INTERNATIONAL TELECOMMUNICATION UNION

CCITT

I.370

THE INTERNATIONAL
TELEGRAPH AND TELEPHONE
CONSULTATIVE COMMITTEE

**INTEGRATED SERVICES
DIGITAL NETWORK (ISDN)**

**OVERALL NETWORK ASPECTS
AND FUNCTIONS,
ISDN USER-NETWORK INTERFACES**

CONGESTION MANAGEMENT FOR THE ISDN FRAME RELAYING BEARER SERVICE

Recommendation I.370

Geneva, 1991

FOREWORD

The CCITT (the International Telegraph and Telephone Consultative Committee) is a permanent organ of the International Telecommunication Union (ITU). CCITT 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 Plenary Assembly of CCITT which meets every four years, establishes the topics for study and approves Recommendations prepared by its Study Groups. The approval of Recommendations by the members of CCITT between Plenary Assemblies is covered by the procedure laid down in CCITT Resolution No. 2 (Melbourne, 1988).

Recommendation I.370 was prepared by Study Group XVIII and was approved under the Resolution No. 2 procedure on the 25th of October 1991.

CCITT NOTES

1) In this Recommendation, the expression “Administration” is used for conciseness to indicate both a telecommunication Administration and a recognized private operating agency.

2) A list of abbreviations used in this Recommendation can be found in Annex A.

Recommendation I.370

CONGESTION MANAGEMENT FOR THE ISDN FRAME RELAYING BEARER SERVICE

1 Congestion management principles

1.1 *Scope*

This Recommendation describes U-plane based congestion management strategy and mechanisms for ISDN frame relay bearer services. It covers both network and end user mechanisms and responsibilities to avoid or recover from periods of congestion. C-plane procedures at the user network interface other than clearing or not accepting calls are not recommended. Procedures, objectives and requirements for C-plane congestion management between networks are for further study. Special provisions for the treatment of continuous bit stream oriented (CBO) traffic is outside the scope of the Recommendation. This strategy is intended to operate for access channel rates up to 2048 kbit/s.

1.2 *Definitions*

access rate

The data rate of the user access channel (D, B or H). The speed of the access channel determines how much data (maximum rate) the end user can inject into the network.

committed burst size, Bc

The maximum committed amount of data a user may offer to the network during a time interval Tc. Bc is negotiated at call set-up.

excess burst size, Be

The maximum allowed amount of data by which a user can exceed Bc during a time interval Tc. This data (Be) is delivered in general with a lower probability than Bc. Be is negotiated at call set-up.

committed rate measurement interval, Tc

The time interval during which the user is allowed to send only the committed amount of data (Bc) and the excess amount of data (Be). Tc is computed.

committed information rate (CIR)

The information transfer rate which the network is committed to transfer under normal conditions. The rate is averaged over a minimum increment of time Tc. CIR is negotiated at call set-up.

forward explicit congestion notification (FECN)

See Recommendation Q.922 for the full definition.

backward explicit congestion notification (BECN)

See Recommendation Q.922 for the full definition.

consolidated link layer management message (CLLM)

See Recommendation Q.922 for the full definition.

discard eligibility indicator

This indicates that a frame should be discarded in preference to other frames in a congestion situation, when frames must be discarded to ensure safe network operation and maintain the committed level of service within the network.

fairness

An attempt by the network to maintain the committed call parameters which the end user negotiated at call setup time. An example of this would be first discarding the frames in excess of the committed information rate (CIR) and refusing to allow new call setups to occur prior to discarding committed data traffic.

offered load

Refers to the frames offered to the network, by an end user, to be delivered to the selected destination. The information rate offered to the network could exceed the negotiated class of service parameters.

congestion management

Includes network engineering, OAM procedures to detect the onset of congestion, and real time mechanisms to prevent or recover from congestion. Congestion management includes but is not limited to congestion control, congestion avoidance and congestion recovery as defined below.

congestion control

Refers to real-time mechanisms to prevent and recover from congestion during periods of coincidental peak traffic demands or network overload (e.g., resource failures). Congestion control includes both congestion avoidance and congestion recovery mechanisms.

