INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

E.681

(10/2001)

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

Traffic engineering – Traffic engineering for IP-networks

Traffic engineering methods for IP access networks based on hybrid fiber/coax system

ITU-T Recommendation E.681

ITU-T E-SERIES RECOMMENDATIONS

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

INTERNATIONAL OPERATION	
Definitions	E.100–E.103
General provisions concerning Administrations	E.104–E.119
General provisions concerning users	E.120–E.139
Operation of international telephone services	E.140–E.159
Numbering plan of the international telephone service	E.160–E.169
International routing plan	E.170–E.179
Tones in national signalling systems	E.180–E.189
Numbering plan of the international telephone service	E.190–E.199
Maritime mobile service and public land mobile service	E.200–E.229
OPERATIONAL PROVISIONS RELATING TO CHARGING AND ACCOUNTING IN THE INTERNATIONAL TELEPHONE SERVICE	
Charging in the international telephone service	E.230–E.249
Measuring and recording call durations for accounting purposes	E.260–E.269
UTILIZATION OF THE INTERNATIONAL TELEPHONE NETWORK FOR NON-TELEPHONY APPLICATIONS	
General	E.300–E.319
Phototelegraphy	E.320–E.329
ISDN PROVISIONS CONCERNING USERS	E.330–E.349
INTERNATIONAL ROUTING PLAN	E.350–E.399
NETWORK MANAGEMENT	
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	
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.649
Traffic engineering for IP-networks	**E.650–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	
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 the list of ITU-T Recommendations.

ITU-T Recommendation E.681

Traffic engineering methods for IP access networks based on hybrid fiber/coax system

Summary

This Recommendation describes generic preferred methods for the traffic control and dimensioning of IP-access networks based on hybrid fiber/coax (HFC) system, taking into account the capabilities and limitations of the cable-modem based access technology. The main focus is on the provision over HFC system of IP telephony service in an integrated voice/data environment. Factors affecting system capacity for the support of voice connections are described. Topics discussed include the dimensioning of a single upstream channel, upstream channel pooling, and interoperability between DOCSIS 1.0 and DOCSIS 1.1.

Source

ITU-T Recommendation E.681 was prepared by ITU-T Study Group 2 (2001-2004) and approved under the WTSA Resolution 1 procedure on 29 October 2001.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. 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 Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 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

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. 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, ITU had 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.

CONTENTS

		Page
1	Scope	1
2	References	1
	2.1 Normative references	1
	2.2 Non-normative references	2
3	Definitions	2
4	Abbreviations	2
5	Introduction	2
6	Grade of Service parameters	4
7	Traffic engineering considerations	4
	7.1 Factors affecting upstream capacity	4
	7.1.1 Channel characteristics	5
	7.1.2 Protocol features	5
	7.1.3 Propagation delay	5
	7.2 Protocol for upstream transmission	5
8	Dimensioning of a single upstream	6
	8.1 Voice capacity	6
	8.2 Available capacity for data	6
9	DOCSIS 1.0 and DOCSIS 1.1 interoperability	7
10	Upstream channel pooling	8
	10.1 Immediate upstream channel change	8
	10.2 Voice slot jitter	8
	10.3 Assignment algorithms	8
11	History	10
Appendix I – Example calculation of upstream channel capacity		10
Appendix II – Example of jitter window		11

ITU-T Recommendation E.681

Traffic engineering methods for IP access networks based on hybrid fiber/coax system

1 Scope

This Recommendation describes generic preferred methods for the traffic control and dimensioning of IP-access networks based on hybrid fiber/coax (HFC) system, taking into account the capabilities and limitations of the cable-modem based access technology. It contains relevant traffic principles for the planning, operation, and management of HFC-based IP-access networks so that quality of service (QOS) objectives to customers can be met.

For traffic engineering, this Recommendation assumes that the network is available: that is, it does not consider network equipment in a failure state.

The first release of this Recommendation is only concerned with the provision over HFC system of IP telephony service in an integrated voice/data environment. While the impact of voice service on the capacity available for data is considered, traffic-engineering methods for TCP/IP-supported data services such as web-browsing, email, file transfer, and high-speed data access are for further study. Other services such as video telephony, video-on-demand, are also for further study.

2 References

2.1 Normative 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; 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.526 (1993), *Dimensioning a circuit group with multi-slot bearer services and no overflow inputs*.

– ITU-T Recommendation E.651 (2000), *Reference connections for traffic engineering of IP access networks*.

