International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

H.243
(10/2005)

SERIES H: AUDIOVISUAL AND MULTIMEDIA SYSTEMS

Infrastructure of audiovisual services – Communication procedures

---

Procedures for establishing communication between three or more audiovisual terminals using digital channels up to 1920 kbit/s

ITU-T Recommendation H.243

ITU-T

ITU-T H-SERIES RECOMMENDATIONS

AUDIOVISUAL AND MULTIMEDIA SYSTEMS

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

ITU-T Recommendation H.243

Procedures for establishing communication between three or more audiovisual terminals using digital channels up to 1920 kbit/s

Summary

This Recommendation concerns the system operation for a conference call between three or more audiovisual terminals conforming to ITU-T Rec. H.320.

This revised version of ITU-T Rec. H.243 introduces a number of enhancements and clarifications to the previous version, primarily the description on the usage of Unicode characters in H.320 systems.

Source

ITU-T Recommendation H.243 was approved on 7 October 2005 by ITU-T Study Group 16 (2005-2008) under the ITU-T Recommendation A.8 procedure.

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.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

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

ITU-T Recommendation H.243

Procedures for establishing communication between three or more audiovisual terminals using digital channels up to 1920 kbit/s

1 Scope

This Recommendation concerns the system operation for a conference call between three or more audiovisual terminals. It is applicable to a single MCU containing an audio mixer and a video switch according to the provisions of ITU-T Rec. H.231, to calls involving two such MCUs, and to calls involving three or more MCUs in a star configuration. It is also applicable to MCUs containing a video mixer. ITU-T Recs H.233 and H.234 address issues concerning the use of the Encryption Control Signal (ECS). ITU-T Rec. H.242 provides for communication between two audiovisual terminals connected point to point, using the frame structure defined in ITU-T Rec. H.221. ITU-T Rec. H.230 defines a number of Control and Indication signals, including those used in the multipoint communication procedures described in this Recommendation.

Three or more audiovisual terminals may be put into communication to form a conference call, by means of one or more Multipoint Control Units (MCU). The general principles of multipoint communication are described in ITU-T Rec. H.231. The means by which digital channels are established between terminals and MCUs, and between MCUs, is outside the scope of this Recommendation (see ITU-T Rec. H.200/AV.420).

It should be noted that the physical realization of an MCU may be such that two or more independent conferences may be set up within the same unit; logically, however, there is no relationship between these conferences; the text herein refers to an MCU only as a logical entity pertinent to the particular call of concern.

This Recommendation concerns only the flow of signals along the fixed digital paths, which may be at 64 kbit/s (56 kbit/s in certain networks) or multiples thereof up to 1920 kbit/s. The flow consists of a multiplex of audio, video, control and indication signals, and optional user data as defined in ITU-T Rec. H.221 which must be handled by the MCU in a way which is satisfactory to the users.

The signal multiplex on each path is fully in accordance with ITU-T Rec. H.221: the BAS commands define explicitly how the demultiplexer at the end of each link shall operate. Likewise, the basic procedures for initialization and mode switching are fully in accordance with those defined in ITU-T Rec. H.242 for point-to-point working. However, the composition of the multiplexed signal transmitted by each terminal and by the MCU is determined by terminal procedures and multipoint system procedures, as follows:

a) Terminal procedures are defined in service-specific system Recommendations, such as ITU-T Rec. H.320 for visual telephony.

b) Multipoint system procedures are defined in this Recommendation and are not of themselves service-specific.

c) ITU-T Rec. T.120: by making use of the T series, MCU and terminal procedures may be greatly enhanced, offering far more sophisticated specific applications to the user. Such enhancement is outside the scope of this Recommendation, although interactions with specific T.120 methods are described here.

For definitions of terms used in this Recommendation, see ITU-T Rec. H.231; for definitions of SBE and MBE symbols, see ITU-T Rec. H.230.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; 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. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

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

– ITU-T Recommendation G.722 (1988), *7 kHz audio-coding within 64 kbit/s*.

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

– ITU-T Recommendation H.221 (2004), *Frame structure for a 64 to 1920 kbit/s channel in audiovisual teleservices*.

– ITU-T Recommendation H.224 (2005), *A real time control protocol for simplex applications using the H.221 LSD/HSD/MLP channels*.

– ITU-T Recommendation H.230 (2004), *Frame-synchronous control and indication signals for audiovisual systems*.

– ITU-T Recommendation H.231 (1997), *Multipoint control units for audiovisual systems using digital channels up to 1920 kbit/s*.

– ITU-T Recommendation H.233 (2002), *Confidentiality system for audiovisual services*.

– ITU-T Recommendation H.242 (2004), *System for establishing communication between audiovisual terminals using digital channels up to 2 Mbit/s*.

– ITU-T Recommendation H.261 (1993), *Video codec for audiovisual systems at p × 64 kbit/s*.

– ITU-T Recommendation H.262 (2000) | ISO/IEC 13818-2:2000, *Information technology – Generic coding of moving pictures and associated audio information: Video*.

– ITU-T Recommendation H.263 (2005), *Video coding for low bit rate communication*.

– ITU-T Recommendation H.320 (2004), *Narrow-band visual telephone systems and terminal equipment*.

– ITU-T Recommendation H.321 (1998), *Adaptation of H.320 visual telephone terminals to B-ISDN environments*.

– ITU-T Recommendation Q.931 (1998), *ISDN user-network interface layer 3 specification for basic call control*.

– ITU-T Recommendation T.120 (1996), *Data protocols for multimedia conferencing*.

– ITU-T Recommendation T.122 (1998), *Multipoint communication service – Service definition*.

– ITU-T Recommendation T.123 (1999), *Network-specific data protocol stacks for multimedia conferencing*.

– ITU-T Recommendation T.124 (1998), *Generic Conference Control*.

– ITU-T Recommendation T.125 (1998), *Multipoint communication service protocol specification*.

– ISO 639-1: 2002, *Codes for the representation of names of languages – Part 1: Alpha-2 code*.

– ISO/IEC 10646 (2003), *Information technology – Universal Multiple-Octet Coded Character Set (UCS)*.

3 Conventions

In this Recommendation, the word "shall" is used to indicate a mandatory requirement, while the word "should" is used to indicate an option or a suggestion.

4 Capabilities and "selected" communication modes

The MCU shall send appropriate capabilities, according to the type of communication intended. Table 2/H.231 lists the MCU types classified as "Listed", showing the capabilities which each will be able to declare if the terminals connected have appropriate capabilities and the modes it will be able to transmit.

For each conference call, a "Selected Communication Mode (SCM)" is identified in the MCU. During the call, the MCU strives to maintain this SCM as that transmitted bidirectionally between itself and all terminals, and between itself and other MCUs. The exceptions are those channels where mode-0 forcing (ITU-T Rec. H.242) has been applied or, in some cases, where data is temporarily being transmitted. Table 1 gives some examples of SCM which could be specified for the listed MCUs of ITU-T Rec. H.231. The SCM may include one or more data channels.

Table 1/H.243 – Examples of selected communication modes

Transfer rate	MCU Type (according to Table 2/H.231)							
	A	B(d)	C	C(d)	D	D(d)	Code	
64/56 kbit/s	a8+v						a1	A-law,0U
		a6+d6400					a2	μ-law,0U
		a7+d8000					a3	A-law,0F
2 × 64/56 kbit/s			a3/4+v	a8+v+d6400			a4	μ-law,0F
			a7+v	a7+v+d8000			a5	Rec. G.722, m1
							a6	Rec. G.722, m2
128 kbit/s					a8+v		a7	Rec. G.722, m3
					a7+v		a8	Rec. G.728
							v	H.261-ON
384 kbit/s					a6+v	a6+v	d6400	LSD
						a6+v+d64k	d8000	LSD
							d64k	HSD
							r	restrict

The following sample methods may be used to determine the SCM; other methods are feasible:

1) The SCM may be fixed as a permanent feature of the MCU as manufactured.

2) The MCU may provide for several possible values of SCM, and one of these is specified by the service provider or at the time of booking the call.

3) The SCM is selected automatically within the MCU according to the capabilities of the terminals connected.

NOTE 1 – The algorithm for this selection is outside the scope of this Recommendation; for example, the SCM is set at the value transmitted by the first terminal to access the MCU; or the highest common mode of all primary terminals is selected; or the SCM is set at the value transmitted by the chair-control terminal, if any.

4) The SCM is set by utilizing procedures effectuated using T.120-series protocols.

In the cases 2, 3 and 4 it is also possible that the SCM may change during the course of the call. It is not mandatory that the MCU have all of these methods available but it should be clear to the user and service provider what method is in force.

NOTE 2 – Where there are several modes common to the connected terminals (e.g., both 16 kbit/s and wideband audio), implementers should document what method is used to determine which one is selected. It should be noted that the power of the MCU to control the mode transmitted by a terminal is limited to a determination of the rates of audio, video and any data paths; it cannot set the video mode CIF or QCIF except by limiting the SCM to QCIF via capability changes, nor the audio mode where there is a choice at one bit rate (for example: ITU-T Recs G.711 or G.722 at 56 kbit/s). Specifically, if the MCU sends the BAS command for ITU-T Rec. G.711 at 56 kbit/s, the terminal may respond with either ITU-T Rec. G.711 at 56 kbit/s or with ITU-T Rec. G.722 at 56 kbit/s without this being considered a violation of MCC. In the situation just described, a terminal that responded with ITU-T Rec. G.728 at 16 kbit/s would be in violation of MCC since G.728 and G.711-56 use different bit rates within the H.221 multiplex.

When advanced video modes such as H.262/H.263 are used, the MCU may make use of Multipoint Mode Symmetrize (MMS), which is supported by all terminals with such advanced video capabilities to require video mode symmetry, including CIF/QCIF symmetry, for all modes. When in receipt of MMS, terminals shall respond to the MCU with whatever video, audio, and/or data modes they receive. If the MCU does not make use of MMS, terminals are free to follow the mode asymmetry allowed by MCC, e.g., to send H.263 while receiving H.261.

The MCU is bound by the hierarchy of advanced video modes as described in Figure A.1/H.320 such that an MCU with H.263 capability shall have H.261 capability, and an MCU with H.262 capability shall have H.263 and H.261 capability. The MCU may choose to exclude a video capability for the purpose of maintaining a particular SCM, during which the hierarchy would be violated, e.g., the MCU might have only H.262 in the capability set. However, MMS/H.230 should be used for maintaining control over video modes. Note that MMS support is required for all terminals that support advanced video modes such as H.262 and H.263.

At the start of a call, the capset transmitted by the MCU shall be that which corresponds to SCM; for brevity this is referred to below as "SCM-cap". Any terminal which is found, from its declared capset, not to be capable of transmitting the SCM will be given secondary treatment (see 13.1). Secondary status may be audio-only, or audio/data only, or some other status as determined by the MCU manufacturer. Later on in the call the MCU must transmit such capabilities as meet its immediate needs with respect to controlling the information transmitted to it; this may still be SCM-cap, or indeed there may be a change of SCM itself during the call.

The conference participants may initiate out-of-band changes if the network and MCU support them. The connection of single- and multiple-channel (e.g., 2B and 128 kbit/s or 6B and 384 kbit/s) terminals into the same multipoint call is provided by SM-comp and 6B-H0-comp as defined in ITU-T Rec. H.221.

The **only** condition that removes a capability from a terminal's capability set is the transmission of a new capability set lacking the capability being removed. Other events, such as a brief loss of framing, initiation of loopbacks, movement to secondary status, etc. shall not be assumed by the MCU to have resulted in a change to a terminal's capability set. See ITU-T Rec. H.242 for a detailed discussion of capability exchanges.

The MCU should not assume that because a terminal supports MBEs that it is capable of supporting all or any particular sub-set of standard MBE messages except as is indicated by its capability set.

Terminals should, to allow the MCU to function more effectively, declare their full capability set and keep that set in effect without unnecessary changes that may result in an oscillating SCM.

A terminal may use the procedure of 9.5/H.242 to request that a particular mode or allocation of bit rate (e.g., ITU-T Rec. G.728) be used in the SCM. The MCU should attempt to meet this request, but is not obligated to do so. A slave MCU shall pass such requests to the master MCU for action. MCUs are not required to provide exactly what the terminal has requested; a "best effort" may be attempted. The method by which an MCU determines such a "best effort" is left to the manufacturer.

5 Initialization procedures for establishing communication between standard terminals and an MCU

The initialization procedure is very similar to that between two terminals, as specified in ITU-T Rec. H.242.

All terminals shall synchronize their outgoing transmissions to the bit rate incoming from the local MCU when they receive MCC from the MCU. Terminals should avoid mode 0 forcing when in receipt of MCC from an MCU unless justified for communication reasons. The following description of the procedure assumes that video and data are included in the communication; however, for conference calls not including video, or data, or both, essentially the same applies.

When the connection has been established, each terminal transmits a signal according to ITU-T Rec. H.221; it sends its capabilities and is waiting to receive frame structure and capabilities, as described in ITU-T Rec. H.242, with transmission in mode 0F only. There may then be additional procedures, such as identity checking, to be carried out before the terminal is added to the conference but after the MCU sends MCC to the terminals.

5.1 First terminal added to conference

The MCU transmits its SCM-cap and the C&I symbols MCC and MIZ (see ITU-T Rec. H.230) immediately after the final cap-mark and command (see Table 2/H.242), indicating that a conference call is being set up, that no other terminals are yet connected and the user should wait.

Until the audio coding law for the terminal is determined and if necessary, set properly, the MCU should not add the terminal's audio signal to the audio mix. At the discretion of the manufacturer, the MCU should also do one of the following:

1) Send the audio-off command to the terminal.
2) Send silence or an optional audio-message to the terminal.
3) Send audio and rely on the terminal to mute the output.

In the above list, 1 and 2 are to be preferred over 3 since not all terminals will reliably mute the output in this circumstance. However, 3 has been included for reasons of backward compatibility with earlier versions of this Recommendation.

The MCU finds incoming frame alignment and registers the capability of this first terminal, which it designates T_A; if the terminal is not capable of transmitting the SCM, a different procedure is followed (see 13.1).

The video seen by the first terminal is up to the discretion of the MCU manufacturer. VIR can be ignored without any ill effect.

Data

– If the terminal has requested that an LSD or HSD channel be opened, the MCU may open the channel at this time at the discretion of the manufacturer. The MCU should provide idle bits to non-H.224-capable terminals and it should connect H.224-capable terminals to the data mixer in loopback mode.

– If the terminal is T.120 or H.224-MLP[1]-capable, the MCU may open the MLP channel to the terminal and connect it to a data conferencing unit.

The MCU may defer opening of MLP data channels to a later time, such as when a predetermined number of terminals is present.

5.2 Second terminal added to conference

The MCU transmits its SCM-cap followed by MCC indicating that a conference call is being set up.

Until the audio coding law for the terminal is determined and if necessary, set properly, the MCU should not add the terminal's audio signal to the audio mix. At the discretion of the manufacturer, the MCU should also do one of the following:

1) Send the audio-off command to the terminal.
2) Send silence or an optional audio message to the terminal.
3) Send audio and rely on the terminal to mute the output.

In the above list, 1 and 2 are to be preferred over 3 since not all terminals will reliably mute the output in this circumstance. However, 3 has been included for reasons of backward compatibility with earlier versions of this Recommendation.

The MCU finds incoming frame alignment and registers the capability of the second terminal T_B; if the terminal is not capable of transmitting the SCM, a different procedure is followed (see clause 13). When the MCU is receiving A = 0 from both terminals, the audio and video paths are set up as follows.

Audio

– Both the (decoded) audio signals are connected to the audio mixer; the symbol Cancel-MIZ is sent to terminal T_A.

– The normal working audio command is transmitted and the appropriate mixer outputs are connected at the next submultiframe.

Video

– If video signals are being received from either or both terminals, they are forwarded via the video switch using the proper mode switching procedure (see ITU-T Rec. H.242) and a fast update request VCU (see ITU-T Rec. H.230) is sent towards the transmitter(s) of those signals.

– If VIR (see ITU-T Rec. H.230) is received from either or both terminals, this is forwarded.

Data

– If one of the terminals has requested that an LSD or HSD channel be opened, the MCU may open the channel at this time.

– If both terminals are T.120 or H.224-MLP-capable, the MCU may open the MLP channel to the second terminal, and connect both to a data conferencing unit.

[1] Note that H.224-capable terminals are required to be able to make use of both the LSD and the MLP channels.

The MCU may defer opening of data channels to a later time, such as when a predetermined number of terminals is present.

5.3 Third terminal added to conference

The MCU transmits its SCM-cap followed by MCC indicating that a conference call is being set up.

Until the audio coding law for the terminal is determined and if necessary, set properly, the MCU should not add the terminal's audio signal to the audio mix. At the discretion of the manufacturer, the MCU should also do one of the following:

1) Send the audio-off command to the terminal.
2) Send silence or an optional audio-message to the terminal.
3) Send audio and rely on the terminal to mute the output.

In the above list, 1 and 2 are to be preferred over 3 since not all terminals will reliably mute the output in this circumstance. However, 3 has been included for reasons of backward compatibility with earlier versions of this Recommendation.

The MCU finds incoming frame alignment and registers the capability of the third terminal T_C; if the terminal is not capable of transmitting the SCM, a different procedure is followed (see 13.1). When the MCU is receiving A = 0 from T_C, the audio and video paths are set up as follows.

Audio

– The (decoded) audio signal is connected to the audio mixer.
– The normal working audio command is transmitted and the appropriate mixer output is connected at the next submultiframe.

Video

– If video signals are being received from either or both terminals T_A, T_B, one of them (default T_A) is transmitted to T_C via the video switch using the proper mode switching procedure (ITU-T Rec. H.242) and a fast update request VCU (see ITU-T Rec. H.230) is sent towards the transmitter(s) of that signal; when video is received from T_C, it may optionally be transmitted to both T_A and T_B, in which case VCU is sent to T_C.
– If VIR is received at the MCU, it may be ignored without ill effect.