congestion avoidance

Congestion avoidance procedures refer to procedures initiated at or prior to point A (see Figures 1/I.370 and 2/I.370) to prevent congestion from progressing to point B. Congestion avoidance procedures operate around point A and within the regions of mild congestion and severe congestion as shown in Figures 1/I.370 and 2/I.370.

congestion recovery

Congestion recovery refers to procedures initiated to prevent congestion from severely degrading the end user perceived quality of service(s) delivered by the network. These procedures are typically initiated when the network has begun to discard frames due to congestion. Congestion recovery procedures operate around point B and within the region of severe congestion as shown in Figures 1/I.370 and 2/I.370.

ingress node

The node that supports the source user-network interface (UNI).

egress node

The node that supports the destination user-network interface (UNI).

1.3 *Objectives of congestion management*

The primary objectives of congestion control mechanisms are to maintain, with a very high probability, specified quality of service (e.g. throughput, delay, frame loss) for each virtual call or permanent virtual circuit.

Congestion in the U-plane of a frame relay bearer service occurs when traffic, arriving at a resource (e.g. memory, bandwidth, processor), exceeds the network engineered level. It can also occur for other reasons (e.g. equipment failure). The impact of network congestion is performance degradation in terms of throughput and delay.

Two levels of congestion in terms of impact on class of service, are defined. Point A is the point beyond which the transit delay in the frame relay network increases at a rate faster than the rate at which offered load is increased. This is due to the network entry into a mild congestion state. This point is the final point on the curve that the network can guarantee the negotiated class of service. A further increase in offered load may cause a degradation in class of service. Point B is the point where the network begins discarding frames to control the existing level of congestion and prevent additional damage to the network provided services.

Point A and Point B are dynamic values determined by the instantaneous condition of the network resources. The end user may perceive the movement from point A to point B without increasing the offered load (e.g, due to resource failure or reconfiguration within the network). Threshold values are determined relative to the U-plane Quality of Service objectives to the end user. Specific networks may define different values, reflecting different performance objectives (e.g. for the support of different grades of services), even within the same network.

Congestion avoidance mechanisms aim to:

– minimize frame discard;

– maintain, with high probability and minimal variance, an agreed Quality of Service;

– minimize the possibility that one end user can monopolize network resources at the expense of other end users;

– be simple to implement, and place little overhead on either the end user or the network;

– create minimal additional network traffic;

– distribute network resources fairly among end users;

– limit spread of congestion to other networks and elements within the network;

– operate effectively regardless of the traffic flow in either direction between end users;

– have minimum interaction or impact on other systems in the frame relaying network; and

– minimize the variance in Quality of Service delivered to individual virtual circuits during congestion. (e.g. individual virtual circuits should not experience sudden degradation when congestion approaches or has occurred).

Congestion recovery mechanisms (in addition to the above) aim to ensure recovery of the network from a severely congested state.

1.4 *Requirements of congestion control mechanisms*

Congestion management mechanisms should have the following characteristics:

– Be part of the U-plane. Explicit congestion notification (ECN) shall be provided for in the U-plane. Note this applies to real-time notification aspects of congestion control and assumes that management functions such as gathering of statistics on congestion (i.e. when, where, why) could be accomplished outside the U-plane.

– Ensure transport of explicit congestion notification across frame relaying networks. The network(s) shall convey the backward ECN (BECN) towards the source end user and the forward ECN (FECN) towards the destination end user. This requires that these indications (if set) shall not be reset as they traverse the network(s) towards source and destination users.

– From a service perspective, call setup negotiations (e.g. throughput) are rate based. This means that from the standpoint of the service provided by the network in a frame relay environment, the rate at which information is offered to the network may be expressed as a number of information units per unit of time and is fundamental to all types of traffic to be carried.

– Reaction by the end user to the receipt of explicit congestion notification (FECN/BECN) is rate based and may be subject to standardization. It is noted that window mechanisms in terminals approximate rate based mechanisms and may be used to control the rate at which traffic is offered to a network.

