ITU Telecommunication Standardization Sector | **Document AVC-830**
STUDY GROUP 15 | **Original: English**

Yokosuka, 24 - 27 October 1995

Question(s): 2/15

Source: TU Berlin / TELES GmbH
Jörg Ott / Peter Kratz Voice: +49 30 314-73389
{jo,kratz}@cs.tu-berlin.de Fax: +49 30 314-25156

Title: Connection Management Procedures for H.323
TUB/TEL-01-C

Purpose: Proposal / Discussion

1. Introduction

This paper is a contribution to draft recommendation H.323. It proposes a connection management procedure for extending the usually ISDN-based videophone and videoconference services — as defined in the H.320 and T.120 series of ITU-T recommendations — to LAN-based end systems (H.323 terminals). The purpose is to allow "videophone" connections between two H.323 systems as well as interoperability between H.323 and H.320/H.324 systems.[1]

This paper provides an overview of the connection establishment for H.323 terminals.[2] For the time being, only point-to-point connections are considered. Multipoint scenarios can be created only by using an MCU. Using the multicast capability of most LANs for more efficient distribution of information is subject to further study.[3]

One of the crucial points for the protocol design is the ability to run identical applications on both ISDN-based and LAN-based end systems. In particular, standardized interfaces such as the CAPI for ISDN applications should be able to provide precisely the same functionality when running in a LAN environment. This assumption implies that all events indicated to an application from the beginning of a connection setup to the connection teardown in the ISDN are available in a LAN environment as well, and that addressing mechanisms work similarly from an application's point of view.

---

[1] In this contribution, we will primarily refer to ISDN but the same procedures are expected to work with other WANs (e. g. PSTN) as well.

[2] Throughout this paper we will use the terms *LAN-based terminal* (or *end system*) and *H.323 terminal* (or *end system*) synonymously. The same applies to *ISDN-based terminals* and *H.320 terminals.*

[3] A second contribution by Technische Universität Berlin and TELES entitled "Transmission of T.120 information in a LAN environment" — that has been submitted to Q2/15 and to Q10/8 of ITU-T — deals with the protocol aspects of running T.122/T.125 on top of a multicast service. The administrative issues of using multicast addresses are discussed in document AVC-831.

2. Interconnection scenarios

The interworking scenarios to be taken into account are shown in figure 1:[4] [5]

a) point-to-point interconnections of two H.323 terminals in a LAN,

b) multipoint conferences using an H.323 MCU among H.323 terminals all located in the same LAN,

c) point-to-point connections between an H.323 terminal and an H.320 terminal via a gateway,

d) multipoint conferences using an H.320 MCU involving any number of H.323 terminals in a LAN and any number of H.320 terminals, and

e) multipoint conferences using an integrated gateway/H.320 MCU (optionally) cascaded with an H.320 MCU with any any number of H.323 and H.320 terminals.