Data

– If an LSD or HSD channel is open in the conference, the channel should be opened to the new terminal.
– If the new terminal is T.120 or H.224-capable, the MCU may open the MLP channel to the third terminal, and connect it to the data conferencing unit.

5.4 Fourth and subsequent calls added to conference

The procedure followed is essentially that of 5.3 above.

5.5 Extension to multiple channels

If the intended SCM of the conference communication involves multiple channels, then the transmitted MCU transfer-rate capability reflects the appropriate rate to all terminals, and the additional channels are set up according to the procedures defined in ITU-T Recs H.242, H.221, and Q.931 and/or 7.2 as appropriate.

Having received MCC, the terminals cannot transmit at the higher transfer rate until the MCU does so, which could be when the other terminals are all ready, or after a time-out, or when at least one terminal has all the requested additional channels available; the MCU itself adopts the higher rate and the terminals shall follow suit. In the event that some connections do not reach the desired number of additional channels, the MCU may downrate the corresponding terminals to secondary status by sending them MIS, or dropping them or by following some other procedure as determined by the manufacturer, and proceed normally with the rest.

5.6 Other standard terminals

The support of non-H.320 terminals by an H.243 MCU is optional.

Ordinary PSTN terminals that do not support ITU-T Rec. H.221 may optionally be supported by the MCU. Such calls may make use of the same network addresses as the H.320 terminals, or different network addresses as determined by the manufacturer. As each call arrives, it is added to the audio mixer. The manufacturer may optionally require a DTMF password prior to attachment to the audio mixer. The manufacturer may optionally transcode to the current audio coding, or change the SCM to include the appropriate G.711 coding. If the PSTN call is first in the conference, treatment is determined by the manufacturer (either loop-back or no connection to the audio mixer).

T.120 terminals that do not support ITU-T Rec. H.221 may optionally be supported by the MCU. Such connections may make use of the same network addresses as the H.320 terminals, or different network addresses as determined by the manufacturer. As each call arrives, it is added to the T.120 data conference. It is the responsibility of the MCU data conferencing module to use T.123 flow control to ensure that the data rate in the non-H.221 connections matches the current T.120 data rate which is less than or equal to the MLP channel rate specified by the conference SCM. Examples of non-H.221 connections include PSTN calls using the PSTN stack of T.123 or LAN connections using one of the various LAN stacks of ITU-T Rec. T.123.

H.321 (H.320 on Broadband ISDN (B-ISDN)) terminals may also be optionally supported by the MCU. Since these terminals are H.320-compatible, the procedures for H.320 terminals shall be followed. The MCU shall support B-ISDN parameters as described in ITU-T Rec. H.321 if this optional feature is provided.

5.7 MCU-MCU interconnections

Cascading of MCUs is an optional feature. Also, master/slave operation is an optional mode of operation for cascaded MCUs which is required for many features, but not for simple cascading. Finally, the construction of a slave-only MCU that does not return MIM is allowed, but note that a cascade of slave-only MCUs is not possible except via simple cascading.

5.7.1 Initialization

The following applies to interconnection of two, three or more MCUs.

After each initial channel between MCUs has been established, each MCU sends its current capability set as it would to a terminal, following this up with MCC as in 5.1. Each becomes aware of the presence of the other MCU by virtue of receiving MCC from it.

During call set-up, the circumstance may arise that one MCU is connected only to another and not to any terminal: it should not then transmit an optional audio message as in 5.1, but the symbol MIZ (see ITU-T Rec. H.230) is included in the outgoing BAS codes. Then, when a terminal is first connected, Cancel-MIZ is transmitted to the other MCU, and the audio from that terminal is forwarded via the mixer.

When an MCU has established communication with a first terminal, and on a second port is receiving MIZ in conjunction with MCC, it does not proceed to 5.2 until either it receives Cancel-MIZ on the inter-MCU port or is itself connected to a second terminal.

In general, each MCU treats each other MCU as though it were a terminal, making decisions as to transmitted mode and BAS values according to the incoming capabilities, and switching video as appropriate. However, it is necessary to ensure that, where there is a choice of bit rate for video, the same is chosen by both MCUs. When one MCU has been designated master, the slave shall symmetrize to the rates transmitted by the master.

In basic operation where no master is specified (that is, apart from the enhanced procedures described in clauses 7, 8 and 9) the following are mandatory for unrestricted operation:

a) If both MCUs have declared G.722-48, the transmission of audio shall be G.722 at 56 kbit/s only.

b) If both MCUs have declared G.728 audio and G.722-48, 56 kbit/s G.722 shall be used.

c) If both MCUs have declared G.728 but not G.722 audio, G.728 shall be used.

d) If the declared video modes are asymmetric, the highest common mode according to the video algorithm hierarchy of Figure A.1/H.320 and Annex A/H.320 shall be used by both MCUs.

For restricted operation, substitute G.722 at 48 kbit/s for G.722 at 56 kbit/s in the above cases.

An MCU M1 receiving VCU from another MCU M2 shall forward this symbol to the terminal or other MCU which is the current source of the video it is transmitting to M2; if M1 receives VCF from M2, it shall forward this symbol to those terminals and/or MCUs (if any) to which it is sending the video signal received from M2. An MCU should respond to VCU and VCF from terminals in the same way as it responds to those symbols from MCUs.

5.7.2 Designation of master MCU

In order for the procedures of clauses 7, 8 and 9 to operate properly in the cascaded case, one of the MCUs must take on the master role. If the procedures of clauses 7, 8 and 9 are not used, a master/slave relationship need not be established for simple operation as described in 5.7.1. It may be desired for operational reasons that for a star configuration of 3 or more MCUs (see clause 6/H.231) the master be that at the centre, but this is not required if all the MCUs are MIH-capable.

MCUs should designate a master in a cascade even if the procedures of clauses 7, 8 and 9 are not being used in order to avoid overuse of the contention resolution principle. MCUs should use the contention resolution principle once to establish mastership, and thus avoid additional use of the contention resolution principle.

5.7.2.1 Master designated prior to the call

When an MCU pre-administered to be master is connected to another (which it recognizes by receipt of MCC therefrom), it sends MIM, ignoring any MIM signal received as a result of procedure of 5.7.2.2. If a call involves two MCUs that have both been configured as master MCU prior to the call, a manual procedure may be followed to resolve the conflict, or the contention resolution principle of 13.2 may be followed as described in 5.7.2.4.

5.7.2.2 Automatic designation of master in the dumb-bell case

If an MCU receives the symbol MCC at one port, and has not also received MIM at that port, it carries out the contention resolution procedure of 13.2; if the result is that it sent a lower number than it received, it assumes the master role and transmits the symbol MIM to the other.

5.7.2.3 Automatic designation of the master when 3 or more MCUs are involved

This subclause is based on the assumption that a master has been established initially, either via pre-administration, or via the procedure of 5.7.2.2. For consideration of the case when two domains of MCUs, each with its own master, are merged, see 5.7.2.4.

In the dumb-bell or star configuration, where each slave is directly connected to the master, as each new MCU attaches to the master and receives MCC and MIM, it becomes a slave.

For situations where MCUs are not connected directly to the master, the procedures are the same, with an MCU that is already a slave sending MIM to an MCU as it connects to the hierarchy. From the viewpoint of the new MCU, the "slave" is its master, but the "slave" MCU merely passes commands up the hierarchy and awaits the true master's response. No more than one MIM shall be processed at the same time by a given slave to avoid confusion.

Note that slave MCUs that are not connected directly to the master must have the Multipoint Indicate Hierarchy (MIH) capability in their capset. This capability allows an MCU to distinguish between an MCU that cannot support multi-level hierarchies and one that can. Since multi-level hierarchies involve additional operations on both the master and each level of slave, all MCUs in a hierarchy of three levels or more must have MIH capability.

Such master/slave networks need not have the master MCU in the centre; the master may be at the "root" of a tree of slave MCUs. This procedure should not be used without the use of ITU-T Rec. T.120 for control for more than a total of 3 layers of MCUs due to the BAS processing delays that accumulate as the hierarchy grows.

5.7.2.4 Domain mergers for master/slave domains

Whenever a group of MCUs with a designated master (perhaps a single MCU pre-administered as being a master) is joined to another such group, the conflict may be removed via administration. However, the conflict may automatically be removed via the use of the contention resolution principle. The following cases exist:

a) The masters are directly connected: Each receives MIM, and the contention is resolved via the use of the contention resolution principle. One MCU and its entire sub-tree of slaves becomes a slave of the winning MCU.

b) The masters are connected via one or more slave MCUs.

   NOTE – These procedures are supported by MCUs with MIH capability.

   When the slave receives a second MIM, it forwards it up the tree to the master, which responds by making use of the contention resolution principle. A slave shall store the path by which MIM came to that slave so that it will be able to conduct the contention resolution principle with the correct port. No more than one application of the contention resolution principle can be under way at a given point in time on a particular MCU. The timers described in the contention resolution principle should be extended by an amount determined by the manufacturer. The amount the timers are extended should be a function of the number of MCUs in the conference.

MCU re-numbering during domain merger is described in 7.3.1.1.

5.8 Closure of conference

If the conference is closed by sequentially dropping terminals, then when only one remains connected it should be sent MIZ to allow the user to understand explicitly the reason for loss of video from the dropped terminal, etc.

5.9 User-initiated connection establishment

Terminals may dial in to an MCU that is appropriately configured for a dial-in conference. The MCU may also be constructed to support dynamic dial-out as well as pre-set dial-out to establish connections. Dynamic dial-out operations may be controlled, for example, either via the use of T.124 operations for T.120-equipped terminals, or via the use of so-called "BAS DTMF" commands using the numerics described in ITU-T Rec. H.230 for non-T.120-equipped terminals, or via some other method. The sequence of numerics used to initiate dial-out during an ongoing conference is beyond the scope of this Recommendation. Methods for dealing with the capacity exhaustion in the MCU that may result from user-initiated dial-out are also beyond the scope of this Recommendation.

6 Video switching and mixing

6.1 Video switching procedure

There are two cases to be considered: in some MCUs the video signal is switched without any processing, whereas in others the video may be processed such that when the switching is carried out there is no discontinuity in the error-correction framing in the outgoing signals.

6.1.1 No video mixing

When it is decided within the MCU that terminal A, currently receiving the video signal from terminal B, should instead be sent from terminal C, the following procedure is used (codes VCF, VCU are specified in ITU-T Rec. H.230).

a) The MCU transmits VCF to terminal A at an appropriate moment, and then switches video such that the picture from C is transmitted towards A.

b) Terminal A receives VCF, and freezes its currently displayed picture; it ignores subsequent decoded video information, but continues to track the error-correction framing, and to monitor Picture Headers for the Freeze Picture Release command.

c) When incoming video to A changes from B-picture to C-picture, error-correction frame alignment is lost, and will take a time T to recover, dependent on the video bit rate and other factors (refer to ITU-T Rec. H.242).

d) After a time greater than T, the MCU transmits VCU to terminal C.

e) On receipt of VCU, terminal C sends its next video frame in "fast-update" mode (see 4.3.2/H.261, H.262 and H.263), together with the Freeze Picture Release command[2].

f) On receipt of the Freeze Picture Release command, terminal A reverts to displaying the incoming decoded picture.

NOTE – Users at other terminals which have been receiving picture C continuously during the above procedure will nevertheless be aware of the switching action because of the use of the fast-update mode: this is the transmission of a single new picture over a period inversely proportional to video bit rate – at 320 kbit/s this period is likely to be about 0.5 seconds.

An MCU capable of video mixing may also be able to speed up video switching by eliminating error-correction framing changes, thereby eliminating the need for resynchronization in the terminal; this requires no additional standardization and is left to the discretion of the manufacturer.

6.1.2 With video mixing

The method of switching for a video mixing MCU is as follows.

When it is decided within the MCU that terminal A, currently receiving the video signal from terminal B via the video mixer, should be sent from terminal C instead of terminal B, the following procedure is used (codes VCF, VCU are specified in ITU-T Rec. H.230):

a) The MCU informs the video mixer that the signal from terminal B is being switched away at an appropriate moment, and then switches video such that the picture from C is transmitted towards the video mixer.

[2] In both ITU-T Recs H.261 and H.263, the Freeze Picture Release is contained in the PTYPE field. In ITU-T Rec. H.262, the Freeze Picture Release is sent as described in Annex A/H.320. In all three cases, the Freeze Picture Release is contained in the video stream.

b) Terminal A continues to receive video framing, and notices no discontinuity from this switch. However, the video mixer may insert FEC fill blocks, or a blank screen or other image (e.g., a loop repeating the previous complete image) while it is finding error-correction frame alignment from terminal C. This recovery will take a time T to occur, dependent on the video bit rate and other factors (see ITU-T Rec. H.242).

c) After a time greater than T or when framing is recovered, the MCU transmits VCU to terminal C.

d) On receipt of VCU, terminal C sends its next video frame in "fast-update" mode (see 4.3.2/H.261, H.262 and H.263), together with the Freeze Picture Release command. The video mixer is responsible for removing the Freeze Picture Release command from the H.261, H.262, or H.263 stream. When freeze picture release arrives from terminal C, the video mixer stops transmitting whatever image it has been sending toward terminal A and sends the image from terminal C.

Note that in the above procedure, switches occur on picture boundaries. More than one sub-part of the mixed image, up to and including the total number of images being mixed, may be switched at the same time. This is accomplished by repeating the above steps in parallel for each terminal being switched.

6.2 Video mixing

A generally desirable feature of audiovisual conferencing is the ability to simultaneously view more than one site other than the viewer's site, optionally including the viewer's site as well. The MCU can provide this feature which will be referred to as video mixing without regard to how many sites are visible at a time. The architecture of this subclause places video mixing at the MCU. While the terminals are in receipt of a mixed video image, they may send a mixed image (e.g., split screen) but note that two levels of video mixing may result in an image of little usefulness.

Figure 1 shows four independent H.221 multimedia multiplexes entering an MCU, being demultiplexed into their composite media streams, and the video streams being mixed by a video mixer to produce a single composite image. There are two general ways the video processor can operate:

– by implementing a full video decode/encode cycle with picture composition in the pel domain;

– by implementing a partial or zero decode combined with picture composition with the intent to lower delay, possibly making use of a modified H.320 terminal.

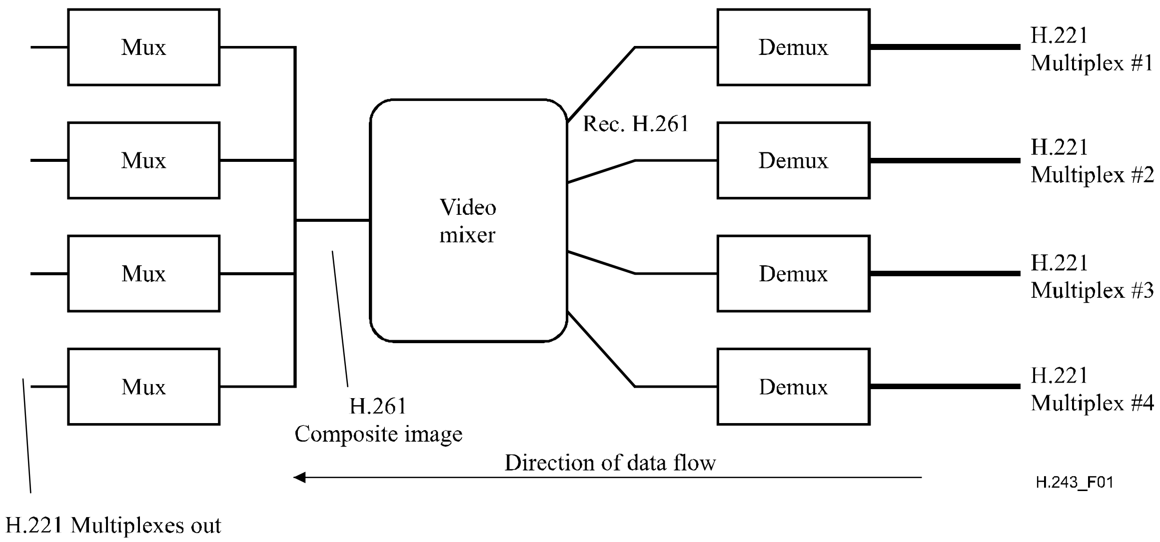

Figure 1/H.243 – Video mixing example: Quadrature

In either case, some MCUs may make use of the asymmetric nature of ITU-T Rec. H.221 to mix incoming QCIF images into an outgoing CIF image. In both cases, a video mixing MCU should operate properly with terminals conforming to the 1990, 1993 and 1996 versions of ITU-T Rec. H.320, and also supporting CIF. QCIF only terminals may not be suitable for video mixing since QCIF image resolution may not be appropriate for the display of a mixed image. Note that standardized mixed picture layouts are only defined for CIF images.

Systems involving modifications (relative to the 1990/1993/1996 versions of ITU-T Rec. H.320) to H.320 terminals are for future study, as are systems based on video mixing in terminals. The mixing procedures described here apply to ITU-T Recs H.261, H.262 and H.263 for any input image format, but only a CIF output format is specified. Mixing procedures for output formats other than CIF are for further study.

Without regard to how the video mixing is done by the MCU, a set of controls are needed analogous to VCB, MCV, and VCS that allow the user control over the images presented by the MCU. These control signals are the same for all methods of picture composition. Since these commands can become complex, they are standardized as part of the T.120 series of Recommendations using the MLP channel. See clause 15 for a discussion of the interactions between H.243 commands and T.120 commands with regard to video mixing.

The MCU may provide only a single mixed image per conference, or it may provide a larger number such as, for example, a different mix for each terminal. Control over multiple mixed images is left for further study as part of the T.120 series of Recommendations.

Terminals not capable of issuing picture composition commands using T.120 will receive a picture mixed according to rules established by the manufacturer. Examples of such rules are:

a) include in the mixed picture the loudest N speakers;

b) include in the mixed picture the first N parties in the conference;

c) include in the mixed picture the most recent N speakers; and

d) include in the mixed picture a pre-determined list of sites.

The details of this mechanism are beyond the scope of this Recommendation, but the MCU with video mixing capability shall provide some method by which spatial arrangement of images is accomplished in the absence of directives from enhanced terminals. Whether the viewing party is included in the composition is left to the discretion of the manufacturer.