– ITU-T Recommendation E.721 (1999), *Network grade of service parameters and target values for circuit-switched services in the evolving ISDN*.

– ITU-T Recommendation E.726 (2000), *Network grade of service parameters and target values for B-ISDN*.

– ITU-T Recommendation G.114 (2000), *One-way transmission time*.

– ITU-T Recommendation G.711 (1988), *Pulse code modulation (PCM) of voice frequencies*.

– ITU-T Recommendation G.728 (1992), *Coding of speech at 16 kbit/s using low-delay code excited linear prediction*.

– ITU-T Recommendation J.112 (1998), *Transmission systems for interactive cable television services*.

2.2 Non-normative references

The following references are listed here for information:

DOCSIS 1 Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification 1.0, SP-RFI-I05-991105, Cable Television Laboratories, Inc., November 1999.

DOCSIS 2 Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification 1.1, SP-RFIv1.1-I07-010829, Cable Television Laboratories, Inc., August 2001.

DOCSIS 3 Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premises Equipment Interface Specification, SP-CMCI-I05-001215, Cable Television Laboratories, Inc., December 2000.

3 Definitions

This Recommendation defines the following terms:

3.1 A **cable modem (CM)** is a modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system.

3.2 A **cable modem termination system (CMTS)**, located at the cable television system headend or distribution hub, provides complementary functionality to the cable modems to enable data connectivity to a wide-area network.

3.3 A **fiber node** is a point of interface between a fiber trunk and the coaxial distribution.

3.4 A **hybrid fiber/coax (HFC)** system is a broadband bi-directional shared-media transmission system using fiber trunks between the headend and the fiber nodes, and coaxial distribution from the fiber nodes to the customer locations.

4 Abbreviations

This Recommendation uses the following abbreviations:

CM Cable Modem

CMTS Cable Modem Termination System

CPE Customer Premises Equipment

GoS Grade of Service

HFC Hybrid Fiber/Coax system

IP Internet Protocol

MTA Multi-media Terminal Adapter

QoS Quality of Service

TCP Transmission Control Protocol

5 Introduction

This Recommendation uses the reference architecture of IP access networks based on hybrid fiber/coax systems as specified in 7.1/E.651. For convenience, Figure 7-1/E.651 is reproduced below as Figure 5-1.