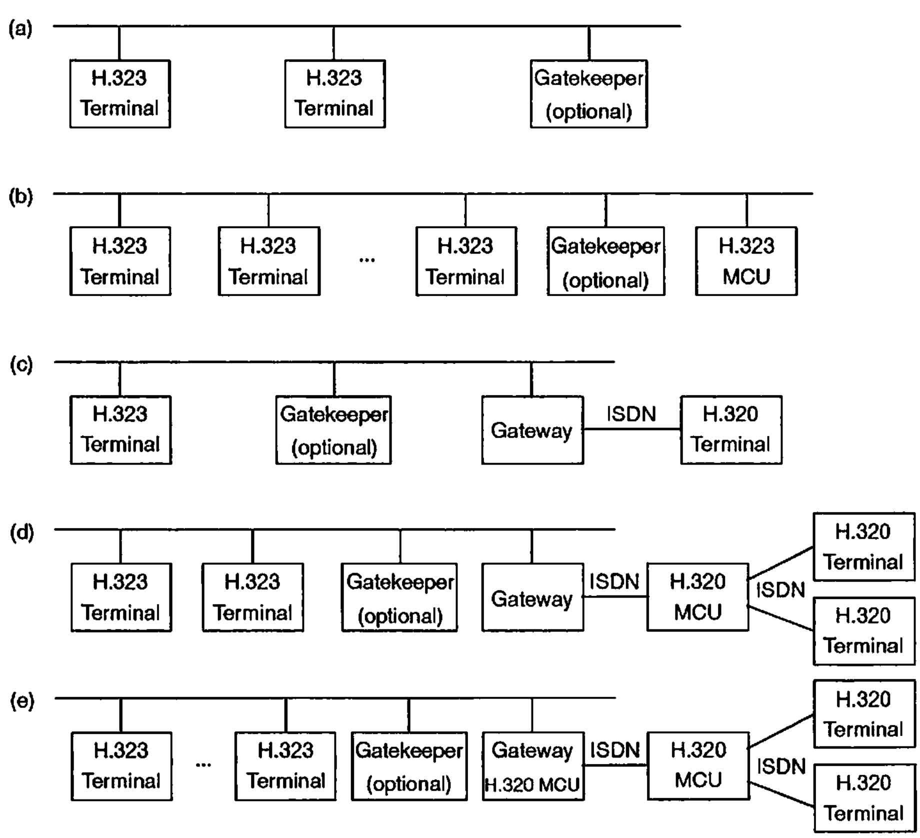

Figure 1: Interconnection scenarios

In scenarios c) through e), a LAN is connected to the ISDN by means of a gateway. The gateway performs all required adaptation from the LAN to the WAN and vice versa. The gateway may be

[4] These figures are taken from Vineet Kumar's contribution.

[5] Further important scenarios that have been identified by the LAN group of the IMTC as well as ITU-T Q2/15 are LAN to LAN interconnection via a WAN and WAN to WAN interconnection via a LAN. These scenarios are not yet covered by this paper.

co-located with an MCU but need not to be, neither does it need to be an MCS node in T.120 teleconferences.

Note, that the gatekeeper shown in figure 1 is an optional component, the functionality of which may also be built into the MCU, in the terminal, or may be provided in an entirely distributed fashion.

3. Connection Management

Connection management includes setup of a connection (with some administrative issues, such as resource allocation), transmission of information, and teardown of a connection. A connection consists of

- a connection control part (the D channel in ISDN) for out-of-band user-to-network signaling and
- one or more information streams, and a control stream (that are multiplexed on the B channel(s) in ISDN).

This LAN counterpart of an ISDN connection is called a H.22Z connection. We use the overall connection setup procedure outline which seems to be consensus of Q2/15.[6]

Connection setup in the ISDN environment is performed via the D channel protocol (Q.931). Q.931 PDUs are used to request a B channel at the BRI or PRI, and to connect the B channel to some remote end point identified by an E.164 address (see below). During the setup the network uses Q.931 PDUs to inform the two communicating entities about the progress of the call to be established. After successful setup, the connected end systems may communicate through a 64kbps transparent layer 1 data channel (i. e. without any higher layer protocols imposed by the ISDN).

H.320 defines five phases for the communication between two H.320 terminals. The protocol used in the LAN basically follows these phases with some additional functionality required for the connection setup and a different multiplexing scheme. The following sections describe how the five phases defined in H.320 are to be refined for a LAN environment and how LAN-based terminals interact with H.320 terminals in the respective phases. Figures 2 and 3 illustrate the connection setup process for H.323 to to H.320 terminals in both directions. In figure 2, replacing the gateway by an H.323 terminal and omitting the right hand side of the figure yields the connection setup process between two H.323 terminals. Note, that clarity the gatekeeper is not depicted in these figures, but the queries sent to it are indicated.

3.1. Phase A: call setup

3.1.1. Phase A1: address translation and authentication

Addressing of H.323 terminals is done using E.164 addresses.[7] For calls coming in through a gateway from ISDN/PSTN/ATM the gateway is the natural place for the address translations —

[6] Proposed by Dale Skran (e-mail dated 13 July 1995).

[7] This seems to be consensus in the group proposed in several contributions and the recent audioconference.

note, however, that the database used for this translation does not necessarily reside on the gateway but may also be maintained by some separate address translator that is considered to be part of the *gatekeeper* functionality.[8] The translation table may be filled and updated in two ways:[9]

1. A static configuration file is used to initialize to the translation table at startup. This prevents startup delays in which the gateway does not have sufficient information to direct incoming calls to the terminals.

