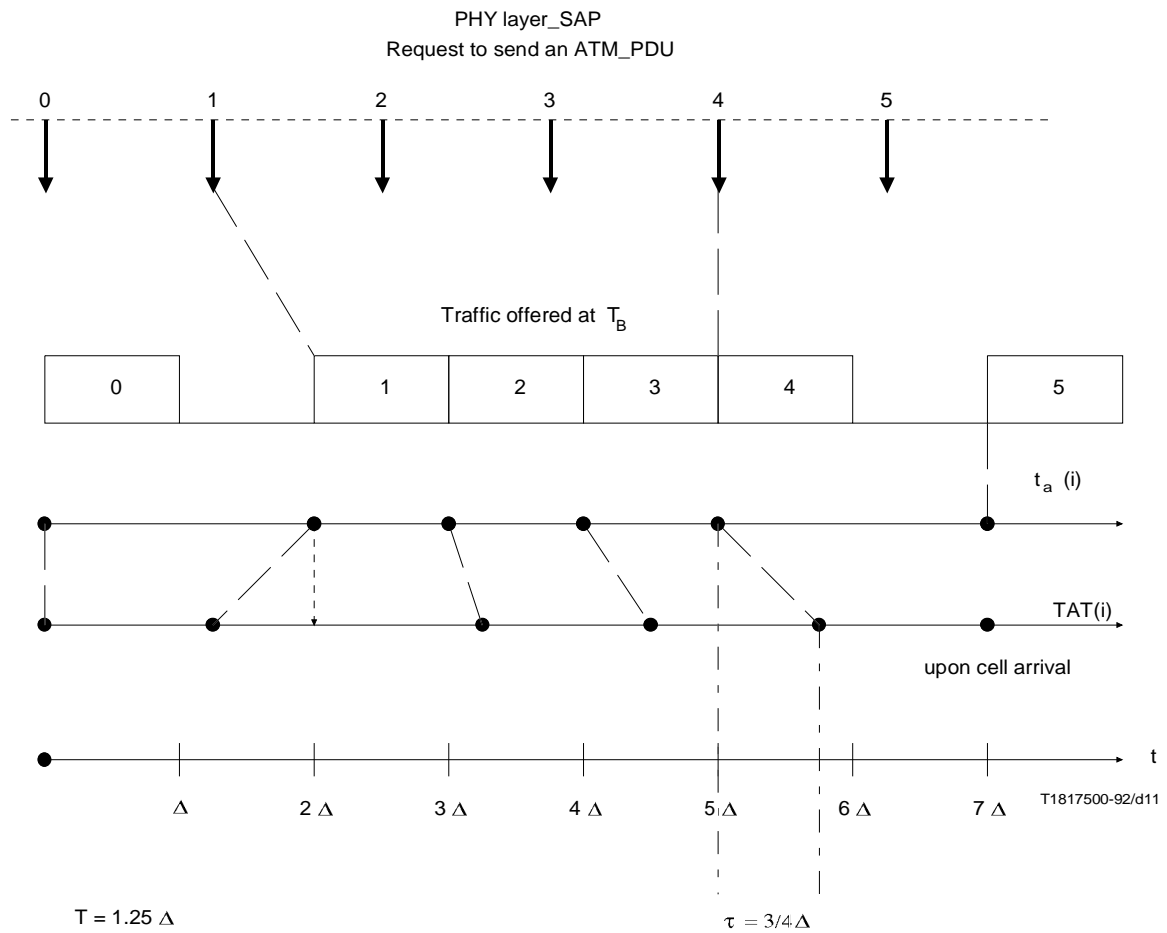


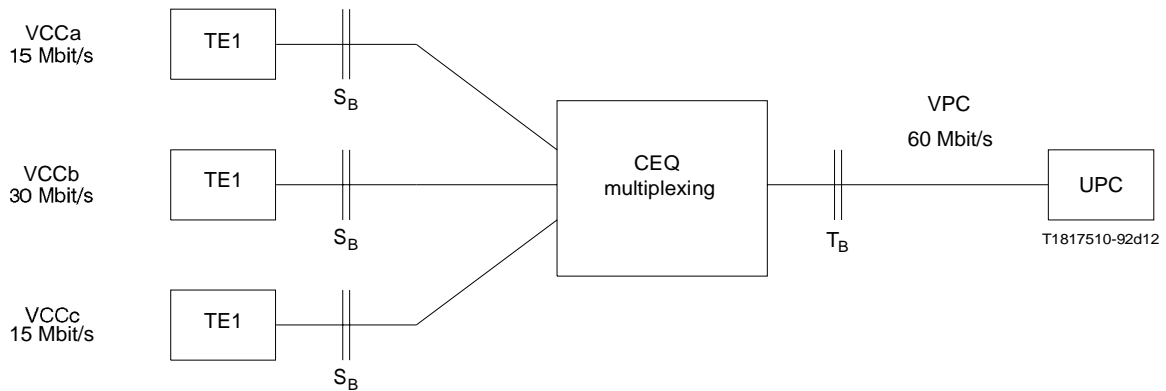
FIGURE I.2/I.371
 Illustration of CDV tolerance τ for traffic configuration 1

Configuration 2

This configuration (Figure I.3) consists of three terminals, each offering traffic on different VCCs. These three VCCs are multiplexed in the CEQ on to a VPC.