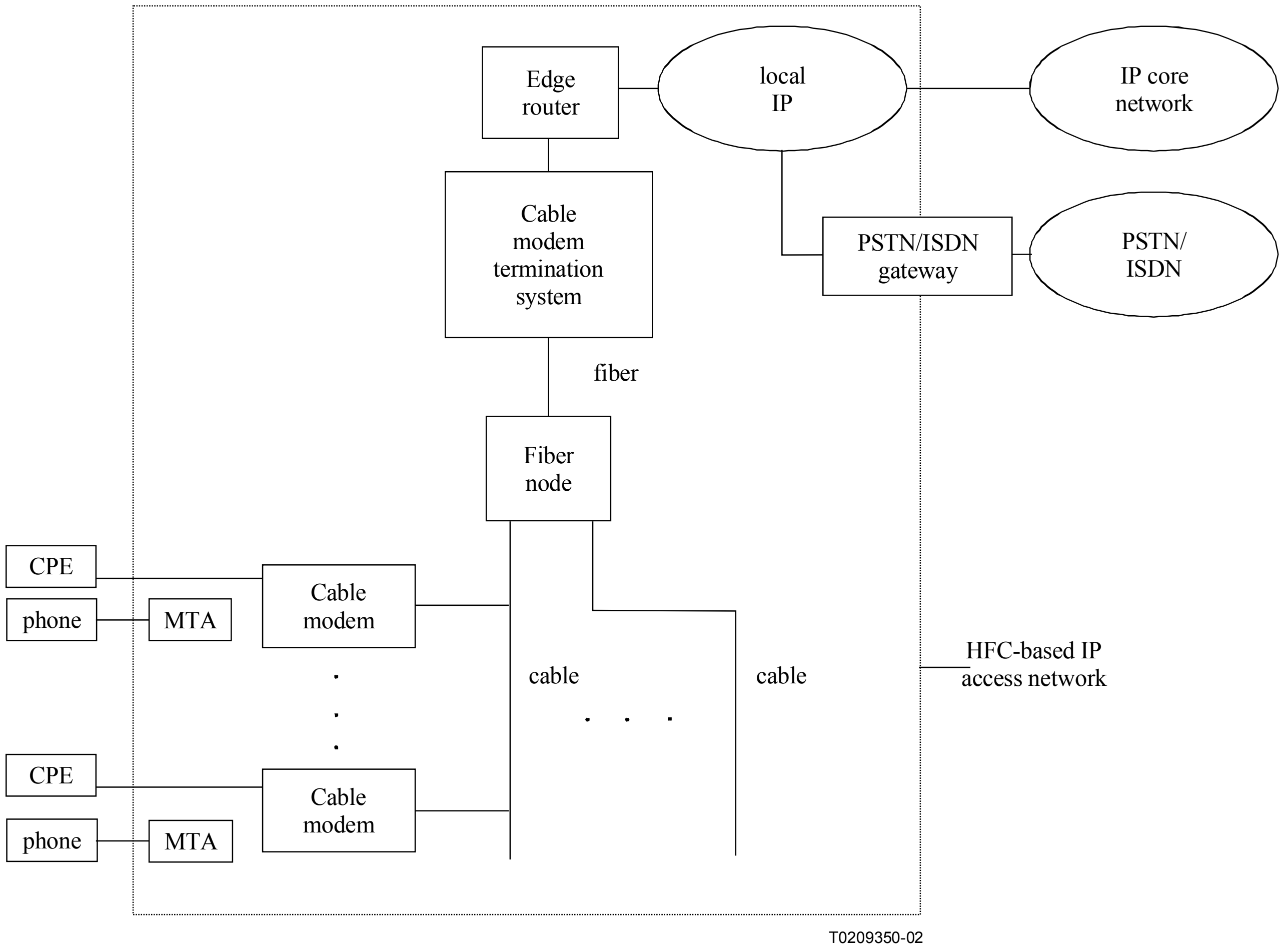

Figure 5-1/E.681 – Reference architecture of an HFC-based IP access network

This figure shows an HFC-based IP-access network with cable modems (CMs) located at customer premises connected to a cable modem termination system (CMTS) at the head-end. The HFC distribution plant includes fiber links between a CMTS and a fiber node, with the latter performing optical/electrical conversion. Coaxial cable is used to connect multiple customer premises in close proximity to the fiber node. Customer-premises equipment, such as a personal computer, can interface directly with a CM. An ordinary telephone is usually interfaced through a multi-media terminal adapter (MTA) to a CM. The MTA may be integrated with the CM. Through a managed local IP network and an edge router, an HFC-access network is connected to an IP core network for either Internet access or all-IP end-to-end telephone calls. Connection is also provided via a gateway for interworking with PSTN/ISDN for telephone calls. The edge router may be integrated with the CMTS and routes the traffic over the local IP network. It may also perform traffic policing and, optionally, admission control functions. The local IP network contains all the network elements/servers required for connection management and call processing.

A CMTS and a CM transfer IP traffic by using the *Data-Over-Cable Service Interface Specifications*, Radio Frequency Interface Specification, Version 1.1 [DOCSIS 2]. This specification is commonly referred to as DOCSIS 1.1. An earlier version, DOCSIS 1.0 [DOCSIS 1], was originally developed to support best-effort high-speed data service. Traffic engineering issues related to the interoperability between these two versions is addressed in clause 9.

6 Grade of Service parameters

In this clause, GoS parameters that are relevant for traffic engineering purposes are listed. Their definitions and target values are the subject of a future Recommendation.

For IP telephony, the following traffic GoS parameters at the call level are recommended:

1) probability of call blocking;
2) post-selection delay;
3) answer-signal delay;
4) call-release delay.

NOTE 1 – These parameters are functionally similar to the corresponding ones specified in ITU-T Recs E.721 and E.726.

NOTE 2 – In the context here, answer-signal delay is functionally equivalent to post pick-up delay. That is, the interval from the time when the terminating phone is picked up (after ringing) to the time when the end-to-end voice path to the originating phone is established. This interval is specified to avoid the "clipped hello".

After call establishment, for information transfer, the term *packet* refers to a packet containing one or more voice samples. In this phase, the following traffic GoS parameters are recommended:

1) speech transmission delay (including packetization delay and one-way packet transfer delay);
2) packet delay variation, also referred to as jitter;
3) packet loss (average and burst loss).

NOTE 3 – These parameters are functionally similar to the corresponding cell-level GoS parameters specified in ITU-T Rec. E.726.