6.2.1 Methods of image arrangement

There are eight defined image arrangements described in Figures 2 to 4 and in Tables 2 and 3. Table 2 shows the numbering of LSD commands in Table A.1/H.221 and Table 3 shows the numbering of HSD commands in Table A.2/H.221. An MCU may support all or any subset of the defined methods.

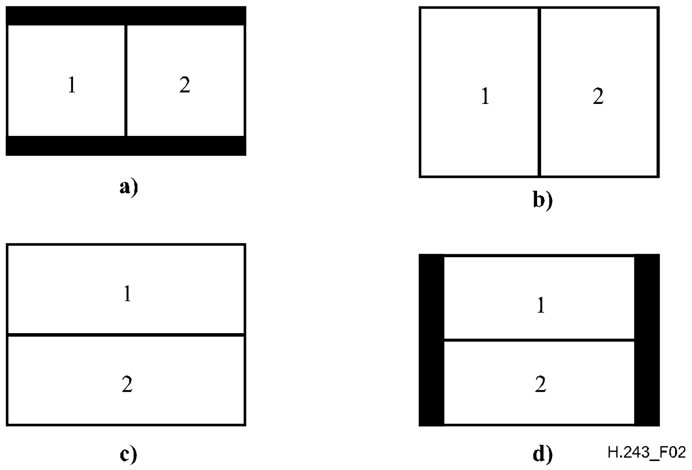

Figure 2/H.243 – Sub-picture location numbering for dual-view

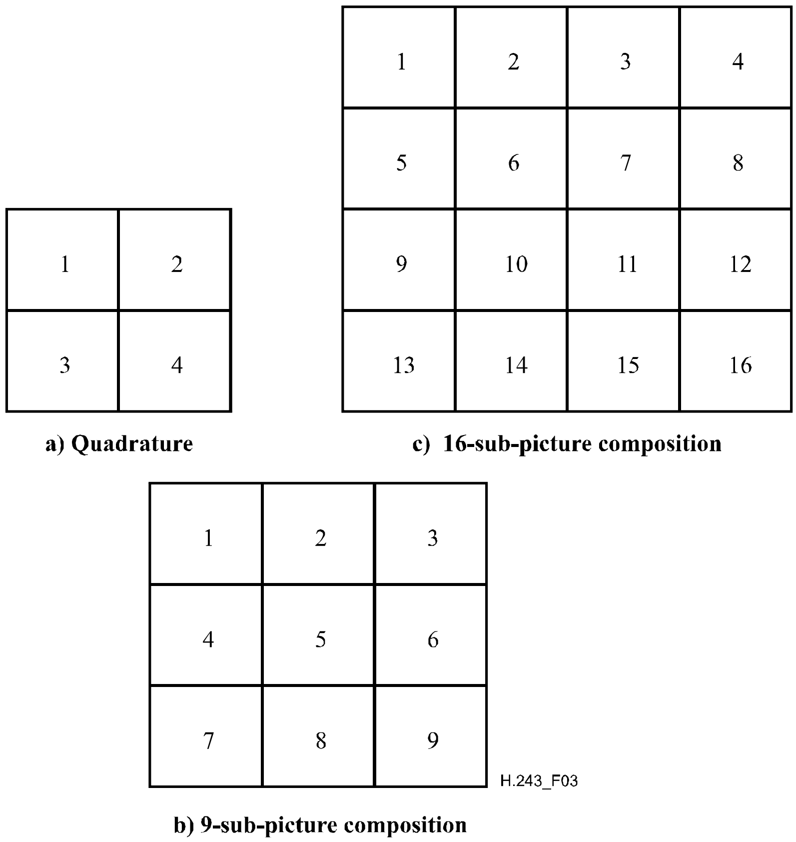

Figure 3/H.243 – Sub-picture location numbering for quadrature, 9-sub-picture composition, and 16 sub-picture composition

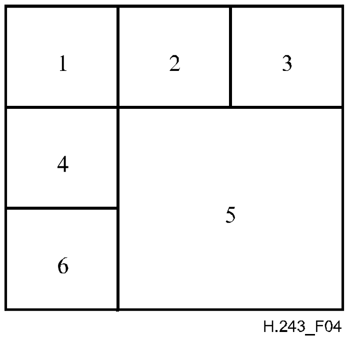

Figure 4/H.243 – Sub-picture location numbering for mixed-view

Table 2/H.243 – Encoding for <B> in {DCA-L, <B>} using SBE/num of ITU-T Rec. H.230

Value of <B>	Requested rate
0	Reserved
1	300 bit/s LSD
2	1200 bit/s LSD
3	4800 bit/s LSD
4	6400 bit/s LSD
5	8000 bit/s LSD
6	9600 bit/s LSD
7	14 400 bit/s LSD
8	16 kbit/s LSD
9	24 kbit/s LSD
10	32 kbit/s LSD
11	40 kbit/s LSD
12	48 kbit/s LSD
13	56 kbit/s LSD
14	62.4 kbit/s LSD
15	64 kbit/s LSD
16-30	Reserved
31	var-LSD
32	MCU will pick highest common rate
33	MCU will pick lowest common rate
34	Use current channel rate
35	T.120 and H.224 in MLP channel using token control
36-255	Reserved

Table 3/H.243 – Encoding for <B> in {DCA-H, <B>} using SBE/num of ITU-T Rec. H.230

Value of <B>	Requested rate
0	Reserved
1	var-HSD (R)
2-16	Reserved
17	64 k HSD
18	128 kbit/s HSD
19	192 kbit/s HSD
20	256 kbit/s HSD
21	320 kbit/s HSD
22	384 kbit/s HSD
23	768 kbit/s HSD
24	1152 kbit/s HSD
25	1536 kbit/s HSD
26	var-HSD
29-31	Reserved
32	MCU will pick highest common rate
33	MCU will pick lowest common rate
34	Use current channel rate
35-255	Reserved

The sub-picture formats shown in Figures 2, 3 and 4 allow:

1) the terminal to overlay terminal identity strings on the correct sub-picture; and
2) the terminal to perform further sub-processing if desired.

The compositions are defined for CIF only, and all sub-picture boundaries are on H.261/H.262/H.263 macroblock boundaries. Pel boundaries can be computed from the macroblock boundaries[3]. For two cases, the picture is cropped before being mixed, and the nature of the crop is defined, the nature of any required visible border, its colour, etc. is left to the discretion of the manufacturer. Unlike in ITU-T Recs H.261/H.262/H.263, this subclause utilizes grid numbering of macroblocks, with (1,1) in the upper left corner and (18,22) in the lower right corner where for (A,B) where A is the row numbering and B is the column numbering as shown in Figure 5.

[3] Note that although macroblocks are the same in ITU-T Recs H.261, H.262 and H.263, H.261 and H.263 groups of blocks (GOB) are not. The H.262 equivalent to the H.261/H.263 GOB is the slice, which is also defined differently from the GOB.

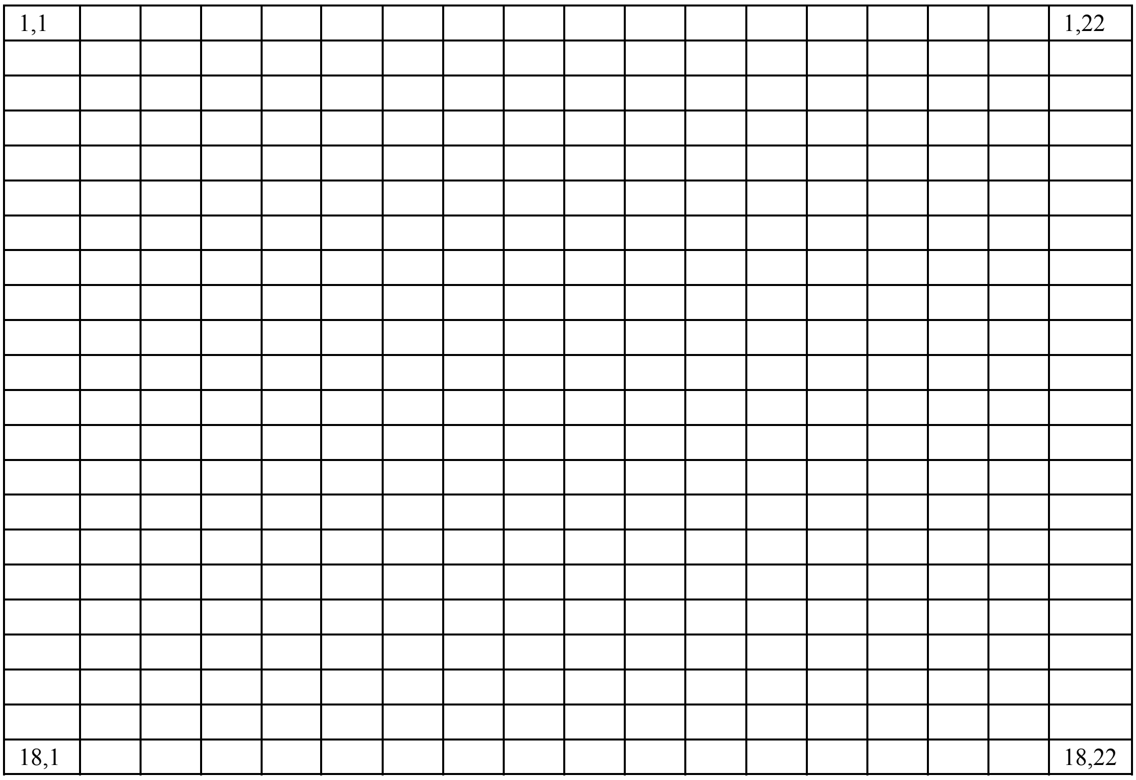

Figure 5/H.243 – H.261 macroblock numbering for video mixing

a) Dual-View as shown in Figure 2-a: An MCU with this capability transmits two images side by side with unaltered aspect ratio, blanking out the top 4 macroblocks and the bottom 5 macroblocks.

b) Dual-View as shown in Figure 2-b: An MCU with this capability transmits two cropped images side by side with an altered aspect ratio; no blanking is present. The area cropped by the MCU on each of the original CIF images is defined as the leftmost 5 macroblocks and the rightmost 6 macroblocks of each image prior to composition.

c) Dual-View as shown in Figure 2-c: An MCU with this capability transmits two cropped images one over the other with an altered aspect ratio; no blanking is present. The area cropped by the MCU on each of the original CIF images is defined as the topmost 4 macroblocks and the bottommost 5 macroblocks of each of the original CIF images.

d) Dual-View as shown in Figure 2-d: An MCU with this capability transmits two images with unaltered aspect ratio one over the other with bordering. The border on the mixed CIF image consists of the leftmost 5 macroblocks and the rightmost 6 macroblocks. There is no cropping of the original images.

e) Quadrature as shown in Figure 3-a: An MCU with this capability transmits four sub-pictures arranged as shown; no bordering or cropping is needed.

f) 9-sub-pictures as shown in Figure 3-b: An MCU with this capability transmits up to 9 sub-pictures arranged as shown; the area cropped by the MCU consists of the rightmost column of macroblocks (#22).

g) 16-sub-pictures as shown in Figure 3-c: An MCU with this capability transmits up to 16 sub-pictures arranged as shown; no bordering or cropping is needed.

h) Mixed-View as shown in Figure 4: An MCU with this capability transmits up to 6 sub-pictures arranged as shown; no bordering or cropping is needed. The area cropped by the MCU consists of the rightmost column of macroblocks (#22). Other variations on this image can be created by local processing.

6.2.2 Procedures for video mixing

When the MCU begins displaying a mixed picture, it sends to all terminals the double SBE Video Indicate Compose (VIC) <M> when M is the composition number used in Table 4. The MCU shall send VIC <M> each time it changes to a new standardized picture composition. If <M> is equal to zero, this indicates that a non-H.243 picture composition method is being used, and that the terminal may require additional information before processing the mixed picture. Values of M above 8 are reserved.

Table 4/H.243 – Sub-picture splits for various compositions

Name & associated figure number	**Vertical macroblocks split points in the mixed image (top to bottom)**	**Horizontal macroblock split points in the mixed image (left to right)**	**Note (MB = H.261/H.262 macroblock)**	**Mix number (used in VIC M)**
Dual-View Figure 2-a	4/5, 13/14	11/12	4 MB border above, 5 MB border below	1
Dual-View Figure 2-b	None	11/12	Cropping of original images required	2
Dual-View Figure 2-c	9/10	None	Cropping of original images required	3
Dual-View Figure 2-d	9/10	5/6, 16/17	5 MB border on left, 6 MB border on right	4
Quadrature Figure 3-a	9/10	11/12		5
9-sub-picture Figure 3-b	6/7, 12/13	7/8, 14/15, 21/22	1 MB border on right edge	6
16-sub-picture Figure 3-c	1/2, 5/6, 9/10, 13/14, 17/18	1/2, 6/7, 11/12, 16/17, 21/22	Drop top and bottommost macroblock, drop leftmost and rightmost macroblock	7
Mixed-View Figure 4	6/7, 12/13	7/8, 14/15, 21/22	1 MB border on right edge	8

When picture composition is in use, VIN is replaced by VIN2, which is followed by an <M><T> pair and a sub-picture number N as indicated in Figures 2 to 4. The number of VIN2s sent is equal to the number of sub-pictures displayed in the mixed image; VIN2s are not sent for a sub-picture until an actual video signal (as opposed to a pre-stored image) from a particular endpoint is displayed in that area of the mixed image. The terminal is responsible for optionally requesting the identity strings associated with the terminals and displaying those strings as an overlay on the mixed image. The terminal may use the knowledge of picture structure in Table 4 from VIC to further process the received image.

When a terminal is removed from the mixed image and not replaced by a new image, a {VIN2 <M> <T> <N>} is sent where M = T = 0, indicating that the sub-picture area is either blank or has been replaced with a message or image provided by the MCU. The MCU should avoid excessively rapid changes of sub-picture location or of the mixed image in general, but the proper strategy for doing so is left to the manufacturer.

If <N>, <M>, and <T> are equal to 0 in VIN2, this indicates that the MCU is overlaying terminal identification strings, and there is no need for the terminal to perform this function.

When the MCU ends picture composition operation, it shall indicate the return to video switching by sending VIN to all terminals.

The MCU should interpret Video Indicate Mixing (VIM) capability as an indication from terminals that they support VIN2 and VIC. While a mixed image is being presented to VIM-capable terminals, the MCU should follow the following rules in sending VIN and the mixed image to non-VIM-capable terminals:

a) When VIC is sent to VIM-capable terminals, VIN <M=0><T=0> should be sent to all non-VIM-capable terminals. Although the precise effect of this notification on 1990/1993 vintage terminals cannot be known with certainty, it represents an undefined terminal identifier, and may have the effect of clearing any terminal name or number currently being displayed by the terminal.

b) When mixed view is being transmitted (Mix number 8 in Table 2), the <M><T> pair should correspond to the largest image (5).

c) When any other mix is being transmitted, the MCU should send the <M><T> pair corresponding to the most recently added or loudest speaker at the discretion of the manufacturer.

If the MCU is overlaying terminal identification strings in the video mix itself, it should not send VIN for non-VIM-capable terminals after it sends VIN <0><0>. Also, note that the MCU may, at the manufacturer's discretion, choose to send only an unmixed image to non-VIM-capable terminals to avoid any VIN ambiguity.

6.2.3 Interactions with MCV

When video mixing is occurring and the MCU is in receipt of an MCV, the MCU should replace the mixed image with the image from the terminal sending MCV. When Cancel-MCV is sent, video mixing shall continue unless it has been de-activated by some external means or internal algorithm of the MCU. Since following the MCV is only a suggestion, some manufacturers may wish to provide MCUs with two modes, one where MCV has priority over mixing, and another where it does not.

6.2.4 Interactions with VCS

When video mixing is in use, the MCU should replace the mixed image with the image requested by VCS for the terminal sending VCS. When Cancel-VCS is sent, the mixed image should be transmitted to the terminal sending Cancel-VCS unless mixing has been de-activated by some external means or internal algorithm of the MCU. Since following the VCS is only a suggestion, some manufacturers may wish to provide MCUs with two modes, one where VCS has priority over mixing, and another where it does not.

6.2.5 Interactions with chair control

When the chair drops a terminal from the conference, the image transmitted from that terminal is dropped from the mixed image. The choice of image (another speaker, a blank screen, etc.) to replace the dropped image is left to the discretion of the manufacturer.

When the chair selects a new broadcaster, this ends video mixing and returns the MCU to video switching operation. Video mixing is restored when the chair sends Cancel-VCB, unless it has been de-activated by some external means or internal algorithm of the MCU.

6.2.5.1 Video composition and cascading

Since the link between two cascaded MCUs can carry only one video image, the interactions between video composition and cascading constitute a special case. A slave MCU shall pass to the master MCU the video image on it most recently selected. The master MCU shall compose an image that is sent to all conference participants on both the master and all slave MCUs. Note that only one video image from each slave may appear in the mixed image at the same time, although more than one video image from terminals connected to the master MCU may appear simultaneously. The possible improvement of this situation by the addition of high-bandwidth inter-MCU cascade links is for further study. The operation for non-master/slave procedures is for further study.

6.3 Automatic switching and visualization-forcing

All video switching actions in this subclause conform to the procedure of 6.1.

When voice-activated switching is used, the automatic switching of video signals is governed by the audio signals transmitted from the terminals, as described in 5.2.4/H.231. This voice-activated video switching is effective from the start of the video transmissions, unless and until overridden by VCB, MCV, VCS, T.120 conference control, the MCU or an out-of-band controller. If optional terminal numbers have been assigned (see clause 7), in each outgoing signal the MCU transmits periodically (with each BAS command cycle) the terminal number of the video it is transmitting, using the symbol {VIN, <M>, <T>}. All terminals having suitable capabilities may thus display an identity (number or name) with the video (see 7.4).

6.3.1 Video Command Broadcast (VCB)

See 9.4.1.