   The format of the configuration file is as follows:[10]

   ```
   extension = network_address (e.g. dot notation for IPv4)
   extension = hostname
   extension = hostname.domainname
   ```

   The table syntax should allow to specify hostnames for they are a) more convenient and b) allow for transparently making use of future mobility concepts that are being discussed e. g. in the IETF. Finally, note that also IPv6 addresses (which are 16 bytes long) should be allowed.

2. The mapping table at the gatekeeper may be updated by means of register/deregister messages.[11] However, the registration process is not mandatory:[12] whether or not a person (with a preconfigured E.164 to LAN address mapping) is currently reachable through his/her H.323 phone can be determined by simply trying to connect to the respective terminal:

   a) A *Destination unreachable* message from ICMP will indicate that the system is switched off, not connected to the network, or the desired protocol or application is not available (at the moment).

   b) Some Q.931 message will indicate that the system is working but the callee is currently not able or not willing to answer the call.

   Register/deregister messages will only be used to change the state of the internal mapping table: to register users who are not (yet) included in the mapping file, to change registrations allowing users to redirect incoming calls to other hosts, and to disable a mapping (i. e. to make a person non-reachable via H.323).

   Obviously, these messages need to be authenticated in some way. This is for further study.

---

[8] Some contributions have addressed the issue of translation table sizes. From our point of view these discussions have shown that a mapping table — even if ASCII-based and containing some thousand entries — is small enough that this approach is feasible.

[9] The following paragraphs assume the usage of TCP/IP, but they should work in a similar way for SPX/IPX.

[10] Partially following what Santo Wiryaman from PictureTel used for some calculations.

[11] As was outlined by Mark Reid in the Pte11 contribution. The messages RRQ, RCF, RRJ, URQ, UCF, URJ, TLR, and TRL described in section 3 of Pte11 would be used.

[12] In the future, registering and deregistering users and their systems is likely to become an issue of local network management which in turn will provide other mechanisms to figure out the current address (if any) of a particular user.
Also, as pointed out by Terry Lyons in the audioconference call on 12 October 1995, other (Internet) services — such as the domain name service — might be usable for this purpose. In this case, no registration protocol would be needed at all but an existing one would be reused. This requires studying the feasibility of such existing services for the purposes of Q2/15: e. g. could hostnames like *extension-number.h323.company.com* be assigned to H.323 terminals leaving the address mapping and resolution management to the Internet domain name service? Resolving could then be done e. g. using *gethostbyname(3N)*.

How an H.323 terminal or a gatekeeper learns about the address of such an *address translator* is a local matter.

In a LAN-internal call the user should also specify an E.164 address but the mapping may take place locally — within the calling terminal — so that contacting the gatekeeper may be avoided. While this basically works only for statically assigned addresses, in "small" LANs also some address resolution protocol using a well-defined multicast address might be employed. However, this should be kept optional and it is up to each enterprise to decide whether their LAN is small and lightly loaded enough and the environment sufficiently trustworthy to make use of such a feature (this puts into question whether or not to standardize such an H.323 address resolution protocol at all).

For incoming calls the caller must be authenticated before being allowed access to the LAN. A first authentication, based on the information conveyed in the D channel, may be done at this point. Which other techniques are employed is for further study.

Phase A1 fails — and thus no connection is set up — if the address translation cannot be performed because

- the address cannot be found in the local translation table; or
- no local translation table exists and
- no address translator can be found; or
- the translator cannot resolve the address.

Phase A1 also fails if the call cannot be authenticated, i. e.

- for an incoming call, the calling user is not allowed to call the particular phone number (or to a terminal on the LAN at all) or
- for an outgoing call, the calling user is not allowed to dial out of the local network (to the specified phone number).

3.1.2. Phase A1a: resource allocation with gatekeeper (optional)

After address translation is completed, a minimum resource reservation must be made with the gatekeeper if this entity is present in the LAN. This "reservation" is basically the query whether a connection of a certain type can be accepted leaving detailed issues of QoS management hidden behind the gatekeeper: admission policies, the reservation protocols used, etc.

If the connection setup is initiated by an H.323 terminal (either to the gateway or to another H.323 terminal) the terminal shall request the maximum QoS parameters it wants to use during the call from the gatekeeper. A terminal that is not sure about its peer's capabilities may decide to go for a lower QoS (in particular, less bandwidth) initially, and change this later on.