– Networks should utilize, and users should react to, explicit congestion notification (i.e. not mandatory but highly desirable).

– Data sources which are unable to respond to explicit congestion notification (i.e. CLLM), can only be controlled by metering and discard.

– The network which perceives congestion should have the option to generate congestion notification using the appropriate congestion control protocols. When ECN is generated, it shall be sent in the appropriate direction(s). The policies for sending ECN will be different for the source control and destination control mechanisms.

– The end users (e.g., private networks), may generate ECNs.

1.5 *Congestion management strategy*

Distributed real-time congestion controls are necessary to prevent, and recover from, congestion during infrequent periods of coincidental peak traffic demands.

Congestion avoidance action is the joint responsibility of both the network and the end user and requires coordination between them. Avoidance procedures seek to return network operation to Region I of Figures 1/I.370 and 2/I.370.

Congestion avoidance with explicit signalling and congestion recovery with implicit signalling are considered to be effective and complementary forms of congestion control in the frame relaying bearer service.

Congestion avoidance and recovery schemes are distributed in that traffic monitoring (e.g. by buffer usage) is most efficient and accurate at congested resources, while traffic rate control is most effective when carried out by an end user. For end users to know when to decrease/increase their traffic rates, there must be a standardized notification mechanism between the network and end user.

Joint responsibility and procedures between the end user and the network should be verifiable by the network.

Congestion recovery initiation is the responsibility of the network. The end user should assist the network by continuing the avoidance procedures. Congestion recovery is used to help control offered load on severely congested networks and move from region III to region I of Figures 1/I.370 and 2/I.370.

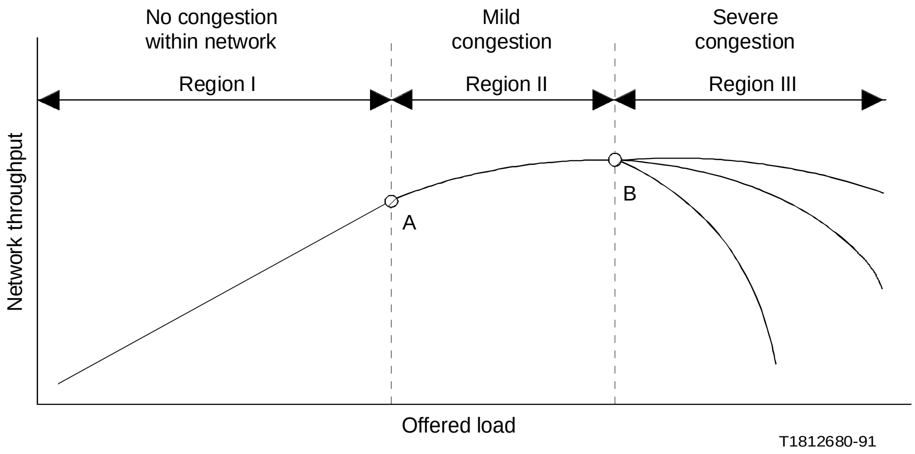

Note – The different lines in the severe congestion region reflect the fact that networks react and degrade differently in the face of severe congestion.