6.3.2 Multipoint Command Visualization (MCV)

By transmitting the symbol MCV (see ITU-T Rec. H.230), a terminal can try to force its MCU to broadcast its video signal to all other ports, overriding the voice-activation mechanism. When the said terminal no longer requires this broadcasting, the terminal shall transmit the symbol Cancel-MCV.

There are two MCV procedures. The first is for use when either end of a terminal-MCU or inter-MCU link lacks the Multipoint Visualization Capability (MVC). The second, and preferred, procedure is for use when both ends of a terminal-MCU or inter-MCU link possess the MVC capability. MVC is effective on a link-by-link basis; note however that since MCUs are entitled to restate their capability set, some implementations may choose to remove the MVC capability if other attached MCUs do not possess MVC, while other implementations may offer the MVC capability regardless, if they have appropriate means to determine when visualization is achieved.

In both procedures, no provision is made for selection of the video to be transmitted to the terminal that is the source of the distributed video. Its local MCU may forward the previous video signal or that from the T_M, if available, or other available signals on a rotating basis (e.g., 20 s at a time), or another basis at the discretion of the MCU manufacturer.

6.3.2.1 Procedure in the absence of Multipoint Visualization Capability (MVC)

On receipt of MCV from a directly-connected terminal or from another MCU, if an MCV is not already in force, the MCU switches the video from that port through to all other ports, including directly-connected terminals and inter-MCU links; it also forwards MCV to any other MCUs other than the MCU that has sent the MCV. Having done this, the MCU transmits MIV to the terminal being broadcast. When the said terminal no longer requires this broadcasting, the terminal transmits the symbol Cancel-MCV. The MCU reverts to voice activation and sends Cancel-MIV to the terminal when the terminal is no longer being seen by others in the conference; the MCU also forwards Cancel-MCV to any other MCUs. Other MCUs should not forward Cancel-MCV to the MCU from which it was received.

If an MCU receives MCV at one port while a visualization phase resulting from receipt of MCV at another port is effective, it shall not act on MCV, instead returning VCR. The action of MCV overrides any VCS commands the MCU may have received before receipt of MCV and until receipt of Cancel-MCV. Once an MCU has assigned its chair-control token by transmitting CIT and a broadcaster has been designated using VCB (see clause 9), or during a chair-control session using

ITU-T Rec. T.120, it shall not act on MCV; instead, the MCU shall respond with VCR. In the event of MCV "collisions", the first MCV received should have priority, and the sender of the second should be sent VCR. This allows, at the manufacturer's discretion, for giving special terminals, such as the chair terminal, priority over other terminals. In a non-master/slave mode, the contention resolution principle must be used to deal with all such collisions.

6.3.2.2 Procedure under mutually declared Multipoint Visualization Capability (MVC)

The following procedures are used when both ends of a terminal-MCU or inter-MCU link have declared Multipoint Visualization Capability (MVC).

This procedure differs from the one used without MVC being in force mostly in the responses to the MCV and Cancel-MCV commands. It provides MVA as a positive response to the terminal that visualization has been achieved, and is therefore the preferred method.

NOTE – MIV will also be received by a terminal (via the normal MIV procedure) which becomes broadcaster, since it is seen by one or more other terminals. However, the reception of MIV shall not be interpreted to mean that the terminal is being seen by all other terminals; only MVA gives this indication.

In these procedures, MCV and Cancel-MCV should be recirculated by the sender until a response is received. MVA or MVR shall be sent a single time in response to each received MCV or Cancel-MCV symbol, as described below.

The action of MCV overrides any VCS commands the MCU may have received before receipt of MCV and until receipt of Cancel-MCV. Once an MCU has assigned its chair-control token by transmitting CIT and a broadcaster has been designated using VCB (see clause 9), it shall act on MCV as follows:

1) If the MCV is received from the currently broadcasting terminal (i.e., the terminal identified in the VCB command), the MCU shall respond with MVA, as visualization is in fact achieved due to the VCB that is in force.
2) Otherwise, MCV shall be rejected by the MCU by responding with MVR.

During a chair-control session using ITU-T Rec. T.120, an MCU shall not act on MCV; instead, the MCU shall respond with MVR. In the event of MCV "collisions", the first MCV received should have priority, and the sender of the second should be sent MVR. This allows, at the manufacturer's discretion, for giving special terminals, such as the chair terminal, priority over other terminals. In a non-master/slave mode, the contention resolution principle must be used to deal with all such collisions.

6.3.2.2.1 MCV actions for master MCUs and MCUs in non-cascaded conferences

On receipt of MCV, if MCV is not already in force on behalf of another port, the MCU shall switch the video from the requesting port through to all other ports, including directly-connected terminals and inter-MCU links, and forward MCV to all MCUs except the one from which MCV was received (if any). The MCU shall then send MVA to the port from which MCV was received.

If MCV is already in force on behalf of another port when MCV is received, or if MCV is withdrawn due to a higher priority switch request (e.g., VCB), MVR shall be sent to the requesting port.

On receipt of Cancel-MCV, the MCU shall revert to voice activation, forward Cancel-MCV to all MCUs except the one from which Cancel-MCV was received, and send MVR to the port from which Cancel-MCV was received.

6.3.2.2.2 MCV actions for slave MCUs

NOTE – In the procedures below, slave MCUs not directly connected to the master MCU shall treat the MCU originating an MCV, Cancel-MCV, MVA, or MVR as a slave or a master MCU, as described in 5.7.2.3.

On receipt of MCV from a master MCU, a slave MCU shall switch the video from the master MCU port through to all other ports, and forward MCV to all MCUs except the one from which MCV was received.

On receipt of MCV from a slave MCU or directly connected terminal, if MCV on behalf of another port or a higher priority switching mode is not already in force, a slave MCU shall forward MCV to the master MCU.

If MCV is already in force on behalf of another port when MCV is received, or if MCV is withdrawn due to a higher priority switch request (e.g., VCB), MVR shall be sent to the requesting port.

On receipt of Cancel-MCV from a connected terminal or slave MCU, a slave MCU shall forward Cancel-MCV to the master MCU. On receipt of Cancel-MCV from the master MCU, a slave MCU shall revert to voice activation and forward Cancel-MCV to all directly connected MCUs except the master MCU.

On receipt of MVR from a master MCU, a slave MCU shall revert to voice activation, forward MVR to the port which has invoked MCV and send Cancel-MCV to all other directly connected MCUs. On receipt of MVA from the master, a slave MCU shall first switch the video from the requesting port which invoked MCV through to all other ports, then forward MVA to the port which has invoked MCV.

6.3.2.2.3 MCV actions for dumb-bell MCUs

On receipt of MCV or Cancel-MCV from a peer MCU, an MCU in a dumbbell configuration shall follow the actions prescribed for master MCUs.

On receipt of MCV or Cancel-MCV from a terminal, an MCU in a dumbbell configuration shall follow the actions prescribed for slave MCUs.

On receipt of MVA or MVR from a peer MCU, an MCU in a dumbbell configuration shall follow the actions prescribed for slave MCUs.

6.3.3 Video Command Select (VCS)

By transmitting the symbol {VCS, <M>, <T>} a suitably-equipped terminal can determine which video signal shall be transmitted to itself. If the local MCU has this optional capability, and if it also has the requested video signal available, it transmits the requested video to this terminal. In the case of contention with a VCB or VCS request from the chair-control terminal T_M, the chair-control request shall take precedence. If the MCU cannot comply, it shall return VCR.

To return to the automatic selection of video (6.3), the terminal transmits Cancel-VCS.

NOTE 1 – This procedure can only be used when numbering of terminals has taken place.

NOTE 2 – It is highly desirable that a terminal equipped to transmit VCS present to the user continuously an indication (reminder) while this facility is activated.

NOTE 3 – Transmission of VCS from a terminal may not secure the desired result, for various reasons: Due to the single link between MCUs, conflicting demands may not be met or, the MCU may not provide for many VCS commands from different terminals simultaneously. There may be other reasons.

NOTE 4 – VCS received by an MCU should not be acted on if such an action would disrupt the video received by any conference participants other than the sender of VCS. An MCU may optionally propagate VCS in a cascade if this rule is followed; such propagation is not required.

In summary, the MCU uses the following rules of precedence for visualization operation:

–	When chair-control token has been assigned:

a) if VCB is in force, refuse all conflicting VCS requests, and override all voice-switching (except that which may be in force to control the return video to the chair terminal); note that the chair terminal may wish to use VCS to survey response to a broadcaster by viewing each non-broadcaster site in turn; this use of VCS is not considered to be in conflict with VCB since the primary goal in the chair control scenario is to accede to the wishes of the chair terminal;

b) if VCB has not been received or Cancel-VCB is in force, accede to VCS from any local terminal requesting to see other local terminal's video.

–	When chair-control token has not been assigned:

c) if MCV is in force, refuse all conflicting VCS requests, and override all voice-switching;

d) if MCV has not been received or Cancel-MCV is in force, accede to VCS from any local terminal requesting to see other local terminal's video.

7 Numbering of terminals

All the provisions of this clause are optional, but it should be noted that they are required for most of the functions available under the chair-control provisions of clause 9.

The assignment of numbers to each terminal can serve the following purposes:

–	association of additional channels with the correct initial channel, when single-number conference service is offered (see 7.2.2 or ITU-T Recs H.242 and H.221);

–	management of chair-control functions (see clause 9).

The following terms are used in this clause:

1) NAN: Network Address Number (similar to telephone number) – to avoid confusion with numbers assigned within the MCU system;

2) per-MCU meet-me NAN: All terminals dial a single NAN to reach a conference on an MCU. This requires the terminals to identify the conference they wish to join once connected to the MCU. This may be done by means such as TCS-3 (see ITU-T Rec. H.230);

3) per-conference meet-me NAN: All terminals in a conference dial a single NAN to reach that conference. Terminals in other conferences dial a different NAN;

4) per-terminal meet-me NAN: Each terminal dials a different NAN. Particular NANs are associated with particular conferences at reservation time.

7.1 Numbering method

All terminals are given a unique number <M><T> in the range <1 to 191> <1 to 191> (192-223 reserved in both cases, values from 224-255 are not used to avoid leading "111" patterns), where the <M> is an 8-bit number allocated to the local MCU (see clause 3/H.231) and <T> is an 8-bit number allocated by the local MCU to the terminal. Both 8-bit numbers are encoded using one of the set of SBE symbols "NUM" (see ITU-T Rec. H.230); however, the pair <M><T> shall always be preceded by another symbol conveying the Control or Indication concerning the terminal of that number.

The value <M = 0> is not assigned except as described in 7.3.1.1. If only one MCU is involved in the call, the value of <M> may be set to any value, default <1>. If 2 or more MCUs are involved in the call, they may be given any unique value in the decimal range <1 to 191 (192-223 reserved)>

(The use of the reserved bits is for further study, as is the question of exhaustion of MCU numbers); MCU numbers may, for example, be assigned sequentially or reserved in advance.

The terminals attached to one MCU may be given any unique value in the decimal range <1 to 191 (192-223 reserved)> (The use of the reserved bits is for further study, as is the question of exhaustion of terminal numbers); they may for example be assigned sequentially or reserved in advance.

If two or more MCUs are connected in the call, it may be necessary to set up a master-slave relationship between them, at least for the generation of a unique set of terminal numbers. One of the MCUs may be designated master prior to the call, or by the in-band procedure of 5.7.2.2. Other MCUs may be connected directly to the master, which will treat them as slaves, or indirectly though other slaves.

7.2 Terminal-MCU interconnection

Two cases are considered: with and without call association. If "per-MCU or per conference meet-me NANs" are in use (see above) and multiple-channel calls are involved, call association is necessary.

In the following cases there is no need for the MCU to associate incoming calls into a single multiplex and the procedures of this subclause apply:

1) when only a single channel is needed for all multiplexes accessed through the same meet-me NAN, e.g., H0, 1B, etc.;
2) when per-terminal meet-me NANs are used;
3) when MCU dial-out operation is in use;
4) others.

7.2.1 Terminal/MCU interactions without call association

When a terminal is first added to the conference and initialization has been completed according to ITU-T Rec. H.242, if the MCU has TIC or CIC capability, it shall transmit to the terminal the symbol {TIA, <M>, <T>}, where <M> is the MCU number and <T> the number assigned by the MCU. If the MCU does not have TIC or CIC capability, this is optional.

NOTE – Terminals not equipped to receive such symbols ignore them, as only SBEs are used.

If the MCU is not, or not yet, connected to a master value of <M>, it is locally assigned (default <1>). If the MCU is subsequently connected to a master and receives a value of <M> from it, {TIA, <M>, <T>} is retransmitted as the new value.

If a terminal leaves the conference or is dropped for any reason, the corresponding value of <T> may be reassigned; a terminal re-added into the conference may be given the same number as before.

If the MCU sends a new {TIA, <M>, <T>} at a later time to a terminal previously assigned an <M><T>, this value replaces the previous value. Before the MCU sends a new <M><T>, the MCU should send MIJ to all terminals in the conference. Each terminal should respond by sending TCU to the MCU. Note TIA may be sent by the MCU without the MCU also sending MIJ. This might be the case if the terminal is being re-assigned a number previously assigned to it.

7.2.2 Terminal/MCU interactions with call association

In the case where the MCU is operated in a meet-me mode using per-conference or per-MCU meet-me NANs, initial and additional channels for each multiplex may be associated using the following in-band signalling procedure.

Terminals and MCUs with BAS-cap TIC are capable of these call association procedures. In the event that a terminal without TIC attempts to join a conference under these circumstances, it may be reduced to secondary status, suffer repeated call failures, or be delayed in joining the conference.

Whenever the MCU has accepted an initial-channel call and carried out the initial capability exchange, it looks for incoming BAS-cap TIC (see ITU-T Rec. H.230); if it finds this, the MCU shall send a TIA value on the I-channel as described in 7.3.1.1. This TIA value consists of <M> (MCU number) and <T> (terminal number), and uniquely identifies the terminal. When the additional-channel calls are made, the terminal sends in the additional channels:

– in FAS, the channel number according to ITU-T Rec. H.221;
– in BAS position, alternately the channel number according to ITU-T Rec. H.221 and the symbol {<TIX>, <M>, <T>}. The MCU is then able to associate the additional channels with the correct initial channels. Note that there is no capability exchange on additional channels and the above values are sent by a terminal upon connection, without waiting for $A_n = 0$ framing response from the MCU.

As calls (channels) arrive at the MCU, it shall begin to send H.221-framed signals towards the terminals. Within the H.221 frame, FAS channel numbering information shall be carried (see 2.2/H.221). Therefore, the MCU transmits FAS with the values L1 = L2 = L3 = 0 until it has made the correct associations, and then supplies FAS with the correct channel number. Terminals having TIC shall be able to accept this condition.

If the MCU sends a new {TIA, <M>, <T>} at a later time, this value replaces the previous value and the terminal sends the new <M>, <T> as part of all subsequent TIX. The MCU should allow a minimum of ten seconds for the terminal to begin reflecting back the new <M><T> before taking corrective action. As noted above, the MCU should send MIJ before the TIA to indicate that the renumbering has occurred and the terminal has joined the real conference.

If the MCU, while using per-conference or per-MCU meet-me NANs, finds a terminal with no BAS-cap TIC, it may do one of the following:

1) maintain a reduced transfer-rate cap, thus keeping the offending terminal in secondary status;
2) drop any additional channels; or
3) send the higher transfer rate-cap only to one terminal at a time, until that connection has been brought up to the desired rate, before proceeding to another terminal. This may lengthen conference set-up time.

7.2.3 Passwords and conference identifiers

The MCU may choose to require a password from terminals using TCS-1. See ITU-T Rec. H.230 for procedures.

The MCU may, in the case where a single network address is used by all terminals joining the conference, request using TCS-3 a conference identifier that is then used to connect incoming calls to the correct conference (see ITU-T Rec. H.230 for the definition of TCS-3).

For both passwords and conference names, the string should be limited to 32 characters. The MCU may request a password or conference identifier at any time after MCC is sent; this may be before or after an additional channel call has been answered by the MCU. Note also that an MCU may request passwords and conference identifiers more than once, as the user moves from one conference to another. As the terminal moves from one conference to another, it should expect to receive a possible MIJ for each conference, and to respond by sending a TCU.

The MCU may choose not to connect the terminal to either the audio mixer or the video mixer until the password and/or conference identifier have been collected (note that a loopback mode does not constitute connection to a mixer). Alternatively, the MCU may choose to connect the terminal immediately into the conference even though a password is needed. When the MCU has connected the terminal to the desired conference (as opposed to a conference where, for example a pre-recorded message is played), the MCU should send the BAS code MIJ (Multipoint Indicate Joined real conference) to the terminal joining the conference. Note that the MCU may ignore any commands from the terminal sent before the terminal has issued the proper password and conference identifier; terminals should re-send any such commands after receipt of MIJ.

7.3 MCU interconnection

7.3.1 Master MCU has been designated

The actions described in this clause are taken after transmission of MIM by the master and receipt thereof by the slave in question.

7.3.1.1 Assignment of MCU numbers

The master transmits the symbol {TIA, <M>, <0>}. The slave recognizes this as coming from the master, registers <M> as its own assigned MCU number, and then transmits to the master the list TIL. Note that the master MCU is always numbered "1".

MCUs with MIH capability operate in the following manner. Once the lower-level slave has received MIM and begun operation as a slave, it shall send TIN <0><0> to the slave it is connected to. This slave will forward the TIN up the hierarchy to the master. The master shall respond with TIA <M><0>, which is passed to the new MCU. M is now the MCU number of the new slave. The newly numbered MCU <M> sends its terminal list to the master with TIL. Slave MCUs are responsible for knowing the sub-tree of MCU numbers connected on each port so that they may properly forward messages. A slave MCU shall not forward an additional TIN <0><0> from a new MCU to the master until the TIN/TIA sequence outstanding has been completed.