7 Traffic engineering considerations

In an HFC system, two-way interactivity is accomplished by augmenting a cable-TV broadcast network with a return path in the upstream direction (i.e. a path from the customer to the network). However, due to the relatively small part of the spectrum allocated for upstream use (for compatibility with TV broadcasts) and the noise characteristics therein, the bandwidths available in the two directions are asymmetrical with the upstream bandwidth being much more limited than the downstream. Because of this asymmetry, a CMTS usually associates one downstream channel with multiple (typically up to eight) upstream channels. However, a CM can only access one of these upstream channels at a time for all its communication needs, i.e. to support several (usually up to four) simultaneous voice calls and/or one "always-on" high-speed data connection.

For high-speed data communications, throughput-intensive downloads and file transfers typically occur in the downstream direction, and the upstream is mainly used for the transmission of small packets containing acknowledgements and commands from the users. Unlike data applications, the bidirectional bandwidth requirements of interactive voice communication are inherently symmetrical. Hence, in supporting integrated voice/data services, the bandwidth-limited upstream channels may be the bottleneck.

7.1 Factors affecting upstream capacity

The capacity of an upstream channel depends on the physical characteristics and the protocol used for communication. The maximum round-trip propagation delay between a CM and a CMTS also affects the maximum capacity for supporting telephony service. Appendix I contains an example that shows how the following factors affect the upstream capacity for voice connections.

7.1.1 Channel characteristics

DOCSIS 1.1 permits operators the flexibility in selecting upstream physical-layer characteristics by providing two modulation formats (QPSK or 16-QAM), a set of five symbol rates, together with a corresponding set of channel widths. The raw bit rate of an upstream channel depends primarily on the channel bandwidth, modulation format, and symbol rate selected for the channel.

7.1.2 Protocol features

The number of voice connections that can be supported on a channel depends on the type of codec encoding algorithms, voice packetization interval, the use of payload header suppression, the size of the forward error correction codeword, and the size of a mini-slot.

DOCSIS specifies a mini-slot as a time unit for the purpose of upstream bandwidth allocation. The upstream channel is therefore modelled as a stream of contiguous mini-slots, with each mini-slot representing the time needed for transmission of a fixed number of bytes. Commonly used mini-slot sizes are either 8 or 16 bytes each.

7.1.3 Propagation delay

DOCSIS specifies that, once every 10 s, upstream bandwidth be dedicated to allow new CMs to join the network and to perform initial ranging. Ranging is the process for a CM to acquire the correct timing offset such that its transmissions are aligned to the correct mini-slot boundary. Referred to as *initial maintenance*, the bandwidth for this ranging activity requires "a long interval, equivalent to the maximum round-trip propagation delay plus the transmission time of the Ranging Request (RNG-REQ) message". During this interval, voice packets cannot be transmitted. To avoid the impact of this on the packet delay variation, upstream channel bandwidth should not be fully utilized for voice. Hence, the longer this interval, the smaller the upstream capacity for voice.

DOCSIS specifies a maximum spacing of 100 miles (160 km) between a CMTS and the most distant CM, although 15-25 km may be typical. Depending on the ratio of media (fiber to copper) length, media propagation delay varies. Suppose that the media speed equals 2/3 of the speed of light in vacuum (see ITU-T Rec. G.114), then the maximum round-trip propagation delay is 1.61 ms.

In practice, cable plants usually operate with shorter distances than the DOCSIS-specified maximum. To get a higher capacity for voice connections, it may be desirable for such cable plants to use a value smaller than the above maximum propagation delay to set the initial maintenance interval.

7.2 Protocol for upstream transmission

Access to the mini-slots for upstream transmission by a CM is controlled by the CMTS. A CM can send a packet only after it has requested and obtained from the CMTS a *grant*, i.e. permission to use some number of contiguous mini-slots. Responding to these requests, the CMTS schedules the upstream payload transmissions through successive grants to different CMs. In data communications, each time a CM is ready to send a packet, it requests a grant from the CMTS.

For constant bit rate (CBR) voice communications under DOCSIS 1.1, to minimize overhead, a CMTS automatically provides *unsolicited grants* to a CM for periodic upstream transmission of voice packets for each CBR connection that has been established for the CM. For simplicity, the term *voice slot* refers to a steady stream of periodic unsolicited grants to use succeeding contiguous sequence of mini-slots to transmit the fixed-sized voice packets generated from a CBR connection. Thus, a voice slot is allocated to a CM for each CBR connection it has established.

8 Dimensioning of a single upstream

The bandwidth of an upstream channel is shared among the payload traffic for both IP telephony and high-speed data, the overhead traffic from telephony signalling and DOCSIS maintenance activities, as well as traffic arising from contention for upstream transmission requests.