FIGURE 1/I.370

Throughput and network congestion

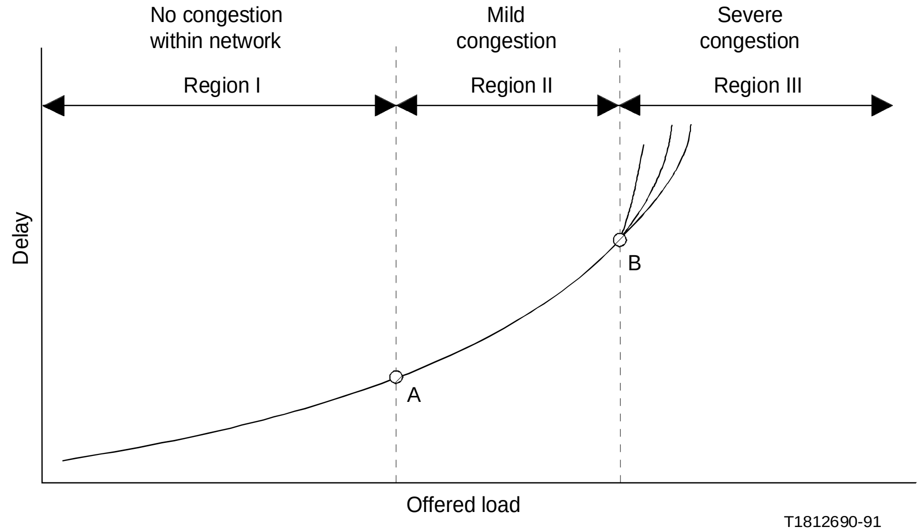

Note – The different lines in the severe congestion region reflect the fact that networks react and degrade differently in the face of severe congestion.

FIGURE 2/I.370

Delay and network congestion

1.5.1 *Congestion control mechanisms*

1.5.1.1 *Explicit congestion notification*

Commonly used end-to-end protocols operate with either source controlled or destination controlled transmit mechanisms for which two optional explicit congestion notification mechanisms for the frame relay bearer service are provided. These mechanisms, when implemented, are independent, not mutually exclusive, and may be used concurrently.

– *Mechanism 1:* For destination controlled transmitters, the FECN is set in the core aspects protocol.

– *Mechanism 2:* For source controlled transmitters, the BECN is set in the core aspects protocol in frames transported in the reverse direction (i.e. toward the transmitter). Alternatively, a consolidated link layer management message may be generated. This provides reverse notification for one or more data link connection identifiers (DLCI's) within a single frame. The CLLM is sent on the layer management DLCI in the U-plane in the backward direction (i.e. toward the source end user). The CLLM and the BECN may be used together or separately to notify the end user.

1.5.1.2 *Discard eligibility*

The use of the discard eligibility indicator by the users and the network is optional. This discard eligibility indicator may be set by the user and/or the network. The discard eligibility indicator determines whether or not this frame should be discarded by the network in preference to other frames. This decision would be necessary when the network is congested, and frames must be discarded to ensure safe network operation and to maintain the committed level of service within the network. Frames offered in excess of the committed size (Bc) may be marked discard eligible by the network.

The discard eligibility indicator is symmetrical, and is passed across both UNI and the NNI.

1.5.2 *Network response to congestion*

The network should in principle generate explicit congestion notification using the appropriate protocol to the source end user and/or the destination end-user around point A (see Figure 1/I.370). All networks must transport the FECN and BECN indications, either unmodified or, if in congested condition, with the appropriate indication set.

Notification in the backward direction can be accomplished using either (or both) of two optional mechanisms:

– A BECN indication is sent with reverse traffic. When there is reverse traffic at the time congestion is noted, then the BECN indication can be carried on an existing frame.

– Consolidated link layer management message. The generation and transport of consolidated link layer management messages by a network is optional. If a network receives this message and does not implement this option, then the consolidated link layer message should be discarded.

The network cannot rely solely on user behaviour to control network congestion (see § 1.5). Therefore the network is expected to protect itself from catastrophic congestion situations, and may do so by monitoring the throughput of each call and invoking the frame discard strategy under congestion conditions for those calls which exceed the lesser of CIR and the information rate currently available to be allocated by the network. Therefore, as congestion can occur even when the calls do not exceed their negotiated throughput (e.g. during network failures), the network should discard frames in a way that assures fairness among users. In certain congestion situations the network may refuse to accept new calls and/or clear existing calls.

1.5.3 *User response to congestion*

End users should in principle reduce their offered load upon receiving implicit or explicit indication of network congestion. Terminals shall have the capability to receive the explicit congestion notification generated by the network even if they are not able to act on the information. Reduction of information transfer rate by an end user may result in an increase in the effective throughput available to the end user during congestion. A user of the frame relaying service should implement some form of congestion-sensitive rate adjustment function that has the following characteristics:

– no blocking of data flow under normal conditions even when the offered load exceeds the CIR;

– reduction to a lower information transfer rate upon detection of network congestion;

– progressive return to the negotiated information transfer rate upon congestion abatement.