Terminals generate ATM_PDU Data_requests every 10Δ , 5Δ and 10Δ respectively, corresponding to 15 Mbit/s, 30 Mbit/s and 15 Mbit/s peak bit rates respectively.

The peak emission interval of the resulting VPC is $T = 2.5 \Delta$, which corresponds to a peak bit rate of 60 Mbit/s.



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

FIGURE I.3/I.371
Traffic configuration 2

Figure I.4 shows the basic events and the needed CDV tolerance τ at T_B corresponding to configuration 2.

This figure and the used terminology are similar to Figure I.2.

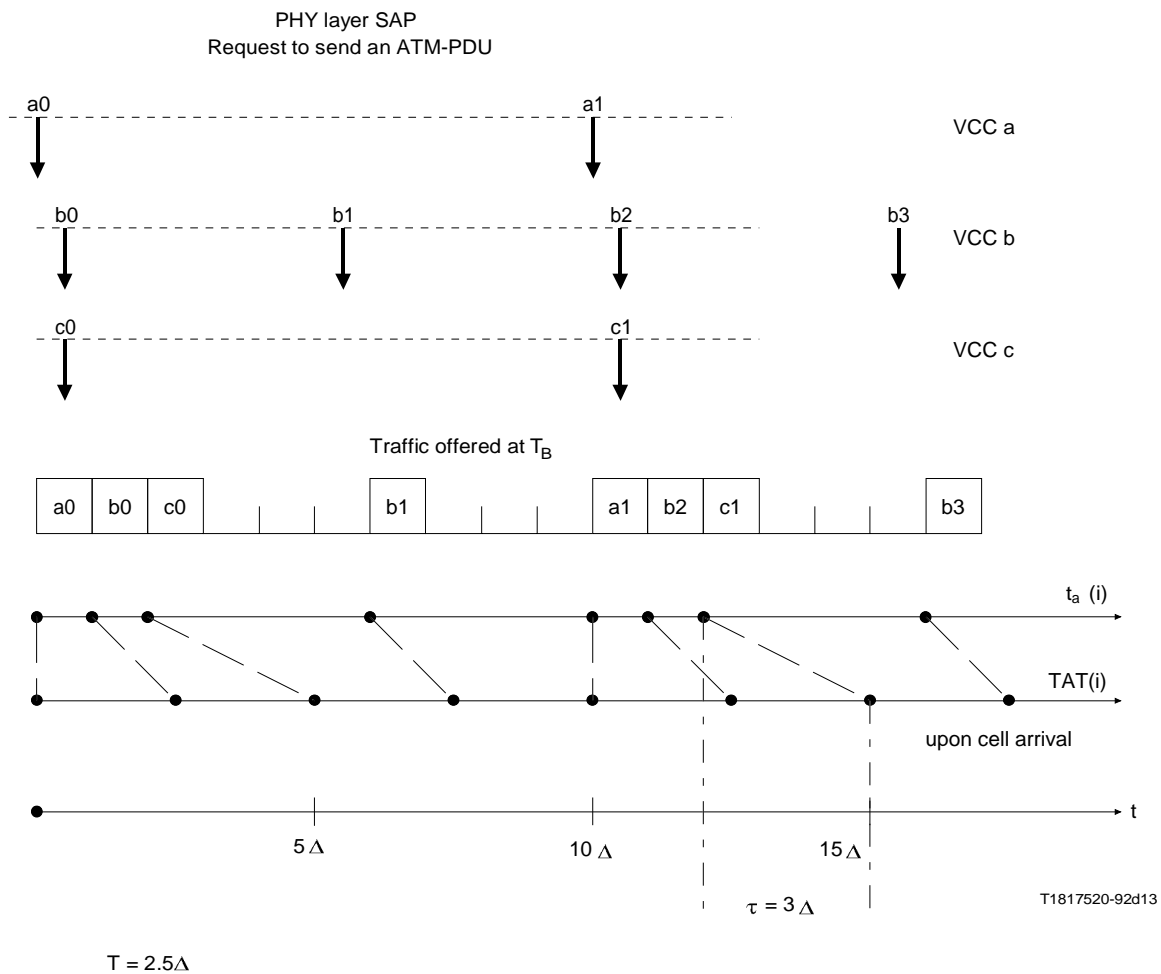


FIGURE I.4/I.371

Illustration of CDV tolerance τ for traffic configuration 2