In this clause, it is assumed that each CM is configured *statically* to use a particular upstream channel at provisioning time (e.g. at installation or power up). Hence, each upstream channel is engineered separately.

8.1 Voice capacity

For a conservative estimate of the capacity of an individual upstream to support voice connections, the Erlang-B formula can be used:

$$\Pi = \left(a^n / n!\right) / \sum_{k=0}^{n} \left(a^k / k!\right)$$

where Π is call congestion (or probability of call blocking), a is the offered load, and n is the maximum number of voice connections that can be supported by an upstream channel. This number n can be determined by using a method similar to that in Appendix I.

To account for finite-source effect, the Engset formula for call congestion can be used:

$$\Pi = \left[\binom{m-1}{n}\hat{a}^n\right] \Bigg/ \left[\sum_{k=0}^{n}\binom{m-1}{k}\hat{a}^k\right]$$

where m is the number of traffic sources, and $\hat{a}$ is the average offered load per idle source. Let $\alpha = a/m$ be the average offered load per source. The relationship $\hat{a} = \alpha / [1 - \alpha\,(1 - \Pi)] \approx \alpha / (1 - \alpha)$ can be used to estimate the quantity $\hat{a}$, when $\Pi << 1$.

When the system supports a mixture of encoding schemes (e.g. both ITU-T Recs G.711 and G.728), different voice calls may require different data rates. In this case, the methods of ITU-T Rec. E.526 can be used for dimensioning.

8.2 Available capacity for data

The upstream capacity for data traffic depends largely on the degree of bandwidth sharing between voice and data services. In the hard partition scheme where voice and data use separate bandwidth regions with no sharing, data traffic gets only what is allocated in its bandwidth region. Any sharing schemes will allow higher effective average data bandwidth compared to the hard partition scheme [CS01, L01].

To facilitate voice/data sharing, it is assumed for traffic engineering purposes that the sequence of mini-slots in an upstream channel is partitioned into *frames*, the length of each being the packetization time for voice sampling. For most packet telephony applications over cable-access systems, this time is typically 10 ms to minimize speech transmission delays. Within each frame, the set of mini-slots is further partitioned into two *fixed* regions. One is designated for data service only, and accommodates variable-sized data packets. This is referred to as a *data-only region*. The other one, referred to as a *voice region*, is for both voice and data services, with voice having priority over data. Voice packets are fixed in size.

To estimate the *maximum* amount of idle capacity in the voice region that is available for use by data service, note that the probability for j busy voice slots, P_j is:

$$P_j = \left[\binom{m}{j}\hat{a}^j\right] \Bigg/ \left[\sum_{k=0}^{n}\binom{m}{j}\hat{a}^k\right] \qquad (j = 0, 1, \ldots, n)$$

The probability for at least i idle voice slots in the voice region is $\sum_{j=0}^{n-i} P_j$. This idle capacity is available for use by data, in addition to the bandwidth in the data-only region.

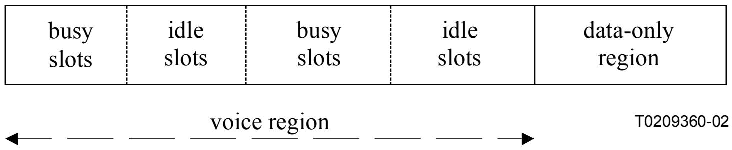

Figure 8-1/E.681 – A voice region with two sections of contiguous idle slots (i.e. two "holes")

Over time, as voice connections are established and cleared, "holes," or sections of contiguous idle voice slots, are created in the channel bandwidth as shown in Figure 8-1. To make full use of this idle capacity in the voice region, DOCSIS 1.1 specifies a data packet fragmentation procedure whereby data packets are fragmented to fill in these holes. A penalty for using this procedure is the extra 16-byte per-fragment overhead that is required to convey fragmentation and reassembly information. In comparison, note that the size of a voice slot under G.711 encoding is about 135 bytes (see Appendix I).

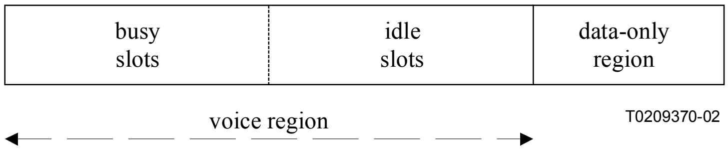

Figure 8-2/E.681 – A voice region with a single section of contiguous idle slots

It is possible to reduce the need for data packet fragmentation by jittering some of the busy voice slots away from the data region as shown in Figure 8-2. By creating the largest possible sequence of contiguous idle voice slots for data transmission, efficiency is improved. However, such jittering of any voice slots assigned to voice connections will introduce additional packet delay variation. To minimize speech transmission delay, a budget of about 2 to 3 ms for packet delay variation is typically allocated to the CMTS. Hence, it is desirable to avoid the jittering of any assigned voice slots as much as possible.