The end user terminal should base the detection of network congestion on implicit congestion detection schemes as well as on explicit congestion notification.

Implicit congestion detection schemes involve certain events available in the Q.922 elements of procedures (e.g., reception of a REJECT frame, detection of frame loss, timer expiration, etc.), or at a higher layer.

1.5.3.1 *Terminals employing destination controlled transmitters*

Reaction, to implicit congestion detection or explicit congestion notification (FECN) bit, when supported, should be as follows: consistent with commonly used destination controlled protocol suites (e.g. the OSI class 4 transport protocol operated over the OSI connectionless network service), rate adjustment is typically a function of higher layer protocols, and end-user reaction is based on the state of the FECN bits that are received over a period of time.

1.5.3.2 *Terminals employing source controlled transmitters*

Reaction, to implicit congestion detection or explicit congestion notification (BECN) bit or CLLM, when supported, should be as follows: rate adjustment is typically a function of the data link layer elements of procedure, and end-user reaction is expected to be immediate when a BECN bit or CLLM is received.

1.6 *Relationships between parameters*

Figure 3/I.370 illustrates the relationships between access rate, excess burst size (Be), committed burst size (Bc), committed information rate (CIR), discard eligibility indicator, and the measurement interval parameters. The CIR, Bc and Be parameters are negotiated at call establishment time for demand establishment of communication or by subscription for permanent establishment of communication. Access rate is established by subscription for permanent access connections or during demand access connection establishment. Each end-user and the network participate in the negotiation of these parameters to agreed values. These negotiated values are then used to determine the measurement interval parameter Tc, even when the discard eligibility indicator (if used) is set. These parameters are also used to determine the maximum allowable end-user input.

The measurement interval is determined as shown in Table 1/I.370. The network and the end-users may control the operation of the discard eligibility indicator and the rate enforcer functions by adjusting the CIR, Bc and Be parameters in relation to the access rate. If both the CIR and Bc parameters are not equal to zero, then Tc = (Bc/CIR). In addition, there are two special conditions:

1) When CIR = access rate, Bc = 0 and Be = 0, both access rates must be equal (i.e., ingress = egress);

2) When CIR = 0 (Bc must = 0) and Be > 0, then Tc = (Be/Access rate).

Figure 3/I.370 is a static illustration of the relationship among time, cumulative user data (bits) and rate. In this example, the user sends four frames during the measurement interval T0 to (T0 + Tc). The slope of the line marked CIR is Bc/Tc. Bits are received at the access rate (by the ingress node) of the access channel. Since the sum of the number of bits contained in frames 1 and 2 is not greater than Bc, the network does not mark these frames with the discard eligibility indicator. The sum of the number of bits in frames 1, 2 and 3 is greater than Bc but not greater than Bc + Be; therefore, frame 3 is marked discard eligible. Since the sum of the number of all bits received by the network in frames 1, 2, 3 and 4 exceeds Bc + Be, frame 4 is discarded at the ingress node. Figure 3/I.370 does not illustrate the case where the end-user sets the discard eligibility indicator. In this case the frames are considered within Be and not CIR.