If the connection setup is initiated by an ISDN or a PSTN terminal, the gateway shall allocate (via a simple request to the gatekeeper) a network specific default value for the incoming connection (which might also depend on the service indicator in the ISDN case). Default values for the bandwidth could be, for example,

- 110 kbps for PSTN (64 kbps for PCM coded audio plus approx. 50 kbps for video) if transcoding G.723 audio is done in the gateway, or

- 170 kbps for ISDN (64 kbps for PCM coded audio plus approx. 100 kbps for video) after transcoding G.728 audio in the gateway if it can be foreseen that an audiovisual connection shall be set up.[13] Note, that bandwidth reservation for data information (e. g. T.120) is not yet considered.
- 64 kbps for ISDN if a simple audio (or data?) connection is expected to be set up.

This reservation can be changed if subsequently more or less than the expected bandwidth is to be used.

Resource reservation and address resolution might be packed into a single request to the gatekeeper.

If the resource reservation fails, the call setup is rejected. This means that the calling terminal (H.323 or H.320) receives a *BUSY* indication.

3.1.3. Phase A2: establishment of a LAN-D channel

Instead of the ISDN D channel layers 1 and 2, a connection-oriented, reliable transport is used in the LAN to transmit the Q.931 PDUs (the LAN-D PDUs). This transport connection is called the LAN-D channel. As Q.931 is an asymmetric protocol (between NT and TE) these roles must be dynamically assigned to the communicating peers: the calling entity (terminal or gateway) will always assume the role of the TE (which besides defining the Q.931 behavior particularly means to care for resource reservations for the call and other administrative issues).

For this transport connection, a well-known transport address (e. g. a port number in a TCP/IP environment) is used. There is at most one LAN-D channel set up between each LAN end system and the gateway or two LAN end systems in each calling direction, respectively. The calling direction describes which terminal is the originator and which terminal is the recipient of a call, i. e. if terminal A sets up a LAN-D channel to terminal B, it creates an A→B LAN D channel. Terminal A re-uses this LAN-D channel for the setup of further connections to the same recipient B. If B decides to call A, it does not make use of a possibly existing A→B LAN-D channel, but creates a new LAN-D channel (in the reverse calling direction, B→A) for this and all subsequent calls. In case of a subsequent call in the same calling direction, phase A2 is skipped.

The Q.931 PDUs are exchanged using the packetization offered by the TPKT header of RFC1006 on top of a reliable LAN protocol (TCP or SPX). Using TPKT guarantees the packet boundaries on which Q.931 relies.[14]

The LAN-D channel connection is established when needed (for the first time) for each calling direction and is released with the last H.323 connection (set up in this calling direction) between two systems. The release of a particular LAN-D channel is always initiated by the terminal or the gateway that created this channel.

---

[13] These values are just provided to give some rough numbers. Any suggestions and comments on this issue are welcome.

[14] There has been a proposal to make use of X.224 for the Q.931 and H.245 channel as well. However, this seems not to be required by the respective recommendations and, therefore, is considered to be unnecessary overhead.

Phase A2 fails if no transport connection based on the LAN protocol can be established.

3.1.4. Phase A3: call setup using Q.931 and establishment of control channel

This phase is somewhat similar to the phase A as defined in H.320. If a gateway is involved it acts as a PBX and forwards the call from the ISDN to the LAN and vice versa. However, as the H.245 control channel also needs to be set up, some further processing is required:

- For an outgoing call from the LAN to the WAN as well as for a LAN-internal call, the calling terminal sends a Q.931 connect request to the gateway or called terminal. The calling entity selects a (for the communicating peers) unique "B channel number" to be used and passes it in the Q.931 connect request.
  Afterwards the called entity (either the called H.323 terminal or the gateway) creates a TCP/SPX connection (again, using the TPKT header of RFC1006 to preserve packet boundaries) to another predefined (i.e. well-known) transport address that is to be used as H.245 control connection. Then the called entity sends an initialization PDU across this connection. The PDU contains the following information: a protocol selector — for robustness purposes — identifying this connection to be an H.245 control connection (which also provides a hook for future extensions)[15] and the "B channel number" received in the Q.931 PDU before to identify — together with the calling and called terminal/gateway addresses — to which call the H.245 control connection belongs.
  - In case of a LAN-internal call, after reception of the initialization PDU for the H.245 control channel the called terminal responds to the call, acknowledging or rejecting it.
  - In case of an outgoing call to an ISDN (WAN) terminal, after having sent the initialization PDU for the H.245 control channel, the gateway issues a connect request on the ISDN D channel, and forwards the called terminal's response to the calling terminal.
- For an incoming call from the WAN, the gateway behaves like the calling H.323 terminal as described above. After completion of the setup with the called terminal, it confirms the connection setup to the WAN terminal using the corresponding Q.931 PDU.