In practice, jittering of busy voice slots to close holes is not always necessary. This is because the holes can be used for the transmission of small control packets such as for maintenance and contention for transmission requests. To estimate the capacity available for data with assigned voice slots left in-place without closing holes, the model presented in [L01] can be used.

9 DOCSIS 1.0 and DOCSIS 1.1 interoperability

DOCSIS 1.0 was originally developed to support best-effort high-speed data service over a cable-based access system. DOCSIS 1.1, a second-generation protocol, enhances upon DOCSIS 1.0 by providing additional functionality so that the performance requirements of delay/jitter sensitive traffic, such as IP telephony, can be met. In particular, this includes protocol features such as unsolicited grant service (see 7.2), data packet fragmentation (see 8.2), and payload header suppression.

To maintain interoperability, a DOCSIS 1.1 CMTS must be backward compatible with a DOCSIS 1.0 CM. Since DOCSIS 1.0 CMs generally do not support packet fragmentation, a DOCSIS 1.1 CMTS must therefore not attempt to fragment the packet transmissions from a DOCSIS 1.0 CM. As a result, there may potentially be some impact on the quality of service for real-time communication needs of DOCSIS 1.1 CMs, when both DOCSIS 1.0 and DOCSIS 1.1 CMs are provisioned to share the same upstream channel. To mitigate this impact, the *packed upward* voice slot assignment algorithm described in [L01] may be used. Also, given a number of DOCSIS 1.0 CMs to be supported on an upstream channel, there is an upper limit to the number of voice sources that can be supported on the same channel, or vice versa, as shown in [L01].

10 Upstream channel pooling

Since the bandwidth of an upstream channel is limited, the number of CMs that can be supported statically by a given channel is small. To increase the overall total number of CMs that can be supported by a CMTS, i.e. to increase the utilization of the upstream channels, it is necessary to pool these channels together for use by any CM. Thus, in addition to provisioning a CM initially to use a particular channel, a CMTS needs to be capable of assigning CMs dynamically to different channels, based on up-to-date traffic conditions.

10.1 Immediate upstream channel change

According to DOCSIS 1.1, when the CMTS determines to move a CM from the currently assigned upstream channel to another, it sends to the CM an Upstream Channel Change Request (UCC-REQ) message. In response, the CM transmits an Upstream Channel Change Response (UCC-RSP) message on the currently assigned channel to signal its readiness to use the new channel. After switching to the new channel, the CM typically performs maintenance functions to make any necessary adjustments to the timing, power, or frequency used. This process is referred to as *re-ranging*.

However, to minimize disruption (e.g. to prevent packet loss) to calls in progress on the currently assigned channel, the CM should be able to use the new channel directly without performing the time-consuming re-ranging. Such a short-cut procedure is referred to as an *Immediate Upstream Channel Change* (*IUCC*). To achieve this, the CM must know about the operating characteristics of the new channel. At periodic intervals (maximum 2 s), the CMTS transmits to all CMs information about each active upstream channel via an Upstream Channel Descriptor (UCD) message. A CM should cache the UCD information regularly to avoid re-ranging in switching channels. If a CM does not cache the UCD information (e.g. due to a lack of memory), it is recommended that the CMTS should send a unicast message to the CM right after the UCC-REQ message. This is so that the CM can obtain immediately the operating characteristics of the new channel.

10.2 Voice slot jitter

Jitter, also referred to as packet delay variation, is the deviation, measured in units of time, from the ideal or anticipated time of receipt of each packet. Voice slot jitter becomes a problem and needs to be controlled when an IUCC is performed. This is because when an IUCC is to be made for a CM, the relative voice slot positions assigned to the CM for existing connections must still be maintained within a maximum tolerated jitter interval between the current and new upstream channels. To meet the jitter constraint, the concept of a *jitter window* is used to define a set of contiguous voice slots in the voice region with acceptable jitter. To keep jitter small, jitter window size must be small. Appendix II provides an example and further discussion on the use of jitter windows.