When master/slave domains merge as described in 5.7.2.4, MCUs on the "losing" side must be renumbered to maintain unique MCU numbers. During the renumbering process, both masters shall refrain from mode changes and other operations that depend on master/slave relationships. However, the procedure just described is not sufficient since the lower layers of the "losing" domain are not aware of the contest; only the former master has this awareness. The "losing" master sends Cancel-MIM to all MCUs in its domain, and also to the "winning" domain where it is forwarded to the "winning" master. This indicates to each MCU in the losing domain that the "direction" to the master has changed, and that they should expect to be assigned new MCU numbers. Each MCU in the "losing" domain will no longer act on master/slave procedures until a new master is designated. Once the "winning" master receives Cancel-MIM from the "losing master", it sends MIM (in the direction that Cancel-MIM came) as a response. The intervening slaves are responsible for remembering the path Cancel-MIM took so that MIM can return on the same path. When the "edge" MCU of the losing domain receives MIM, it sends TIN <0><0> to the MCU it just received MIM from while holding the MIM until it receives a TIA <M><0> from the direction it sent TIN <0><0>. The "winning" edge MCU is responsible for passing the TIN <0><0> toward the master while not passing on any other TIN <0><0>s it may receive. As it is possible that the TIN <0><0> may be lost, the MCU sending it should re-send it after a time-out of at least 5 seconds; it may be desirable to increase this time-out for very large domains. Once the "losing" edge MCU has been assigned an MCU number, it sends MIM to any MCUs it is connected to (other than toward the master). It then waits to receive TIN <0><0> and passes the first one received toward the master, while ignoring the rest. When TIA <M><0> arrives, the MCU sends it to the MCU that originated the TIN <0><0>, and awaits the next TIN <0><0>. This procedure continues recursively until the entire losing domain has been renumbered. Note that the renumbering process is identical to that used when a new MCU joins a multi-level hierarchy and is not connected directly to the master.

7.3.1.2 Forwarding of numbers for terminals added or dropped

If a new terminal is connected subsequently to any MCU, the local MCU shall send {TIN, <M>, <T>} to all its ports. If a terminal is dropped, the local MCU shall send {TID, <M>, <T>} to all its ports. If an MCU receives TIN and/or TID values from another MCU, it shall forward these to all its ports. Thus, information about terminals added or dropped is rapidly disseminated to all parties.

7.3.1.3 Storage and dissemination of terminal numbers

All numbers of terminals added and dropped shall be stored in the master MCU and optionally at any other. The symbol TCU may be used at any time after MCC is received and any password/conference identification procedure is complete and by any terminal to elicit a list of the terminal numbers currently participating in the conference. TCU may be transmitted from any terminal or MCU to any MCU. If that MCU holds the complete list, it responds with TILs followed by TIE (H.230); if not, it shall forward TCU to the master and the master MCU shall respond accordingly. Note that the slave must remember the path of the TCU so that the TIL may be sent to the requestor, and that the terminal number of the requesting terminal will be included in one of the TILs. Other BAS commands may be interspersed between TILs, but the last TIL should be followed by TIE.

7.3.2 No master MCU has been designated

For further study.

7.4 Terminal identity information

The procedures in this clause are optional for terminals and MCUs.

Provision is made for the transmission of personal – or terminal-identity numbers, names, or other information between a terminal and its local MCU, if both entities are suitably-equipped. This procedure applies only to directly-connected terminals.

7.4.1 Terminal identity using Table 3/H.230 (ASCII) defined characters

One entity transmits the symbol TCI or TCS-2 (ITU-T Rec. H.230).

An entity which receives TCI responds with a sequence of symbols {TII, A-N} (see Note 2 below), where A-N represents one of a set of values defined for alphanumerics in ITU-T Rec. H.230. The sequence shall be terminated by the end-marker TIS.

An entity which receives values of TCS-2 responds with the MBE message IIS. Values for the IIS message can be found in ITU-T Rec. H.230.

NOTE 1 – The same identity string shall be used to reply to TCS-2 and TII. The length of this identity string should be limited to 32 characters.

MBE-cap is required for TCS-2 procedures; TCS-2 is ignored by a terminal not having this capability.

NOTE 2 – The TII sequence used to send the string "XYZ" is {<TII><X>}, {<TII><Y>}, {<TII><Z>}, {<TIS>}. Other BAS codes may be interspersed within this sequence but not between <TII> and the following symbol.

To associate a terminal number with the respective terminal's identification string as obtained via TCI or TCS-2, a requesting terminal transmits the symbol TCP (Terminal Command Personal-Identifier) followed by a terminal number encoded as an SBE NUM as described in ITU-T Rec. H.230 (e.g., {TCP, <M>, <T>}).

If the terminal whose personal identifier is requested is attached to the local MCU, the MCU responds with the MBE symbol TIP containing a terminal number followed by the identification string encoded as described in ITU-T Rec. H.230.

If the MCU does not support personal identifiers, it ignores the TCP request. If the terminal number is invalid, or if the terminal does not support personal identifiers and hence has not provided an identification string, the MCU responds with the MBE symbol TIP containing a null ASCII identity string. If the MCU supports personal identifiers but has not requested the personal identifier string for the specified terminal, the MCU requests the personal identifier by using the TCI or TCS-2 symbol.

If the terminal whose personal identifier is requested is not attached to the local MCU, or if the local MCU does not keep lists of personal identifiers, the TCP request is forwarded to the master MCU. If the master MCU is not connected to the appropriate terminal or does not know the appropriate personal identification string, the master MCU forwards the request to the MCU identified by the <M> portion of the terminal number. The TIP is returned by the same path used by the TCP request (note that the response may be the null identity string if the terminal number was invalid). If the destination MCU identified by <M> is not valid, the master MCU returns the TIP with the requested terminal number and a null ASCII identity string.

7.4.2 Terminal identity using ISO/IEC 10646 (Unicode) defined characters

The ISO/IEC 10646 (Unicode) set offers a greatly extended set of characters.

7.4.2.1 Encoding of Unicode characters within MBEs

For the purposes of this Recommendation, each Unicode character shall be encoded within MBE as follows:

1) The character from the four-octet canonical form of the ISO/IEC 10646 character set (UCS-4), consisting of the Group-octet (G), Plane-octet (P), Row-octet (R), and Cell-octet (C), shall be converted to an integer by the following formula:

$$\text{integer} = (G * 16\,777\,216) + (P * 65\,536) + (R * 256) + C$$

2) The resulting integer, which represents the character, shall be carried in the MBE according to the non-negative integer procedure of Annex A/H.239.

NOTE 1 – This procedure both avoids emulation of the MBE BAS code, and has the effect of encoding any ASCII character in a single octet, and any Basic Multilingual Plane (BMP) character in three or less octets.

NOTE 2 – H.320-H.245 gateways may translate BMP UCS-4 characters to BMP UCS-2 characters by dropping the Group (G) and Plane (P) octets, which both have a value of zero for BMP characters, and may translate UCS-2 characters to UCS-4 by adding Group (G) and Plane (P) octets of value zero.

7.4.2.2 Procedure

One entity transmits the symbol TCS-5 (ITU-T Rec. H.230).

An entity which receives values of TCS-5 responds with an MBE message IIS with a value of n=5.

NOTE – The length of the identity string in IIS-5 (value of n=5) and TIP-5 should be limited to 32 octets.

MBE-cap is required for TCS procedures; a terminal not having this capability or not recognizing TCS-5 ignores TCS-5.

To associate a terminal number with the respective terminal's identification string as obtained via TCS-5, a requesting terminal transmits the symbol TCP-5 (Terminal Command Unicode Personal-Identifier) followed by a terminal number encoded as an SBE NUM as described in ITU-T Rec. H.230 (e.g., {TCP-5, <M>, <T>}).

If the terminal whose personal identifier is requested is attached to the local MCU, the MCU responds with the MBE symbol TIP-5 containing a terminal number followed by the identification string encoded as described in ITU-T Rec. H.230 and 7.4.2.1.

If the MCU does not support Unicode personal identifiers, it ignores TCP-5 requests. If the terminal number is invalid, or if the terminal does not support Unicode personal identifiers and hence has not provided a Unicode identification string, the MCU responds with the MBE symbol TIP-5 containing a null identity string. If the MCU supports Unicode personal identifiers but has not requested the Unicode personal identifier string for the specified terminal, the MCU shall request the Unicode personal identifier by using the TCS-5 symbol.

If the terminal whose personal identifier is requested is not attached to the local MCU, or if the local MCU does not keep lists of personal identifiers, the TCP-5 request is forwarded to the master MCU. If the master MCU is not connected to the appropriate terminal or does not know the appropriate personal identification string, the master MCU forwards the request to the MCU identified by the <M> portion of the terminal number. The TIP-5 is returned by the same path used by the TCP-5 request (note that the response may be the null identity string if the terminal number was invalid). If the destination MCU identified by <M> is not valid, the master MCU returns the TIP-5 with the requested terminal number and a null identity string.

7.4.2.3 Use of Language ID field

The languageID field in the messages IIS-5 and TIP-5 is a hint to receiving terminals as to which form (ASCII or Unicode) of the identity or personal identifier string should be displayed to the human end-user. LanguageID consists of 2 octets representing a two-letter code as defined in ISO 639-1 encoded using the letter values given in Table 3/H.230.

If the language represented by the received languageID field matches the language of the terminal's user interface, or a language otherwise known as likely to be understood by the human end-user (for example, based on the location the terminal is installed), the terminal should display the Unicode version of the string. Otherwise, the terminal should display the ASCII version of the string.

For example, if a user located in Japan calls a terminal located in Switzerland, the Swiss user may not understand Japanese characters. The languageID field would be set to "Japanese". Since this languageID does not match the Swiss user's language, the ASCII version of the string would be displayed.

7.4.3 Backward compatibility requirements

The following are required in order to ensure that all terminals are able to receive some form of terminal identity.

All devices which respond with IIS-5 shall also be capable of responding with IIS-2. Implementations which support IIS-5 should request users to supply personal identifier information suitable for transport in both IIS-2 and IIS-5 formats.

All devices which transmit TCS-5 shall also transmit TCS-2, since some older terminals support only TCS-2 and IIS-2.

All terminals which transmit TCP-5 should also transmit TCP, since some older terminals support only TIP.

8 Mode switching and data broadcast procedures

8.1 General mode switching

The provisions of this clause are mandatory for all MCUs.

8.1.1 Bit-rate symmetry

In a point-to-point call, a terminal is free to change modes at any time, within the constraint of the capabilities it has received from the other end. However, in a multipoint call, there are additional, temporal constraints:

a) Because the output frames from the MCU cannot be synchronous with all input frames, there will usually be at least a partial-submultiframe delay in transmitting a necessary BAS code; in a more extreme case the MCU may already be engaged in a capability exchange with another terminal, and so be unable itself to mode switch for some time.

b) Time is needed for the MCU to process BAS capabilities and commands to ensure that the resulting modes are acceptable to all primary terminals (see ITU-T Rec. H.231) and are imposed in coordination, without corruption of any video being transmitted.

To ensure that an MCU has adequate control, and in particular that it can drive video signal transmission to a common rate (noting that in the case treated here the MCU has no power to transcode the video), bit-rate changes are initiated solely from the MCU. Terminals, after having received MCC and MCS from the MCU, shall not change bit rates except in response to such a change incoming from the MCU, so as to maintain the symmetry commanded by MCC and, if relevant, MCS for each component signal. This applies to bit rates for audio, data (LSD, HSD, MLP, H-MLP), video, Encryption Control Signal (ECS) channel and the transfer rate; audio and video mode changes not involving bit-rate changes may still be initiated by terminals. When the bit rate incoming from the MCU changes, the terminal shall follow suit as promptly as other procedures allow, as any delay may preclude the terminal's transmission from being received by the other parties in a conference.

8.1.2 Video mode switching without bit-rate switching

When in receipt of MCC but not MMS, terminals may initiate video mode switches that do not change bit rates. Thus, a terminal may switch from sending H.261 to H.263 at any time unless in receipt of MMS. This switch shall involve sending the sequence Video-Off followed by H.263-ON.

Since this kind of unconstrained video mode switching may be disruptive, the MCU can prevent it by issuing MMS. Once in receipt of MMS, the following procedures for video mode are used; they are similar to those for bit-rate changes. The MCU shall perform the following steps:

1) Send Video-Off to all parties in the conference.

2) Send H.26x-ON to all parties in the conference.

It is the responsibility of each receiving terminal to display a frozen image or blank screen during this switch.

8.1.3 Restricted networks and bit-rate symmetry

If some calls to the MCU are on restricted networks, and others on unrestricted networks, the following situations can arise:

1) All the parties currently in the conference may be operating in an unrestricted fashion. When a call arrives from a terminal on a restricted network the MCU may choose to:
   a) reject the call;
   b) accept the call: Any terminal that has signalled compatible restricted capabilities shall be forced into the restricted mode by use of the restrict command. Any terminal that has not signalled compatible restricted capabilities shall be put into secondary status;
   c) accept the call: All terminals except those that have sent the No_Restrict cap, shall be forced into the restricted mode by the restrict command. Any terminal that cannot successfully establish a restricted connection shall be put into secondary status;
   d) accept the call and relegate the terminal to secondary status.

2) All the parties currently in the conference may be operating in a restricted fashion due to the inclusion of the Restrict_Required cap in the SCM for the conference. When a call arrives from a terminal on an unrestricted network the MCU may choose to:
   a) reject the call;

b) accept the call, and use the Restrict_Required cap and the restrict command to inform the terminal that it must operate in the restricted mode. The terminal shall respond with the restrict BAS command;

c) accept the call and if the No_Restrict cap is declared by the terminal, relegate the terminal to secondary status.

3) All the parties currently in the conference may be operating in a restricted fashion (situation 1 b). When all of the terminals using restricted networks are dropped from the conference, the MCU may choose to:

a) use the Derestrict BAS command to inform the parties that remain on the conference to switch back to unrestricted operation. Each terminal shall respond with the Derestrict BAS command;

b) do nothing.

In the multipoint case where an MCU has sent the MCC code to a terminal or slave MCU, the terminal or slave MCU shall not initiate the restrict mode on its own unless it is connected via restricted networks and has Restrict_required in its capset. When connected to an MCU, it is not allowable to operate in asymmetrical restricted modes. A terminal or slave MCU shall respond by sending a Restrict command as soon as possible.

The three BAS indication symbols below shall be used in cascaded operations between master and slave MCUs that have Restrict_P and/or Restrict_L, or No_Restrict in their capability sets. If only Restrict_Required is present, the MCU is an older unit that cannot be expected to support RIR, RIU, or RID. The three symbols are as follows:

– **Restrict_Request** (RIR): This indication is sent by a slave MCU to a master MCU to request restricted operation.

– **Restrict_Indicate_Unrestricted** (RIU): This indication is sent by a slave MCU to a master MCU to request unrestricted operation.

– **Restrict_Denied** (RID): This indication is sent by a master MCU to a slave MCU to deny an earlier request for restricted operation from the slave MCU.

A slave MCU operating in the unrestricted mode shall send Restrict_Request (RIR) if a terminal joined which signalled Restrict_Required. The master MCU shall respond either by sending the Restrict command to all of its terminals and slave MCUs, or by sending Restrict_Denied (RID) to the requesting slave MCU. In the former case, all of the slave MCUs shall send the Restrict command to all of their directly connected terminals. In the latter case, the requesting slave MCU should have put the restricted terminal into secondary status while awaiting a response from the master.

If all Restrict_Required sites left from the slave MCU, it should signal this to the master MCU by sending Restrict_Indicate_Unrestricted (RIU). The master MCU should respond by either sending a Derestrict command, or if there were other Restrict_Required nodes still attached, it would not respond. The master MCU should use this indication to keep track of the fact that this particular slave MCU no longer needed restricted operation. Note that the master MCU is not required to return to unrestricted operation.

If the master MCU is operating restricted, and it decides to move to the unrestricted mode, the secondary MCU shall relegate any terminals that are not capable of unrestricted operation to secondary status (audio only).

8.1.4 Changing the video bit rate

As a consequence of changing the bit rate of other signals, the video rate will also change, since it occupies all bits not designated for other signals. The procedure to be used is similar to that for video switching:

a) The MCU transmits VCF and Video-Off to all terminals before transmitting the BAS commands that establish a new rate for video.

b) Until a terminal that is a video source has responded by symmetrizing transmission, its outgoing video will be at the wrong rate and can no longer be forwarded to other terminals which may previously have been receiving it; if the terminal does not adjust its rate promptly, the MCU may switch to distribution of another video source pending restoration of the proper symmetrical condition.

c) When the new bit rates have been established for the other signals, video is again switched ON, forwarding from the same sources as before unless b applies or there is another overriding change.

d) After a time sufficient for video receivers to recover error-correction-frame alignment, the MCU transmits VCU to all video source(s).

8.1.5 Mode changes in multi-MCU calls

8.1.5.1 Master/slave working

On master-slave interconnections, bit rate mode changes shall only be initiated by the master; the slave shall respect the MCC and MCS commands in the same way as a terminal would do. In the case of multi-level hierarchies of MCUs, the lower levels shall treat the upper levels as masters, and the lower levels as slaves. It should be noted that in large hierarchies with many layers, mode changes will take longer periods of time, and the master should increase the time allowed for a mode change accordingly.

8.1.5.2 No master MCU has been designated

In this case, all MCUs may reject the literal interpretation of MCC and MCS if relevant, and seek only to coordinate mode changes with their peers, by "dynamic compliance": on each interconnecting link, each MCU adopts a mode change initiated by the other unless it is in the process of carrying out a contrary change. In the case of contention, the contention resolution principle (see 13.2) should be applied.

8.2 Mode switching for data distribution in multipoint conferences

The provisions of 8.2.1, 8.2.2 and 8.2.3 are mandatory for those MCUs that support data distribution using HSD/LSD. Procedures for MLP related mode switching are described in 8.2.4.

8.2.1 General HSD/LSD provisions

8.2.1.1 Range of data channel provisions

In this clause, the term "data" is used generically to refer to either of two types of data channel permitted by ITU-T Rec. H.221 and designated there as LSD and HSD. These are managed independently, and may be effective simultaneously: LSD may be sent by one terminal while HSD is sent by the same or a different terminal.