Note, that in a LAN-internal call the called terminal acts as an NT towards its peer (concerning the Q.931 protocol) while acting as a TE towards its local applications.

Phase A3 fails if

- the H.245 control connection cannot be established or
- the called terminal rejects the call.

3.2. Phase B

---

[15] If an initialize PDU is also used for the LAN-D channel (identifying it as the Q.931 connection), the same transport address may be used for both types of connections (and, consequently, only a single well-known port is required).

3.2.1. Phase B0: Frame synchronization

As defined in H.320, H.221 frame synchronization is established on the ISDN side. If no H.221 synchronization can be achieved, the call shall proceed as a simple audio connection (plain-audio call).

This phase is not applicable to LAN-internal connections.

3.2.2. Phase B1: Mode initialization on initial channel

In a LAN-internal call, the two H.323 terminals exchange their capability sets using the messages defined in H.245.

Otherwise, the two end points exchange their capability sets with the gateway. The gateway is responsible for converting H.242 BAS codes into H.245 messages and vice versa as necessary.

Two types of capability negotiations need to take place when an H.323 terminal talks to a WAN terminal:

- H.323 terminal to gateway
  The H.323 terminal indicates its capabilities to the gateway and the gateway indicates the WAN terminal's capabilities to the H.323 terminal (in case of a plain-audio call, the gateway signals audio only capabilities to the H.323 terminal). The gateway may also declare capabilities for the WAN terminal that are achieved by doing format conversions in the gateway (e.g. for audio encodings). This capability exchange is based on H.245.
- Gateway to WAN terminal
  The WAN terminal indicates its capabilities to the gateway and the gateway indicates the H.323 terminal's capabilities to the WAN terminal. The gateway may also declare capabilities for the H.323 terminal that are achieved by doing format conversions in the gateway (e.g. for audio encodings). This capability exchange is based on H.242 (ISDN terminal) or H.245 (ATM or PSTN terminal).

From these negotiations the terminals can derive which media and which encoding formats they may use to communicate (taking into account that the gateway may be able to do encoding format conversions for some media types). Based on the outcome they shall proceed as if the capability exchange was performed without interference of the gateway.

In addition to the capabilities currently defined in H.245, some LAN-specific capabilities are to be defined. For H.22Z, further transport address(es) to be used for this particular B channel (and its information streams) is/are exchanged. This information is used to identify the various information streams that belong to the same (aggregated) B channel(s):

- Two transport addresses for an end point using RTP and RTCP as protocol on top of an unreliable (connectionless) transport protocol are allocated for audio information.
- Two transport addresses for an end point using RTP and RTCP as protocol on top of an unreliable (connectionless) transport protocol are allocated for video information.
- A single transport address is needed to associate a set of T.120 reliable transport connections — being carried in the MLP channel in the ISDN case — (that use a fixed, predefined TCP port/IPX transport address) with a particular LAN-ISDN connection (refer to section 3.4.3.).

Further authentication procedures might be performed at this point — before data connectivity is established.

3.3. Phase CA

3.3.1. Phase CA3: Call setup of additional channels

This phase works as described in H.320. However, for additional B channels on the ISDN side, no additional Q.931 connect request PDUs are sent on the LAN. Instead, only the bandwidth to be used is increased which is signaled via the H.245 control channel.

In the presence of a gatekeeper a change of the reservation might be required:

- If the calling terminal of a connection is an H.323 terminal, it needs — in the presence of a gatekeeper — to increase its bandwidth reservation for the connection if the bandwidth utilization is to be extended beyond the initial reservation. If the gatekeeper rejects this additional reservation, no H.245 message to increase the bandwidth is issued.
- If the calling terminal is a WAN terminal and it sets up additional channels that exceed the previously allocated reservation, the gateway needs to extend this reservation by contacting the gatekeeper. If the additional reservation is rejected, the setup of the additional connections is rejected as well.