10.3 Assignment algorithms

A CMTS assigns upstream channels and idle voice slots to CMs, on a call-by-call basis, in response to their requests for call establishments. A CMTS performs channel assignments for a CM on two occasions. Either:

1) the first call request from the CM; or

2) subsequent call requests whenever all the voice slots in the voice region of the channel currently assigned to the CM are occupied, thereby necessitating the IUCC procedure to be performed.

In either case, a channel is *assignable* only if it can accommodate the request, i.e. there are enough idle voice slots for any established calls plus the new call, all within the jitter constraint. When an IUCC needs to be performed, this means that the following two conditions must be satisfied. First, the number of idle voice slots in each jitter window in the new channel must be no less than the number of voice slots allocated to the CM in the corresponding jitter window in the currently assigned channel. Second, at least one of the jitter windows in the new channel has room to accommodate the new call in addition to existing calls.

In selecting a channel for assignment, a CMTS first searches all the channels under its control to find the ones that meet the above assignment criterion. Suppose that a CMTS has c upstream channels. Typically, c is no more than eight. Assuming that the channels are numbered consecutively and identified as 1 through c. The CMTS can conduct the search in one of two ways:

1) *Search in opposite directions*, i.e. upward from channel 1 for the first call from a CM, downward from channel c for a subsequent IUCC; or vice versa.

2) *Search in the same direction*, i.e. upward from channel 1 (or downward from channel c) for both the first call from a CM and a subsequent IUCC.

By using one of these search procedures, the CMTS marks a subset of the c channels that meets the assignment criterion. If this subset is empty, the new call is blocked.

Assuming a successful search, the CMTS selects one channel out of the marked subset for assignment to the CM with a new call. One of the following procedures may be used:

1) *Packed with first fit*: Assign the first channel in the specified search direction. For example, if the search upward procedure is used, then the lowest channel that meets the assignment criterion is selected.

2) *Minimally packed*: Assign a channel with the minimum number of idle voice slots, tie breaking by the specified search direction. Thus, suppose there are multiple channels that meet the assignment criterion and have the same minimum number idle voice slots. If the search upward procedure is used, then the lowest such channel is selected.

3) *Maximally spread*: Assign a channel with the maximum number of idle voice slots, tie breaking by the specified search direction.

4) *Random*: Assign a channel probabilistically.

These four procedures can be combined in different ways. For example, the "packed with first fit" procedure may be used for selecting a channel for the new first call from a CM, while the "the maximally spread" procedure for IUCCs in new subsequent calls from the same CM.

After the selection of a channel for assignment, the CMTS does the following. If an IUCC is involved, then all existing voice connections established on the currently assigned channel for the CM are first moved to the new channel. In this move, the relationship of jitter windows in the current and new channels is preserved. Thus, existing voice connections assigned to one jitter window in the currently assigned channel are maintained in the same jitter window in the newly selected channel. After moving all existing connections in this manner, an idle voice slot in the new channel is selected and assigned to the new call. Of course, in the case of a new first call from the CM, IUCC is not needed and the CMTS simply selects an idle voice slot in the channel and assigns it to the CM for the new call. Assuming that the voice slots in the voice region of a channel are numbered consecutively and identified as such, one of the following procedures can be used for the selection of an idle voice slot:

1) *Packed*: Select the lowest (or highest) idle voice slot.

2) *Random*: first select randomly a jitter window with at least one idle voice slot and then select the lowest (or highest) idle voice slot in the selected jitter window.

For a given maximum jitter tolerance, the different algorithms achieve different blocking probabilities [L00]. The maximally spread and random algorithms tend to give more favorable performance results.

11 History

This is the first issue of ITU-T Rec. E.681.

Bibliography

[CS01] CHOUDHURY (G.L.), SEGAL (M.): Cable Infrastructure to Carry Voice and Data – Performance Analysis, Traffic Engineering and Cable Head-end Node Design Design, *Proc. 14th ITC Specialists Seminar on Access Networks and Systems*, Girona, Spain, 25-27 April 2001.

[L00] LAI (W.S.): Upstream Bandwidth Allocation for Packet Telephony in Hybrid Fiber-Coax Systems, *Proc. 2000 Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS'2000)*, Vancouver, Canada, 16-20 July 2000, pp. 96-100. (Sponsored by the Society for Computer Simulation International.)

[L01] LAI (W.S.): DOCSIS-Based Cable Networks: Impact of Large Data Packets on Upstream Capacity, *Proc. 14th ITC Specialists Seminar on Access Networks and Systems*, Girona, Spain, 25-27 April 2001, pp. 57-67.