LSD and/or HSD may be sent by one terminal to its MCU, whence it is broadcast to all other terminals and MCUs in the call. Selectivity of destinations and multiple simultaneous transmissions of either LSD or HSD are for further study.

The following procedure shall be followed when the MCU has declared a suitable data capability: this can only occur if the MCU includes the requisite data distribution unit(s), if the service provider has previously agreed to its use, and if at least two terminals have declared the same capability.

Having received MCS from an MCU, a terminal shall not open a data channel of its own volition, but may make a request to the local MCU and await the outcome, as described below.

8.2.1.2 Idle bits

After a data channel is opened and before the data token is assigned, idle bits are broadcast by the MCU. For a period of time after the data token is assigned, the contents of the data channel may be undefined (i.e., consists of whatever a terminal stuffs in the channel before sending data). The only transmitter of data of each of the two types is the one terminal which has been assigned the requisite data token. An idle bit is a binary one: this is both a stop bit for asynchronous serial transmission and a permitted inter-frame time fill for HDLC-based protocols.

8.2.1.3 Terminals without data capability

Some of the terminals connected may not have the data capabilities to be used (but see Note 2 below) and, therefore, no data channel will be opened to them; the audio may not be affected. The following options are available to deal with this situation:

a) If they are not transmitting or receiving video, there is no change of service to the user.

b) If they are transmitting a video signal, this will no longer be at the same rate as those terminals to which the data channel has been opened; therefore, their video cannot be forwarded to those terminals, neither can they receive video from them, during data transmission; they may, however, continue to exchange video with other non-data-capable terminals if the MCU provides for this.

c) The MCU may choose not to open any data channel.

NOTE 1 – Since only SBE codes are used in this procedure, such terminals may ignore these symbols without misoperation.

NOTE 2 – It is desirable for a terminal to declare an LSD or HSD capability upon receipt of MCC/MCS and have such a channel opened to it, even though it may have no actual data equipment attached, provided only that its video rate conforms to the MCC and MCS condition. It is advisable to include such capabilities in terminals which may be used in multipoint calls involving LSD or HSD. Simple terminals may declare the nil-data cap to indicate that no use will be made of the declared LSD/HSD rates. More advanced terminals may use the procedures of 12.5/H.242 to indicate during the capability exchange what LSD/HSD rates are real and which are not.

8.2.2 Data tokens for HSD/LSD operations

The control of data distribution is governed by means of data tokens, one for each type of data; all tokens are assigned independently. Tokens may optionally be reserved in advance, or otherwise assigned by an out-of-band method. Tokens for LSD and HSD may be assigned to two different terminals.

Possession of a data token confers the right to transmit data for distribution to all other terminals having sufficient data capability; however, the token can be released by one terminal and taken up by another without the MCU closing the data channel or changing its rate. This subclause pertains to the management of LSD. Exactly the same process applies to the management of HSD, using the codes DCA-H, etc. (see ITU-T Rec. H.230). The data paths may be managed independently, and more than one may be effective simultaneously. An MCU may restrict the data transmissions to one type if appropriate, by withholding/withdrawing the other token and declaring a new capability set with omission of that type.

8.2.2.1 Assignment of the token

8.2.2.1.1 A terminal T_D wishing to transmit LSD may proceed to claim the requisite token if its currently-registered capset from the MCU includes the appropriate LSD value.

T_D requests assignment of the LSD broadcast token by sending {DCA-L, <B>}, where <B> represents the desired data rate according to the values in Tables 3 and 4. If it does not receive a response (see below) from the MCU, it should repeat the request in a reasonable time.

8.2.2.1.2 On receipt of {DCA-L, <B>} from terminal T_D, the local MCU acts as follows:

a) IF {OR (it has already assigned the token to a terminal or MCU other than T_D, (having transmitted DIT-L and not received DIS-L)), OR (received another request to do so from a directly-connected terminal or MCU), OR (is in the process of closing a data channel, or making a conflicting mode change), OR (if the requested data rate is not in the current common capset), OR (if the MCU is in a state of resource exhaustion)} THEN the MCU shall respond with DCR-L.

b) If the MCU has previously assigned the token to T_D, two cases exist:

   i) If the data channel is already open at the requested rate, the MCU shall respond with DIT-L, and T_D retains the token.

   ii) If this is as a result of T_D requesting a different rate, the MCU responds with DCR-L and T_D no longer holds the token. T_D should send another DCA-L to request the token at the new rate. (The preferred method of requesting a new data rate is for T_D to release the token by sending DIS-L and then ask for the new rate.) Since the MCU now holds the token, it responds as in c, i and ii below.

c) IF {it has neither assigned the token nor received another request to do so, nor have any of the other refusal conditions of a occurred} THEN:

   i) If it is the only MCU, it proceeds to perform any needed mode changes according to the procedures of 6.1. After the channel has been assigned and any appropriate mode change has taken place, the MCU sends DIT-L to T_D. At this point the terminal may begin to transmit data.

   ii) If it is one of two or more interconnected MCUs, three cases must be considered, keeping in mind that only a single LSD token exists in a master/slave network, and that the master controls that token:

      – A master has been assigned, and the local MCU is a slave. The slave MCU forwards {DCA-L, <B>} to its master and awaits DIT-L. When the slave MCU receives DIT-L or DCR-L, it forwards the code to T_D.

      – A master has been assigned, and the local MCU is the master. The master MCU acts on {DCA-L, <B>} by treating its subordinate MCUs like terminals.

      – If no master has been assigned, the operation is left for further study.

8.2.2.2 Release and reassignment of data token

A change of data control should be negotiated between conference participants; the terminal holding the token, and having ceased to transmit data, may release the token by sending DIS-L or DCC-L. This allows the terminal to request either that the channel be left open for future use (DIS), or that it be closed (DCC) to maximize video bandwidth. Several cases exist:

a) If the MCU is singular, it sends DCR-L to T_D and in the case of DCC-L closes the channel.

b) If the MCU is a slave, it forwards DIS-L or DCC-L to the master MCU, and awaits DCR-L. On receipt of DCR-L from the master, the slave forwards DCR-L to T_D and in the case of DCC-L closes the channel after the close has been initiated by the master MCU.

c) If the MCU is a master, it acts on DIS-L or DCC-L itself while treating its slaves as terminals.

d) If no master has been selected, operation is left for further study.

After the receipt of DCR-L or DCC-L, T_D is free to request the token again, perhaps at a different data rate.

An MCU receiving DIS-L or DCC-L from any directly connected terminal other than the one to which it previously assigned the token shall respond with DCR-L. In this case, if the data channel is open, it shall not be closed in response to DCC-L.

After receiving DIS-L or DCC-L from T_D, the MCU shall revert to transmitting idle bits if the channel is left open. Terminals receiving data may experience a brief period during which the state of data in the channel is undefined (between when DIS-L is received at the MCU and when the MCU begins sending idle bits). After sending DIS-L or DCC-L, a terminal shall not resume sending data without again claiming and receiving DIT-L.

8.2.2.3 Withdrawal of data token

The withdrawal of the data token may be invoked using the chair-control facility (see 9.6). In case of necessity (e.g., to resolve a fault condition), any MCU may itself withdraw the data token. In both cases, the withdrawal should be understood as the correction of an error condition, not a request to the token holder. Generally, terminals should hold the token only as long as needed to transmit their data.

An MCU making the withdrawal transmits DCR-L on the path on which it had sent DIT-L during the assignment of the token. An MCU receiving DCR-L shall forward this on the path on which it had sent DIT-L during the assignment of the token, and the MCU itself shall send DIS-L or DCC-L on the path from which it received DCR-L. This description applies to both master-slave and MCU-terminal paths.

A terminal receiving DCR-L while possessing the token shall cease data transmission within the LSD path consistent with proper operation of protocols within the data; it shall then transmit DIS-L or DCC-L to the MCU. If the MCU receives DCC-L, it shall close the data channel or allocate the token to a new broadcaster as appropriate. If the MCU receives DIS-L, it shall leave the channel open for future use.

If T_D does not return DIS-L or DCC-L within a reasonable time, the MCU to which it is connected may force the token release according to the above procedures. It is understood that this procedure may cause some loss of data. In general, the master time-out should be greater than the slave time-out.

8.2.3 Opening/closing/rate changing the LSD/HSD data channel

Opening a channel refers to moving from a mode where no data channel exists to a mode where one does exist. Closing the channel refers to the reverse operation. Rate change refers to moving from one rate to another in an already open channel (e.g., from LSD-300 to LSD-9600 or from LSD-2400 to var-LSD). All of these operations are mode changes, and are done according to the procedures of 8.1.

8.2.3.1 Master/slave or single MCU

The data channel shall only be opened, closed, or its rate changed when the master MCU or single MCU possesses the relevant data token. If any T_D is broadcasting when this necessity arises, the appropriate token shall either be released voluntarily by T_D when it completes its broadcast or forcibly withdrawn from T_D.

The MCU with the token changes the modes of all its ports in the same conference according to the procedures of 8.1.

Receiving the mode change from the MCU while the MCC and MCS are in effect, each terminal shall respond by symmetrizing its transmission; that is, the terminal shall open an identical data channel in the direction of the MCU, using the mode switching procedure (8.1) with the requisite BAS command(s). The terminal should be prepared to receive data from the time of the mode change; there is no warning of a transition from idle to undefined data to actual data, except perhaps asynchronously by a preceding data application BAS command. Use of the data channel shall take this into account, recognizing that different terminals may send different fill when they are not sending data.

The MCU awaits symmetrization of the changed channels; if any terminal is tardy in symmetrizing or does not support the data rate, the MCU should downgrade this terminal to secondary status.

In the master/slave arrangement, mode changes originate with the master and rate symmetrization radiates outward from the master.

When all primary terminals have established symmetry with the MCU, the MCU sends DIT-L to T_D, and begins broadcasting the data from T_D to all other connections.

T_D may now begin data transmission. If after receipt of DIT-L T_D sends a data application BAS command (Table A.3/H.221), the MCU will forward this on to all other ports, and after doing so will echo the same back to T_D. The forwarded commands cannot be synchronous to the broadcast data stream, nor can the echoed command be taken as a guarantee that all terminals have received it; the commencement of real-data broadcast should take these limitations into account.

After a period of disuse, during which no data token has been claimed, an MCU may close the data channel. Until a channel is closed, the MCU shall transmit idle bits when it holds the token.

8.2.3.2 No master

For further study.

8.2.4 Opening/closing/rate changing the MLP data channel

Opening a channel refers to moving from a mode where no data channel exists to a mode where one does exist. Closing the channel refers to the reverse operation. Rate change refers to moving from one rate to another in an already open channel (e.g., from MLP-4k to MLP-6400 or from MLP-6400 to var-MLP). All of these operations are mode changes, and are done according to the procedures of 8.1.

8.2.4.1 Master/slave or single MCU

The MLP channel shall only be opened, closed, or its rate changed by the master MCU or a single MCU. This MCU may initially open the MLP channel based on any of a number of rules, including:

a) a pre-set MLP rate for the conference;

b) an MLP rate (e.g., 6400 bit/s) that is common among the terminals in the conference;

c) an MLP rate that has been requested by one or more terminals according to the capability exchange procedures of 12.5/H.242, and is supported by other terminals in the conference;

d) maximize conference participation from non-MLP-equipped terminals by using the MLP rates of 32 K or 40 K, which allow the video rates to be matched with non-MLP terminals given that the MLP terminals support ITU-T Rec. G.728;

e) a request from a terminal using the DCM command from ITU-T Rec. H.230;

f) a mode preference request from one or more terminals using the method of clause 9/H.242.

The choice of method for deciding at what MLP rate to open the MLP channel is left to the manufacturer, but note that 6400 bit/s is mandatory on all terminals and MCUs supporting MLP. The MCU initiating the MLP mode operation changes the modes of all its ports in the same conference according to the procedures of 8.1. It should be noted that simple terminals may declare nil data to indicate that the declared MLP rates can be opened, but that ITU-T Rec. T.120 is not supported. Also, the capability exchange procedures of 12.5/H.242 may be used to indicate that ITU-T Rec. T.120 is supported on one or more MLP rates, while ITU-T Rec. T.120 is not supported on other MLP rates.

Receiving the mode change from the MCU while the MCC and MCS are in effect, each terminal shall respond by symmetrizing its transmission; that is, the terminal shall open an identical data channel in the direction of the MCU, using the mode switching procedure (8.1) with the requisite BAS command(s). The MCU awaits symmetrization of the changed channels; if any terminal is tardy in symmetrizing or does not support the common data rate, the MCU should downgrade this terminal to secondary status. In the master/slave arrangement, mode changes originate with the master and rate symmetrization radiates outward from the master.

When the MCU decides it is appropriate, the use of the MLP channel by T.120 data will be signalled by the T.120-on BAS command. When T.120-on is received, the MCU should connect the T.120 data streams to the T.120 data processor. The MCU shall not connect terminals with nil-data capability declared to the T.120 data processor. All terminals with T.120 capability or H.224_MLP capability (see A.14/H.221 and A.15/H.221) shall be connected to the T.120 data processor assuming they have established rate symmetry. Procedures for turning on ITU-T Rec. H.224 in the MLP channel are analogous to those for turning on ITU-T Rec. T.120. Note that ITU-T Recs T.120 and H.224 can co-exist in the MLP channel at the same time or operate independently. It should be noted that the MLP data application BAS commands such as ITU-T Recs T.120 and H.224 do not represent multiplex changes and, as such, are not required to take effect 20 ms after receipt as is the case for BAS commands such as an MLP rate command or ITU-T Rec. G.711.

There may exist terminals that are T.120-capable but do not implement the T.120 capability or make use of T.120-on/off; these terminals may be connected to the T.120 data processor at the discretion of the manufacturer.

Each terminal may now attempt to establish a connection to the T.120 data processor according to the procedures of ITU-T Recs T.122/T.123/T.125.

The MLP data channel may be closed for a variety of reasons, including:

– a pre-set or administered change;
– a request from a terminal using MLP facilities.

Once the MLP data channel has been closed, ITU-T Rec. T.120 cannot be used to request that it be opened again. DCM (see ITU-T Rec. H.230) may be sent by a terminal that desires an MLP channel to be opened by the MCU. The MCU is not obligated to open the channel in response, nor is it obligated to follow any particular rule in deciding at what rate to open the MLP channel.

A terminal may request a change of MLP rate by means of a mode preference indicator as described in 9.5/H.242. The syntax for this command is <MLP rate><M><T><Request Modifier> where <MLP> is a mode preference request, <M><T> is the terminal number of the terminal requesting the change, and <Request Modifier> is an SBE number, the meaning of which is defined in Table 5. If an MCU is requesting the rate change, the value of <T> shall be zero. The inclusion of the terminal number allows the MCU to differentiate between repeated requests from one terminal and requests from different terminals. In a master/slave cascade, the entire sequence, including the <M><T> pair, is forwarded to the master for processing. Note that although the MLP preference request is expected mainly to be used between terminals and MCUs, use in the point-to-point case is allowed as well.

Table 5/H.243 – MLP request modifier values

Bit position (as per H.221)	Value	Request for
0	x	Reserved (avoid leading 111)
1	0	audio off
	1	audio on
2	0	video off
	1	video on
4-8	0	6.4 kbit/s MLP (Note)
	1	8 kbit/s MLP
	2	14.4 kbit/s MLP (Note)
	3	at least 20 kbit/s MLP in initial channel
	4	32 kbit/s MLP (Note)
	5	40 kbit/s MLP (Note)
	6	at least 45 kbit/s MLP in initial channel
	7	var-MLP
	8	14.4 kbit/s H-MLP
	9	62.4 or 64 kbit/s H-MLP
	10	128 kbit/s H-MLP
	11	192 kbit/s H-MLP
	12	256 kbit/s H-MLP
	13	320 kbit/s H-MLP
	14	384 kbit/s H-MLP
	15	at least 100 kbit/s aggregated in both initial channel and additional channels
	16	at least 150 kbit/s aggregated in both initial channel and additional channels
	17	at least 200 kbit/s aggregated in both initial channel and additional channels
	18-31	Reserved
NOTE – Notes a preferred MLP rate.		

8.2.4.2 No master

For further study.

8.2.5 Data tokens for H.224 in the MLP channel (new)

As described in 6.2/H.224, both H.224 and T.120 packets may be sent in the MLP channel if the receiving terminal supports H.224-sim capability. When the receiving terminal has the H.224-token capability, H.224 and T.120 packets may be sent in the MLP channel using the following procedures:

1) The sending terminal shall request token access to the MLP channel using the appropriate value of <B> from Table 2 and the {DCA-L, <B>} command. Other than this, procedures are the same as if access was being granted to an LSD channel. Note that the token control only applies to the H.224 data, and not to the T.120 data.
2) After the MCU has granted the token and opened the MLP channel (if necessary), the MCU shall send H.224-token-on to indicate the start of H.224-token operations in the MLP channel.

3) When in this mode, the MCU may choose to either broadcast all H.224 datagrams, or to selectively broadcast datagrams based on H.224 header addresses to reduce MLP channel congestion.

4) When the token is given up, the MCU shall send H.224-token-off to indicate the end of H.224-token operations to all parties in the conference.

The MCU controls the SCM of the conference, and may choose to manage H.224 operations in the modes allowed by H.224-sim, H.224-MLP, or H.224-token. LSD operations may be optionally suppressed by the MCU when H.224-token-based operations are in use; thus H.224-token does not imply support by the MCU for simultaneous MLP and LSD operations.

9 Chair-control procedure using BAS codes

The provisions of this clause, except those of 9.4.2, are mandatory if the MCU supports chair-control.

9.1 General

See clause 15 for a description of the interactions between T.120 control and chair-control.

This option requires the MCU to have certain software and hardware provisions, and at least one terminal shall be suitably enhanced, as detailed below.

The MCU possessing CIC can:

– assign a number to each terminal;
– assign a chair-control token;
– disconnect a terminal from the conference on command from the token-holder;
– switch video signals according to commands from the token-holder;
– stop data transmission by all other terminals;
– drop the entire conference.