No additional functions need to be performed on the LAN, as the address translation has already been done and the H.245 control channel channel has already been established. Also, no new audio/video/data channels are established; the existing ones are re-used (as are the transport addresses).

In parallel, on the ISDN connection the actions of phase CB1 (as defined in H.320) are performed.

3.4. Phase C: Visual telephone communication

After the establishment of all B channels, a certain multiplexing mode is chosen. This mode defines how the p*64 kbps are shared among the media to be transmitted simultaneously. The various media streams are multiplexed using the different transport addresses that have been allocated and exchanged during phase B1.

3.4.1. Transmission of Control Information

Mode switches, bandwidth allocation, and other control commands are signaled according to H.242 (ISDN) and H.245 (LAN). In the LAN, the reliable transport connection is used.

Note, that some of the H.245 commands need to be synchronized with the flow of the media streams — in H.320, for example, it is defined that some commands signaled via a certain BAS code are to take effect a fixed number of frames after the command code was sent/received. This, however, requires synchronization between the two information streams. As no fixed relationship exists (as defined by the frame structures of H.320 or H.324), some other means for providing this timing needs to be done.[16]

[16] The main difficulty is that it is not known how long the transmission of a particular H.245 command

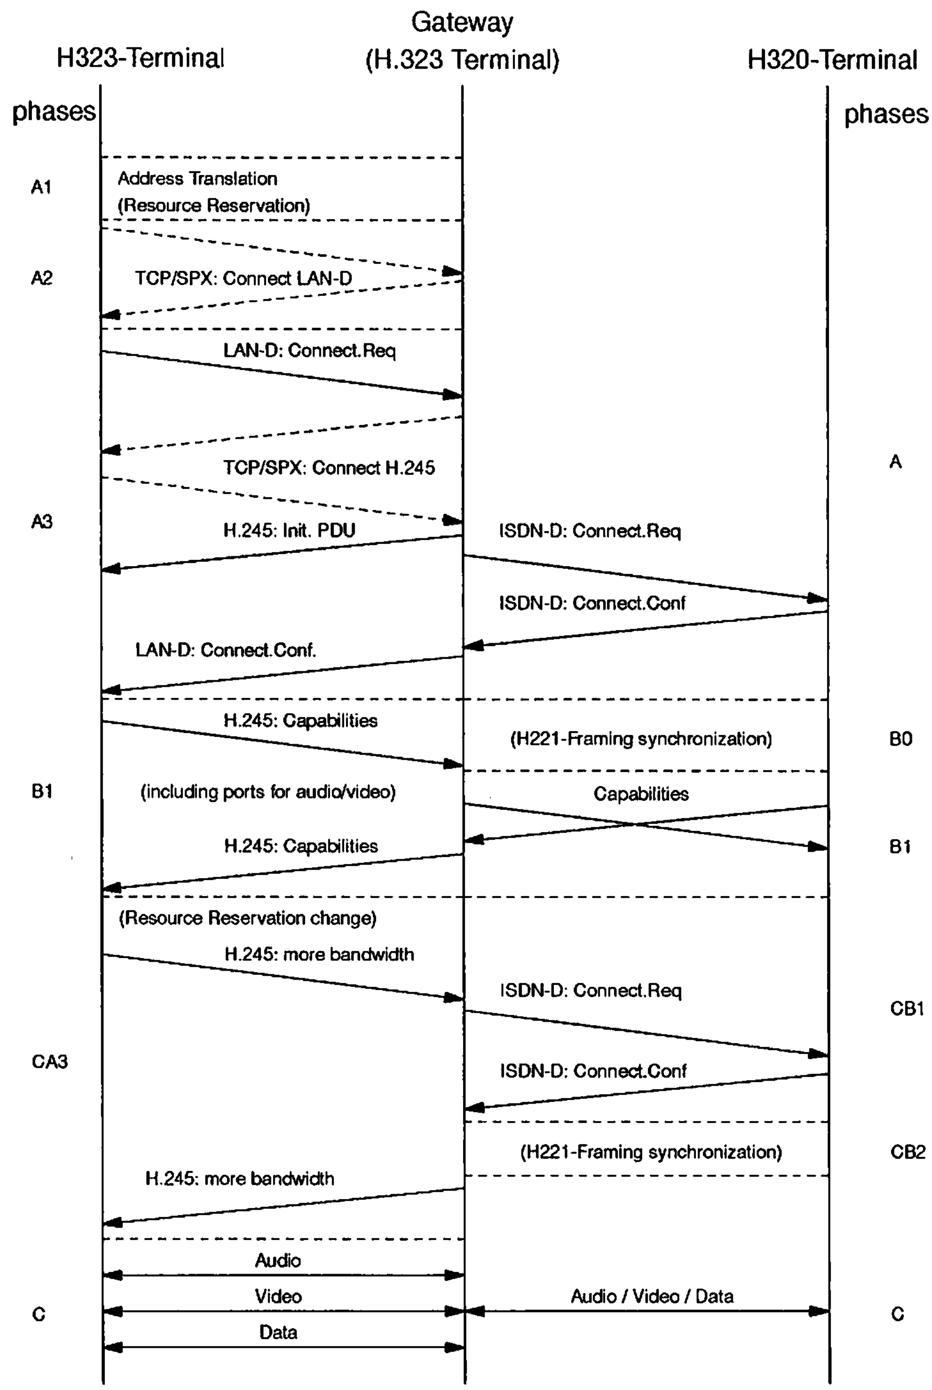

Figure 2: Connection setup initiated by H.323 terminal

over a reliable LAN connection may take: a packet may be dropped and thus needs retransmission, and the transmission delay may vary due to varying queuing delays in routers in the LAN.