Appendix I

Example calculation of upstream channel capacity

As an example, the following is assumed:

1) *Channel characteristics*: A 3.2 MHz upstream channel operating with QPSK at a symbol rate of 2560 ksymb/s, resulting in a 5.12 Mbit/s channel. This is the raw bandwidth for upstream transmission.

2) *Protocol features*: G.711 encoding at 10 ms packetization interval with payload header suppression.

   In this case, a voice packet has a voice payload of 80 bytes. Including the overheads in various protocol layers and using the payload header suppression feature of DOCSIS 1.1, a voice packet may come up to 135 bytes in length. For 8-byte mini-slots, a voice slot will require 17 mini-slots, at a rate of 108.8 kbit/s. A 5.12 Mbit/s channel then yields a maximum of 47 connections, when the full channel bandwidth is used to support voice only. Similar calculations show that there are 44 connections for 16-byte mini-slot at 115.2 kbit/s. (The use of lower bit rate encoding schemes such as G.728 and others will give a higher number of connections.)

3) *Propagation delay*: Transmission time for the DOCSIS Ranging Request (RNG-REQ) message at 5.12 Mbit/s is 0.1 ms. Together with the maximum round-trip propagation delay of 1.61 ms, there is a maximum total of 1.71 ms.

For a 10-ms packetization interval, this means that 17.1% of the upstream bandwidth is needed for this purpose. As a result, in a pure-voice environment, the per-channel maximum allowable voice capacity is (100% – 17.1%) 5.12 Mbit/s = 4.24 Mbit/s. For G.711 encoding with payload header suppression, this translates to a maximum of 38 connections for 8-byte mini-slot at 108.8 kbit/s, or 36 connections for 16-byte mini-slot at 115.2 kbit/s.

4) In an integrated voice/data environment, the maximum allowable voice capacity is further reduced by the amount of bandwidth that needs to be dedicated for the support of data services. For example, when voice can use only up to 60% of the channel bandwidth, then for the above two scenarios, the corresponding maxima are 28 and 26 voice connections, respectively.

Since initial maintenance is a low-level activity, its bandwidth is derived from the data bandwidth in this case.

Appendix II

Example of jitter window

As an example, Figure II.1 shows a 10-ms frame with a 6.6-ms voice region followed by a 3.4-ms data-only region. It illustrates a typical scenario whereby the bandwidth in a frame is split in a ratio of approximately 2:1, for sharing between voice and data services. The voice region has fifteen voice slots, which is an odd number. The sizes of the two non-overlapping jitter windows are approximately equal, with seven voice slots in the lower window and eight in the upper window. The duration of the lower window is $6.6 \times (7/15) = 3.08$ ms, while that of the upper window is $6.6 \times (8/15) = 3.52$ ms.

By splitting the voice region into two approximately equal non-overlapping windows, and maintaining calls within the same jitter window in an IUCC procedure, jitter is limited to the duration of the jitter window. This duration is smaller than that of either the voice region or the frame. Furthermore, by arranging the jitter windows to be both non-overlapping and collectively covering the whole voice region, jitter is maximally reduced while access by a CM to channel bandwidth is maximized (and so call blocking is minimized). Also, all CMs will use the same set of jitter windows, hence simplifying system management.

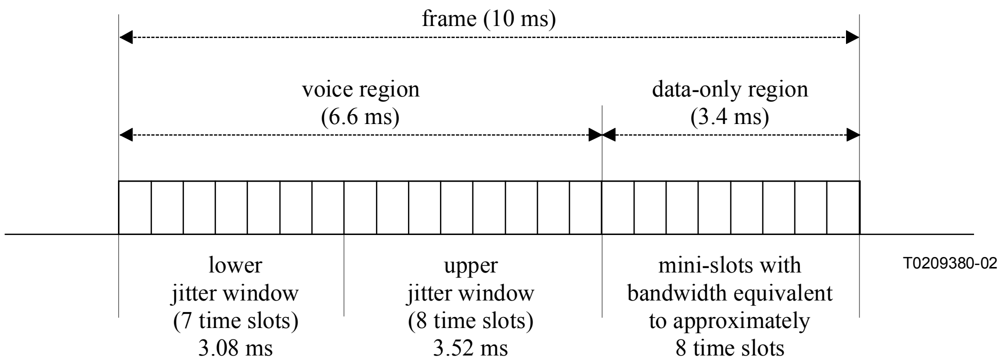

Figure II.1/E.681 – Example frame structure with two jitter windows

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of 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	Cable networks and 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 communications
Series Y	Global information infrastructure and Internet protocol aspects
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2002