If two or more MCUs are to be involved in the conference, all shall declare CIC to support cascaded chair-control. Note that it is possible for an MCU to have CIC, and not support cascaded chair-control, since cascading is an optional feature separate from chair-control.

The terminal to be used for chair-control shall have means to:

– send the BAS values CCA, CIS, CCD, CCK, VCB, Cancel-VCB, and SBE numbers;
– display terminal numbers or other identifiers with associated video (or audio);
– accept user input regarding video switching and terminal disconnection, etc.;
– accept CIC, CIR, and TIF from the MCU.

It is not essential that other connected terminals have any special capabilities. Since only SBE codes are used in this procedure, such terminals may ignore these symbols. Note that the chair terminal itself need not declare CIC in its capability set.

The facilities provided by CIC may be presented to a single user, or the capabilities may be split at the terminal level to provide for two persons to act respectively as controller and chairman as specified in ITU-T Rec. F.702.

9.2 Assignment, release and withdrawal of chair-control token

9.2.1 Assignment

9.2.1.1 A terminal T_M wishing to assume chair-control may proceed to claim the requisite token if its currently-registered capset from the MCU includes CIC.

T_M requests assignment of the chair-control token by sending CCA. If it does not receive a response (see below) from the MCU in a reasonable time, it may repeat the request. The MCU may optionally provide a mode in which the chair token is pre-allocated at reservation time. In this case, the MCU refuses all requests for the token unless they come from the pre-allocated chair.

9.2.1.2 On receipt of CCA from terminal T_M, the local MCU acts as follows:

a) If it has already assigned its token to a terminal or MCU other than T_M (having transmitted CIT and not received CIS) or received another request to do so from a directly-connected terminal or MCU, the MCU shall respond with CCR.

b) If the MCU has previously assigned its token to T_M, the MCU shall respond with CIT and T_M retains the token.

c) If it has neither assigned its token nor received another request to do so nor have any of the other refusal conditions of a occurred, then:

   i) if it is the only MCU, the MCU sends CIT to T_M. At this point the terminal may begin to issue chair-control commands. The chair terminal may give some indication to the user that the chair token has been received;

   ii) if it is one of two or more interconnected MCUs, three cases must be considered, keeping in mind that only a single chair token exists in a master/slave network, and that the master controls that token:

   – A master has been assigned, and the local MCU is a slave. The slave MCU forwards CCA to its master and awaits CIT. When the slave MCU receives CIT or CCR, it forwards the code to T_M. If the master receives two or more CCAs at the same time, one is picked at random and the rest receive CCR.

   – A master has been assigned, and the local MCU is the master. The master MCU acts on CCA by treating its subordinate MCUs like terminals.

   – If no master has been assigned, the operation is left for further study.

9.2.2 Release of chair token

A change of chair-control should be negotiated between conference participants; the terminal holding the token may release it by sending CIS to the MCU.

Several cases exist:

a) If the MCU is singular, it sends CCR to T_M as a confirmation of the withdrawal of the token.

b) If the MCU is a slave, it forwards CIS to the master MCU, and awaits CCR. On receipt of CCR from the master, the slave forwards CCR to T_M.

c) If the MCU is a master, it acts on CIS by itself while treating its slaves as terminals.

d) If no master has been selected, operation is left for further study.

After the receipt of CCR, T_M is free to request the token again, or another terminal may request the token.

An MCU receiving CIS from any directly connected terminal other than the one to which it previously assigned the token shall respond with CCR.

9.2.3 Withdrawal of chair-control token

The chair-control token may be withdrawn by the MCU. One possible example of this procedure is that two MCUs which have both assigned chair tokens are subsequently connected and one becomes a slave MCU. The slave chair token shall be withdrawn.

An MCU making the withdrawal transmits CCR on the path on which it had sent CIT during the assignment of the token; this will therefore propagate through to T_M. This description applies to both master-slave and MCU-terminal paths. If the MCU making the token withdrawal is a slave MCU, it shall inform the master of the release with CIS after sending CCR to T_M. The master confirms the CIS from the slave with CCR.

A terminal receiving CCR while possessing the token shall cease chair operations immediately; it shall then transmit CIS to the MCU, subsequent operation being as for token release above.

If T_M does not return CIS within a reasonable time, the MCU to which it is connected may act on its behalf in effecting token release according to the above procedures. For the best operation, the master time-out should be longer than the slave time-out.

When the chair-control token has been released or withdrawn, the control of video switching reverts to voice activation (see 6.3). Whether this reversal affects an operative VCS or MCV is left to the manufacturer.

9.3 Information available to the chair-control terminal

The following information is available to a chair-control terminal T_M, provided that the facilities indicated are present in the MCU to which it is connected:

a) the assigned numbers of terminals and MCUs that have been connected – {TIN, <M>, <T>};

b) the numbers of any terminals which have been dropped from the call – {TID, <M>, <T>}. Note that TID is sent only at the time the terminal drops;

c) the terminal number associated with the incoming video – {VIN, <M>, <T>}; (a to c are also available to the other terminals – see below);

d) requests from the floor {TIF, <M>, <T>}.

The values of {TIN, <M>, <T>} and {TID, <M>, <T>} are forwarded by the master as they are received, after this information has been collected from all MCUs (see 7.3.1.2); alternatively, T_M should extract a list of terminal numbers currently in the conference by transmitting {TCU} to the master. The MCU procedures for this operation are described in 7.3.1.3.

9.4 Video selection

9.4.1 Chair-control of broadcast video

The terminal numbers <T> can be acquired at the chair-control terminal T_M by the transmission of TCU or conversationally (bringing each video source up using the voice switching action of the MCU), or by using VCB. By transmitting the symbol {VCB, <M>, <T>} the chair-control terminal determines which video signal will be transmitted to all video-capable parties except the source of that video. On receipt of this symbol, an MCU first inspects the <M> part of the number; if this is not its own value, it sends the video from the master or relevant connected slave to all its video-capable ports; if <M> is its own value, it sends the video signal from the relevant local terminal to all its ports. In addition, it forwards the VCB value to any connected MCU, except that if it received that value from another MCU, it does not reflect the value back there.

T_M may order a return to automatic video switching (see 6.3) by transmitting Cancel-VCB (forwarded to other MCUs). No provision is made for selection of the video to be transmitted to the terminal which is the source of the distributed video. Its local MCU may forward the previous video signal or that from the T_M, if available, or other available signals on a rotating basis (e.g., 20 s at a time), or another basis at the discretion of the manufacturer.

9.4.2 Chair-control of video received at T_M

This clause is optional for chair terminals.

By transmitting the symbol {VCS, <M>, <T>}, T_M determines which video signal will be transmitted to itself (see 6.3.3). If the local MCU has this (optional) capability, and if it also has the requested video signal available, it transmits the requested video to the T_M. If the MCU cannot comply, it returns VCR. To return to the automatic selection of video, the terminal transmits Cancel-VCS.

9.5 Terminal dropping by chair-control

The terminal numbers <T> can be acquired as described in 9.4.1. If it is then desired to disconnect a terminal from the conference, the symbol {CCD, <M>, <T>} is transmitted to the MCU.

NOTE – It has become common practice for a computer to seek the user's confirmation before doing a requested unrecoverable action, such as deleting a file; it is suggested that this precaution be included in the software of the chair-control terminal.

On receipt of this symbol, an MCU first inspects the <M> part of the number and acts thus:

– If <M> is its own value (the terminal is directly connected to the MCU), it disconnects that terminal and transmits the symbol {TID, <M>, <T>} to the port on which it received CCD; it repeats this to all other connected MCUs and terminals.
– If the terminal is connected to another MCU, it repeats the symbol on the inter-MCU link.

On receipt of {TID, <M>, <T>} at one cascaded port of an MCU, it repeats this to all other connected MCUs and terminals, and to T_M if directly connected.

This process results in the correct terminal being dropped even if the token has been assigned to a terminal connected to a slave MCU.

If an MCU receives a command to disconnect a terminal which does not exist or which has already been disconnected, it shall send {CIR} in the direction whence the command came.

9.6 Withdrawal of data tokens by chair-control

The chair-control terminal may transmit the code DCR-L and/or DCR-H, causing the local MCU to send the codes DCR-L/H either to the local terminals holding the relevant data tokens or to the master MCU as appropriate; the effect of this is to cause the cessation of all appropriate data transmission. Subsequent closure of the data channels is according to 8.2.3. This assumes that the MCU itself supports HSD/LSD; if this is not the case, the codes shall be ignored.

9.7 Request for floor

Any suitably equipped terminal may input a "request for the floor", using the symbol TIF.

An MCU receiving TIF forwards it to the chair-control terminal if T_M is locally connected, otherwise it is forwarded to the master MCU to forward it to the chair-control terminal.

9.8 Dropping the entire conference

When an MCU receives the BAS code CCK from the chair-control terminal, it drops the connections related to the conference on which T_M was a participant, including the port used for T_M. The chair terminal T_M should be dropped last to allow the TIDs to arrive to confirm conference drop. Upon receiving CCK, the slave MCU sends CCK to the master (unless it received CCK from the master) and drops all local connections except the link to the master. The TIDs received from the master constitute confirmation of the success of the procedure.

Upon receiving CCK either from a terminal or a slave MCU, the master forwards CCK to all slave MCUs except the one that sent it CCK originally, and then drops all local connections, leaving the links to the other MCUs up. Received TIDs constitute confirmation of the success of the CCK procedure. Note that TIDs shall be forwarded to all slave MCUs.

The inter-MCU links are considered, for the purposes of CCK, to be part of the "cascaded MCU" and are thus left up until the TIDs have been distributed. After this, the links may be dropped at the discretion of the manufacturer.

9.8.1 Identification of token assignment

This clause is optional for all terminals, including chair terminals.

Any suitably equipped terminal may request information on which terminals have been assigned the data and chair-control tokens by using the symbol TCA (Token Command Association). If the connected MCU knows the terminal numbers of the assigned tokens, it responds with the MBE symbol TIR (Token Indicate Response), which includes the terminal numbers {<M>, <T>} of the current holders of LSD, HSD, and chair tokens in this order. The terminal number used when a token is unassigned, or when a capability is not supported is a terminal number of {<M> = 0, <T> = 0}.

If the MCU to which the requesting terminal is connected is a slave supporting this feature and does not know the address of the terminals possessing the tokens, the TCA symbol is forwarded by the slave to the master MCU. The master MCU may need to query slave MCUs to find the terminal numbers associated with particular tokens. The TIR from the master is forwarded by the slave to the requesting terminal.

10 BAS sequencing

The principles of clause 14/H.242 should be followed, with the additions described below.

The MCU transmits the C&I symbol MCC and MCS if relevant to all terminals along with the normal repetitions of BAS commands, to ensure that they remain aware of participation in the multipoint call.

11 Capability exchange during a call

Capability exchanges may be initiated by terminals in the same way as for point-to-point calls (see ITU-T Rec. H.242), and by an MCU when necessary to accommodate the different capabilities declared by the terminals connected (see clause 4).

12 Procedure for loop detection at an MCU

This clause is optional.

NOTE 1 – This clause is not concerned with digital loopback within connected terminals (this is a maintenance function and should not normally occur in conferences, but the MCU may periodically transmit LCO to be sure).

When a loop is placed on a line connected to an MCU (be this within a terminal or elsewhere in the network), the MCU is effectively communicating with itself: an indication of a looped port can be obtained by transmitting a symbol sequence which is sufficiently unique that emulation is very improbable, and looking for that same sequence to appear within a reasonable time in the received signal at that port. Such a test can be made at any or all ports as required according to circumstances (for example, normally every few seconds) provided that the port is not involved in a dynamic mode switch or in a capability exchange.

Either of two sequences may be used, according to circumstances:

1) If the MCU has been numbered, the sequence {MIL, <M>} may be used, since this cannot be generated by any other MCU.
2) Alternatively, the sequence may be as described below.

The sequence consists of {MIL, <N>} where <N> is a random SBE number between 0 and 223 (see ITU-T Rec. H.230). After transmission, the incoming BAS position is monitored for 2 seconds (for example): if the same sequence is returned within that time, the conclusion is reached that the port is indeed looped (but see Note 2 below), and further action depends on the internal software (e.g., disconnecting the port from a conference if one is in progress, perhaps timing the loop delay for diagnostic purposes). The test could also be repeated for greater confidence, using a different random number.

NOTE 2 – Where there is a possibility that the test is being carried out simultaneously by another connected equipment (e.g., when MCUs are connected together), it needs to be established that the received sequence could not have been generated elsewhere; the test should be repeated twice with different random numbers; the probability of a false indication is then reduced to a very low value.

13 Exceptional procedures

13.1 A terminal connected does not indicate capability for the SCM

The MCU transmits to this terminal a reduced capset, consisting of cap-mark and at least one audio capability code.

Communication proceeds as in clause 5, except that the transmission mode between this terminal and the MCU is in a lower mode. The MCU transmits MIS to this terminal, indicating that it has been accorded secondary status (see ITU-T Rec. H.231).

13.2 Contention resolution principle

In a master/slave situation, the slave shall adopt the choice made by the master, and the master shall ignore the action taken by the slave, expecting this to be corrected within a short time.

Where both MCUs on an inter-MCU link happen to transmit conflicting commands at about the same time, then instead of acting on the incoming value, each transmits a random SBE number (see ITU-T Rec. H.230). The MCU which receives a higher number than it transmitted maintains the decision already taken, while the other shall adopt the action chosen by the first. If both numbers happen to be the same, the process is repeated.

If one or both of the random SBE numbers sent in this procedure are lost, confusion can occur. Therefore, each MCU shall set a 5-second timer. If, at the end of five seconds after sending a random SBE, the MCU has not received a random SBE, the MCU shall send a new random SBE. If after three attempts no response is received, the sending MCU should consider the other MCU to be incapable of supporting the contention resolution principle or to be broken. Action taken at this point is at the discretion of the manufacturer.

When operating in a non-master/slave mode, the contention resolution principle should be used once, and the outcome of that contest then serves to guide any future contests. The purpose of this suggestion is to attempt to avoid a large number of uses of the contention resolution process which may greatly slow MCU operations. The contention resolution process should not be run on two or more conflicts at the same time.

Note that it may take different MCUs different lengths of time to detect the "command collision". Thus, one MCU might receive the random number BEFORE it detects the "collision". It could potentially "cheat" and always end up as the master MCU. However, since the purpose of the contention resolution principle is to arbitrarily resolve conflicts, this lack of strict fairness is not of concern.

14 MCU loopback procedures

MCU loopback procedures are optional. Furthermore, any particular MCU may implement any one or any combination of the procedures described in this clause. A loopback command from a particular terminal should not affect any other terminal in the conference, except indirectly in the sense that the current video source may change. Note that in the point-to-point loopbacks of ITU-T Rec. H.242, LCA and LCV involve looping an analogue signal, but the MCU equivalent will not generally involve an analogue signal. It is suggested that the loopback commands be repeated several times for more robust operation. The MCU shall not pass any loopback command through to any other terminals in the conference.

The procedures of this clause follow the principle used in the point-to-point loopback of ITU-T Rec. H.242 that the far end not initiating the loopback should continue to see/hear the looped media. For an MCU, this means that the looped media will continue to be a part of the conference.

On receipt of LCD in the BAS stream for a particular port, the MCU shall loop the entire multiplex from that terminal back toward the terminal. The MCU shall continue to pass the multiplex from the looped terminal into the conference. The MCU shall continue to monitor the BAS stream from the terminal in loopback for Loopback Command Off (LCO), and deactivate the loopback condition when it is received. Also, the MCU shall not participate in any capability exchanges or mode changes on the looped back port while the digital loopback is active. The MCU should not allow any SCM changes while the loopback is in operation since those changes cannot be communicated to the looped terminal.

On receipt of LCA in the BAS stream for a particular port, the MCU shall loop the incoming (toward the MCU) audio for that port back to the requesting terminal, while continuing to send the audio to the conference audio mixer. Video, data, and BAS processing will be unaffected. While in this state, audio from the port in loopback should continue to be used to drive the operation of the voice activated switching feature. The MCU shall continue to monitor the BAS stream from the terminal in loopback for LCO, and deactivate the loopback condition when it is received. During the loopback interval, the MCU should continue to react to capability exchanges on the Audio looped port as it normally would.

On receipt of LCV in the BAS stream for a particular port, the MCU shall loop the incoming (toward the MCU) video for that port back to the requesting terminal while continuing to send the video to other conference participants. Audio, data, and BAS processing will be unaffected. While in this state, commands such as MCV, VCS, and VCB directed at the port in loopback may be ignored by the MCU at the manufacturer's discretion. The MCU shall continue to monitor the BAS stream from the terminal in loopback for LCO, and deactivate the loopback condition when it is received. During the loopback interval, the MCU should continue to react to capability exchanges on the Video looped port as it normally would.

When a loopback command is deactivated, the MCU should impose the current SCM on the previously looped terminal with appropriate commands. The looped terminal should not change its capability set while in the looped state.

15 Interactions with T.120 control

15.1 Chair-control interactions

Chair-control may also be assigned using ITU-T Rec. T.120; when terminal numbers and chair-control token have been assigned using the T.120, and all terminals in the conference are T.120-equipped, those assignments take precedence and the procedures of this clause shall not be invoked. When a T.120 channel is open between a terminal and its local MCU in such a conference, then the BAS codes referenced in the definition of CIC (see ITU-T Rec. H.230) shall not be transmitted. Note that based on the conference SCM, a terminal may move between these two control modes over a period of time as the MLP channel is opened and closed.

However, in a conference with both T.120-capable and non-T.120-capable terminals, the MCU should continue to assign terminal numbers according to this clause using TIA since those are the only terminal numbers that the non-T.120 terminals will understand. The MCU should also collect identity strings from the non-T.120 terminals using BAS procedures. Non-T.120 terminals may continue to use the procedures of this clause to request terminal identity strings as well. Non-T.120 terminals should also continue to receive TIN, TID, and VIN. In this case, T.120 chair-control may make use of H.243 terminal numbering to control non-T.120-equipped terminals. However, in all such mixed conferences, T.120-based chair-control should be given precedence over H.243 chair-control, and the H.243 chair token should not be granted to any terminal. Non-T.120-equipped terminals should not experience a reduction in function (other than being unable to become chair) when participating in T.120 conferences.