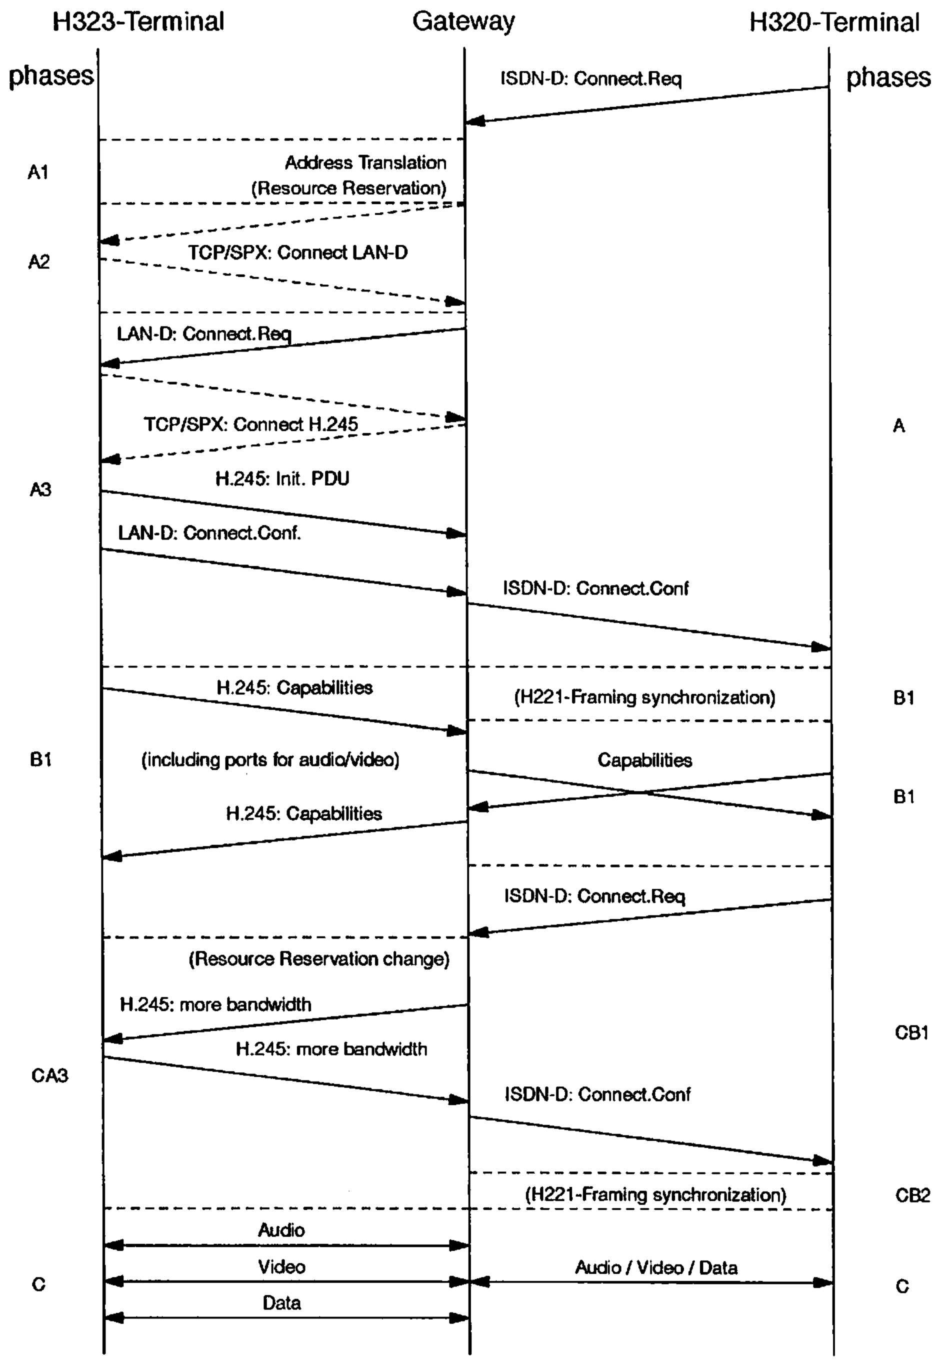

Figure 3: Connection setup initiated by H.320 terminal

Basic synchronization should be done by providing references to the RTP media time stamps of audio and.or video. However, this does not prevent a control packet from being delayed and arriving after the first media packet it should take effect on. To accommodate these cases as well, some kind of handshake is required — at the cost of some extra delay being added. Note, that RTP already provides some fields in its packet header to synchronize important information with the flow of audio and video streams, e. g. which encoding format is used for a particular packet, whether this packet is encrypted or not. Nevertheless does this issue need further study.

3.4.2. Transmission of Audio and Video Information

Audio and video information are transmitted using an unreliable, connectionless protocol.

An additional transport protocol is required for audio/video information to cater for the real-time requirements of these information streams. Such a protocol should

- provide timing information, jitter control, synchronization of different media streams, and the like, and
- define a packetization format for the various audio and video encoding formats.

This functionality will be provided by RTP/RTCP as defined in the respective Internet Proposed Standard.

It might be required that this protocol provides for clock synchronization among the LAN-based terminals and the ISDN — it is suggested that the gateway, who knows the ISDN clock, uses this protocol to synchronize the LAN terminals with the ISDN clock. For pure LAN connections this function is not needed.

3.4.3. Establishment and Association of T.120 Transport Connections