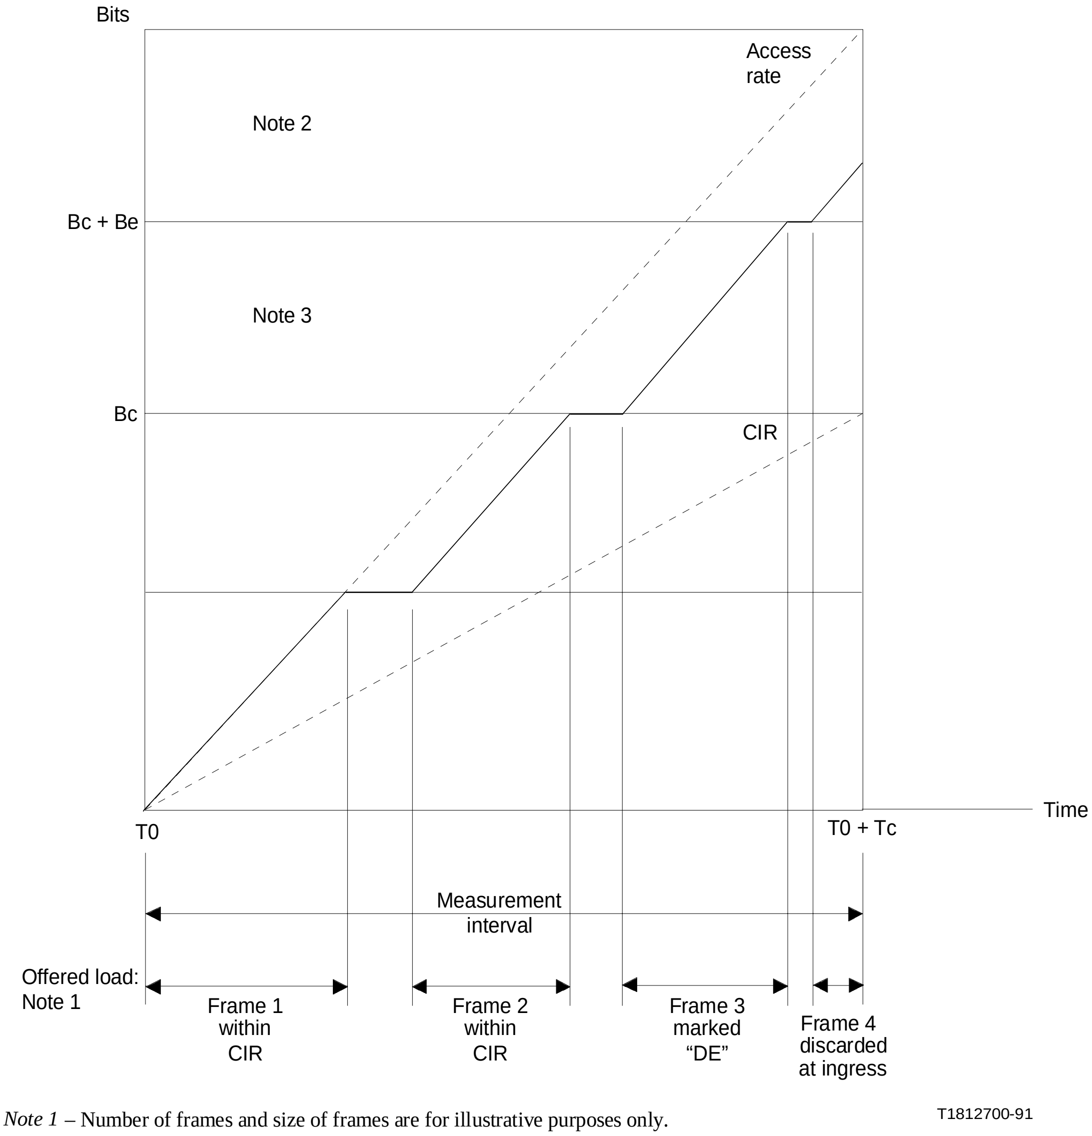

Note 1 – Number of frames and size of frames are for illustrative purposes only.

Note 2 – Frames discarded at ingress node. This is a region of rate enforcement.

Note 3 – Frames marked as discard eligible.

FIGURE 3/I.370

Illustration of the relationships between parameters

TABLE 1/I.370

Congestion parameter states

CIR	Bc	Be	Measurement interval (Tc)
> 0	> 0	> 0	Tc = (Bc/CIR)
> 0	> 0	= 0	Tc = (Bc/CIR)
= 0	= 0	> 0	Tc = (Be/Access rate)

Note – Table 1/I.370 contains the known valid parameter configurations. Other parameter configurations are for further study.

ANNEX A

(to Recommendation I.370)

Alphabetical list of abbreviations used in this Recommendation

BECN Backward explicit congestion notification

CBO Continuous bit stream oriented

CIR Committed information rate

CLLM Consolidated link layer management message

DLCI Data link connection identifier

ECN Explicit congestion notification

FECN Forward explicit congestion notification

UNI User-network interface