If T.120 control is not desired for the mixed conference, the MCU should be appropriately commanded to remove T.120 from its capability set, with the result that H.243 control will be used. This might be the case if the chairperson's site supports only H.243 control, and it is imperative for that site to be in control.

15.2 Interactions with passwording

The H.243 password allows entrance into the audio/video conference. The meaning of the T.124 password is defined in ITU-T Rec. T.124.

In general, for a T.120 (non-H.221) terminal, only the T.120 password can be collected, and for non-T.120-capable H.320 endpoints, only the H.243/H.230 password can be collected. However, for H.320 terminals with T.120 capability, the MCU may collect the password twice, once at the H.243 level to allow entrance to the audio/video conference, and once at the T.120 level to allow entrance to the data conference. It is strongly suggested that MCU manufacturers consider shielding the user from the existence of the two log-in procedures by requesting the password once unless some operational purpose is served, such as a higher level of security for the data conference. Note that since the data conference may not start immediately after the T.120-capable H.320 terminal joins the conference, a clear need exists to collect H.243/H.230 passwords from T.120-capable H.320 terminals. The manufacturer has the option of foregoing the H.243-level password or the T.120-level password, or both as deemed appropriate.

H.243 passwords may be required of T.120-capable H.320 terminals.

15.3 Interactions with TIX/TIA

Since there is no call association procedure in ITU-T Recs T.120/T.124/T.128, TIX/TIA operations shall apply to all terminals whether T.120-capable or not.

15.4 Interactions with SCM management

A situation may arise where a terminal is not capable of meeting the MLP/T.120 SCM for the conference, but is still MLP/T.120-capable. The MCU may, at the manufacturer's discretion:

a) exclude the terminal from the T.120 conference and use BAS commands to collect passwords;

b) open the MLP channel at the mandatory 6.4 kbit/s rate, and rely on ITU-T Rec. T.120 to negotiate the situation. In this case, it is at the manufacturer's discretion whether or not the terminal receives audio/video before T.120 negotiations are complete.

15.4.1 Interactions with video mixing

When the methods of T.120 are in use by an MCU in a conference consisting of mixed T.120 and non-T.120 terminals, the following rules shall be followed:

a) In general, the MCU will provide H.243 "indications" such as VIC and VIN2 to allow the H.243 terminals to understand the mixed image. In the event that a more complex composed image is being used than is described in ITU-T Rec. H.243, this is indicated by VIC <0> and VIN2 <M><T><0>. The H.243 terminal may use this information to inform the user that it is incapable of processing the more complex image other than to simply display it.

b) Control signals such as MCV and VCS from the H.243 terminals will receive appropriate negative responses such as VCR.

c) T.120-type command will be used to control the type of picture composition, as well as the location of each picture in the composition.

If all terminals are T.120-equipped, the MCU does not need to issue VIC or VIN2, and may instead make use of the appropriate T.128 signals.

15.5 Rate matching in a cascade

Various situations can exist in a cascade involving H224_MLP and H224_LSD-equipped terminals and MCUs, which are documented in Table 6:

Table 6/H.243 – Allowed H.224 modes in a cascade

Case number	Capability for terminal attached to MCU #1	H.224 caps for MCU #1	H.224 caps for MCU #2	Capability for terminal attached to MCU #2	Allowed mode for far-end camera control using ITU-T Rec. H.224
1	H224_MLP, H224_LSD	H224_MLP	H224_MLP	H224_MLP, H224_LSD	H224_MLP
2	H224_MLP, H224_LSD	H224_MLP	None	H224_MLP, H224_LSD	Not allowed
3	H224_MLP, H224_LSD	H224_MLP	H224_LSD	H224_MLP, H224_LSD	Not allowed
4	H224_MLP, H224_LSD	H224_MLP	H224_LSD, H224_MLP	H224_MLP, H224_LSD	H224_MLP
5	H224_MLP, H224_LSD	H224_LSD	H224_MLP	H224_MLP, H224_LSD	Not allowed
6	H224_MLP, H224_LSD	H224_LSD	None	H224_MLP, H224_LSD	Not allowed
7	H224_MLP, H224_LSD	H224_LSD	H224_LSD	H224_MLP, H224_LSD	H224_LSD
8	H224_MLP, H224_LSD	H224_LSD	H224_LSD, H224_MLP	H224_MLP, H224_LSD	H224_LSD
9	H224_MLP, H224_LSD	H224_MLP, H224_LSD	H224_MLP, H224_LSD	H224_MLP, H224_LSD	Either H224_MLP or H224_LSD
NOTE – It is noted that several scenarios exist where interoperability problems exist, while there is one scenario in which full interoperability is always possible.					

Appendix I

C&I signals defined in ITU-T Rec. H.230

Abbreviation	Notes	Description
AggIN	S	Aggregation indicating the number n as per H.244
AIA		Audio Indicate Active
AIM		Audio Indicate Muted
AMC-open	S	Additional Media Channel Open
AMC-close	S	Additional Media Channel Close
CCA		Chair-control Command Acquire
CCD	S	Chair-control Command Disconnect
CCK		Chair-control Command Kill
CCR		Chair-control Command Release/Refuse
CIC		Chair-control Indicate Capability
CIR		Chair-control Indicate Release/Refuse
CIS		Chair-control Indicate Stopped-using-token
CIT		Chair-control Indicate Token
DCA-L	S	Data (LSD) Command Acquire
DCA-H	S	Data (HSD) Command Acquire
DCC-L		Data (LSD) Command Close
DCC-H		Data (HSD) Command Close
DCM		Data Command MLP
DCR-L		Data (LSD) Command Release/Refuse
DCR-H		Data (HSD) Command Release/Refuse
DIS-L		Data (LSD) Indicate Stopped-using-token
DIS-H		Data (HSD) Indicate Stopped-using-token
DIT-L		Data (LSD) Indicate Token
DIT-H		Data (HSD) Indicate Token
h239ControlCapability	C	Capability – indicates H.239 support
IIS	M	Information Indicate String
LCA		Loopback Command, Audio Loop Request
LCD		Loopback Command, Digital Loop Request
LCO		Loopback Command Off
LCV		Loopback Command, Video Loop Request
MCC	C	Multipoint Command Conference
MCN		Multipoint Command Negating MCS
MCS		Multipoint Command Symmetrical Data-transmission
MCV	C	Multipoint Command Visualization-forcing
MIH		Multipoint Indicate Hierarchy
MIJ		Multipoint Indicate Joined Real Conference

Abbreviation	Notes	Description
MIL	S	Multipoint Indication Loop
MIM		Multipoint Indicate master MCU
MIS	C	Multipoint Indication Secondary-status
MIV	C	Multipoint Indication Visualization
MIZ	C	Multipoint Indication Zero-communication
MMS	C	Multipoint command Mode-Symmetrize
MVA		Multipoint Visualization Achieved
MVC		Multipoint Visualization Capability
MVR		Multipoint Visualization Refused/Revoked
NCA-i		Network Command send address-initial
NCA-a		Network Command send addresses-additional
NIA-s		Network Indicate addresses – using SBE
NIC		Network Indicate consecutive addresses
NID		Network Indicate double addresses
NII		Network Indicate incompatible-aggregators
NIQ-s		Network Indicate query address – using SBE
NIQ-m	M	Network Indicate query address – using MBE
NIR		Network Indicate refuse-address
NIS		Network Indicate same addresses
RAN	S	Random Number
RID		Restrict Indicate Denied
RIU		Restrict Indicate_Unrestricted
RIR		Restrict Indicate Request
TCA		Token Command Association
TCI		Terminal Command Identity
TCP	S	Terminal Command Personal-Identifier
TCP-5	S	Terminal Command Unicode Personal-Identifier
TCS-n		Terminal Command String
TCU		Terminal Command Update
TIA	S	Terminal Indicate Assignment
TIC		Terminal Indicate Capability
TID	S	Terminal Indicate Dropped
TIE		Terminal Indicate End of Listing
TIF	S	Terminal Indicate Floor-request
TII	S	Terminal Indicate Identity
TIL	M	Terminal Indicate List
TIN	S	Terminal Indicate Number
TIP	M	Terminal Indicate Personal-Identifier
TIP-5	M	Terminal Indicate Unicode Personal-Identifier
TIR	M	Token Indicate Response

Abbreviation	Notes	Description
TIS		Terminal Indicate identity-Stop
TIX		Terminal Indicate additional-channel-X
VBMBC	C	Video Capability "videoBadMBsCap"
VCB	S, C	Video Command Broadcast
VCF		Video Command Freeze-Picture Request
VCR		Video Command Release/Refuse
VCS	S, C	Video Command Select
VCU		Video Command Fast Update Request
VIA		Video Indicate Active
VIA2		Video Indicate Active 2
VIA3		Video Indicate Active 3
VIC	S	Video Indicate Compose
VIM		Video Indicate Mixing (capability)
VIN	S	Video Indicate Number
VIN2	S	Video Indicate Number 2
VIR		Video Indicate Ready-to-activate
VIS		Video Indicate Suppressed
VSTRDEL		Video Indicate "video spatial temporal tradeoff encoder level"
1997Recs.		Indicate conformance to 1997 versions of ITU-T Recs H.221, H.242 and H.230
S Followed by an SBE number or alphanumeric. M Uses start-MBE. C Cancel signals are also defined.		

Appendix II

Mandatory and optional codes for MCUs

The C&I functions are defined such that, under various appropriate circumstances, the audiovisual system will operate in a fault-free manner and also such that sympathetic presentation to users is possible. Some functions are therefore mandatory, others optional.

CM Conditionally mandatory: If the terminal (or MCU) is capable of entering the given state, then it shall transmit the given code and, when leaving that state, the complementary code. If it has no such capability it can ignore both.

M Mandatory: For all equipments of either terminal or MCU type.

X Non-mandatory: On receipt of such a code, it may be unrecognized, or recognized but not acted upon, or recognized and acted upon, entirely at the discretion of the manufacturer or user.

NA The code is not applicable in that case.

In this Appendix, only the most simple case is shown, either terminal-MCU or MCU-MCU with no consideration of any attached terminals, i.e., for MIM. In general, in the cascade case, the slave MCU takes on the role of a terminal to the master MCU, and thus the terminal requirements for send and receive codes are also placed on the slave MCUs. To avoid duplication, the codes covered in ITU-T Rec. H.230 are marked with a (#).

Code first 3 bits	Code last 5 bits in decimal form	Abbreviation	Transmit		Receive		Reference for procedures
			Terminal	MCU	Terminal	MCU	
Code (000)	[0,1]	Reserved					
	[2]	AIM	#	#	#	#	H.230
	[3]	AIA	#	#	#	#	H.230
	[4]	ACE	#	#	#	#	H.230
	[5]	ACZ	#	#	#	#	H.230
	[6]-[7]	Reserved for audio-related symbols					
	[8]	TCI	NA	X	X	NA	H.243
	[9]	TII*	X	NA	NA	X	H.243
	[10]	TIS	X	NA	NA	X	H.243
	[11]-[15]	Reserved					
	[16]	VIS	#	#	#	#	H.230
	[17]	VIA	#	#	#	#	H.230
	[18]	VIA2	#	#	#	#	H.320
	[19]	VIA3	#	#	#	#	H.320
	[20]	VIC*	NA	CM	CM	NA	H.243
	[21]	VSTRDEL	#	#	#	#	H.230
	[22]	VIN2***	NA	CM	CM	NA	H.243
	[23]	VIM	X	X	X	X	H.243
	[24]	VBMBC	#	#	#	#	H.230
	[25]-[30]	Reserved for video-related symbols					
	[31]	VIR	#	#	#	#	H.320

Code first 3 bits	Code last 5 bits in decimal form	Abbreviation	Transmit		Receive		Reference for procedures
			Terminal	MCU	Terminal	MCU	
Code (001)	[0]	MCC	NA	M	M	NA	H.243
	[1]	Cancel-MCC	NA	X	M	NA	H.243
	[2]	MIZ	NA	X	X	NA	H.243
	[3]	Cancel-MIZ	NA	CM	X	NA	H.243
	[4]	MIS	NA	X	X	NA	H.243
	[5]	Cancel-MIS	NA	CM	X	NA	H.243
	[6]	MIM	NA	CM	NA	CM	H.243
	[7]	TIC	X	X	X	X	H.243
	[8]	TIX**	CM	NA	CM	NA	H.243
	[9]	RAN	NA	X	NA	X	H.243
	[10]	MIH	NA	X	NA	X	H.243
	[11]	TIA**	CM	CM	CM	CM	H.243
	[12]	TIN**	NA	CM	X	NA	H.243
	[13]	TID**	NA	CM	X	NA	H.243
	[14]	TCU	X	NA	NA	CM	H.243
	[15]	TCA	X	NA	NA	X	H.243
	[16]	MCV	X	NA	NA	X	H.243
	[17]	Cancel-MCV	CM	NA	NA	CM	H.243
	[18]	MIV	NA	X	X	NA	H.243
	[19]	Cancel-MIV	NA	CM	X	NA	H.243
	[20]	MCS	NA	M	M	NA	H.243
	[21]	MCN	NA	X	M	NA	H.243
	[22]	VIN**	NA	CM	X	NA	H.243
	[23]	VCB**	X	NA	NA	CM	H.243
	[24]	Cancel-VCB	CM	NA	NA	CM	H.243
	[25]	VCS**	X	NA	NA	CM	H.243
	[26]	Cancel-VCS	CM	NA	NA	CM	H.243
	[27]	VCR	NA	CM	X	NA	H.243
	[28]	MMS	NA	X	CM[a)]	NA	H.243
	[29]	Cancel-MMS	NA	X	CM	NA	H.243
	[30]	Cancel-MIM	NA	X	NA	CM	H.243
	[31]	MIL*	X	X	CM	CM	H.243
Code (010)	[0]	CIC	NA	X	X	NA	H.243
	[1]	CCD**	X	NA	NA	CM	H.243
	[2]	CIR					H.243
	[3]	CCK	X	NA	NA	CM	H.243
	[4]	CCA	X	NA	NA	CM	H.243
	[5]	CIT	NA	CM	CM	CM	H.243
	[6]	CCR	NA	CM	CM	NA	H.243
	[7]	CIS	CM	NA	NA	CM	H.243
	[8]	TIF**	X	NA	NA	CM	H.243
	[9]	TIE	NA	CM	CM	NA	H.243
	[10]-[11]	Reserved					
	[12]	MVC	X	X	X	X	H.243
	[13]	MVA	NA	CM	CM	CM	H.243

Code first 3 bits	Code last 5 bits in decimal form	Abbreviation	Transmit		Receive		Reference for procedures
			Terminal	MCU	Terminal	MCU	
Code (010) *(cont.)*	[14]	MVR	NA	CM	CM	CM	H.243
	[15]	MIJ	NA	X	X	NA	H.243
	[16]	DCA-L	X	NA	NA	CM	H.243
	[17]	DIT-L	NA	X	CM	NA	H.243
	[18]	DCR-L	X	X	CM	CM	H.243
	[19]	DIS-L	CM	NA	NA	CM	H.243
	[20]	DCC-L	CM	NA	NA	CM	H.243
	[21]-[23]	Reserved					
	[24]	DCA-H	X	NA	NA	CM	H.243
	[25]	DIT-H	NA	X	CM	NA	H.243
	[26]	DCR-H	X	X	CM	CM	H.243
	[27]	DIS-H	CM	NA	NA	CM	H.243
	[28]	DCC-H	CM	NA	NA	CM	H.243
	[29]-[30]	Reserved					
	[31]	DCM	X	NA	NA	CM	H.243
Code (011)	[0]	Reserved					H.243
	[1]	TCS-1	NA	X	CM	NA	H.243
	[2]	TCS-2	NA	X	X	NA	H.243
	[3]	TCS-3	NA	X	CM	NA	H.243
	[4]	TCP**	X	NA	NA	CM	H.243
	[5]	AggIN*	#	#	#	#	H.244
	[6]	NCA-i	#	#	#	#	H.242
	[7]	NCA-a	#	#	#	#	H.242
	[8]	NIS	#	#	#	#	H.242
	[9]	NIC	#	#	#	#	H.242
	[10]	NID	#	#	#	#	H.242
	[11]	NII	#	#	#	#	H.244
	[12]	TCP-5**	#	#	#	#	H.243
	[13]	NIA-s	#	#	#	#	H.242
	[14]	NIQ-s	#	#	#	#	H.242
	[15]	NIQ-m	#	#	#	#	H.242
	[16]	NIR	#	#	#	#	H.242
	[17]	TCS-4	#	#	#	#	H.242
	[18]	TCS-5	NA	X	X	NA	H.243
	[19]-[28]	Reserved					
	[29]	RIR	NA	CM	NA	CM	H.243
	[30]	RID	NA	CM	NA	CM	H.243
	[31]	RIU	NA	CM	NA	CM	H.243

Code first 3 bits	Code last 5 bits in decimal form	Abbreviation	Transmit		Receive		Reference for procedures
			Terminal	MCU	Terminal	MCU	
Code (101)	[0]	1997Recs.	#	#	#	#	H.230
	[1]	H239ControlCapability	#	#	#	#	H.239
	[2]	AMC-open**	#	#	#	#	H.239
	[3]	AMC-close*	#	#	#	#	H.239
	[4]-[31]	Reserved					
Code (111)	All values forbidden						
Codes listed in Annex A/H.221							
		VCF	#	#	#	#	H.230
		VCU	#	#	#	#	H.230
		LCV	#	#	#	#	H.230
		LCA	#	#	#	#	H.230
		LCD	#	#	#	#	H.242, H.320
		LCO	#	#	#	#	H.242, H.320

a) If the terminal supports H.262 and/or H.263 reception is MMS mandatory; hence CM appears in the column.

* The number of * indicates how many SBE number or SBE character values must follow the symbol.

# Indicates the direction in which the symbol is transmitted.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
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	Telecommunication management, including TMN and network maintenance
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, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2006