According to T.123, T.120 connections are established within the MLP channel of the H.221-multiplexed information stream. Within an ISDN connection, the T.120 connection is therefore implicitly associated with the corresponding control, audio, and video channels. The setup procedure for the MLP channel, Q.922-, Q.933-, and X.224 connections is defined in T.123.

For a LAN environment, in the annex of T.123 the usage of TCP/IP with RFC1006 on top is suggested (where TCP/IP might also be substituted by SPX/IPX). The selection of the MLP channel on the ISDN side is conveyed to the H.323 terminal via H.245. The bandwidth allocation may be refined. The X.224 connection setup on top of MLP is terminated at the gateway and a (set of) X.224 connection(s) is created in the LAN (or vice versa).

In the LAN, an MCS connection is set up on top an X.224 class 0 connection that makes use of of a reliable transport connection (TCP or SPX); the TCP setting is described in T.123. The TCP/IP transport connection is set up using a well-known port (port 1503 is assigned to T.120 by the IANA; refer to RFC1700). For SPX/IPX such a transport address needs to be allocated as well.

The X.224 transport part of the address (the OSI TSAP) is used to distinguish transport connections belonging to different LAN B channels. This may done by using the transport address of the LAN control connection (for H.245) exchanged in phase A4. This address is then used as the *calling TSAP-ID* and the *called TSAP-ID* for the T.120 LAN connection setup, respectively.

3.5. Phase D: Termination phase

Similar to H.320.

Previously allocated resources (such as transport addresses, bandwidth, etc.) are released.

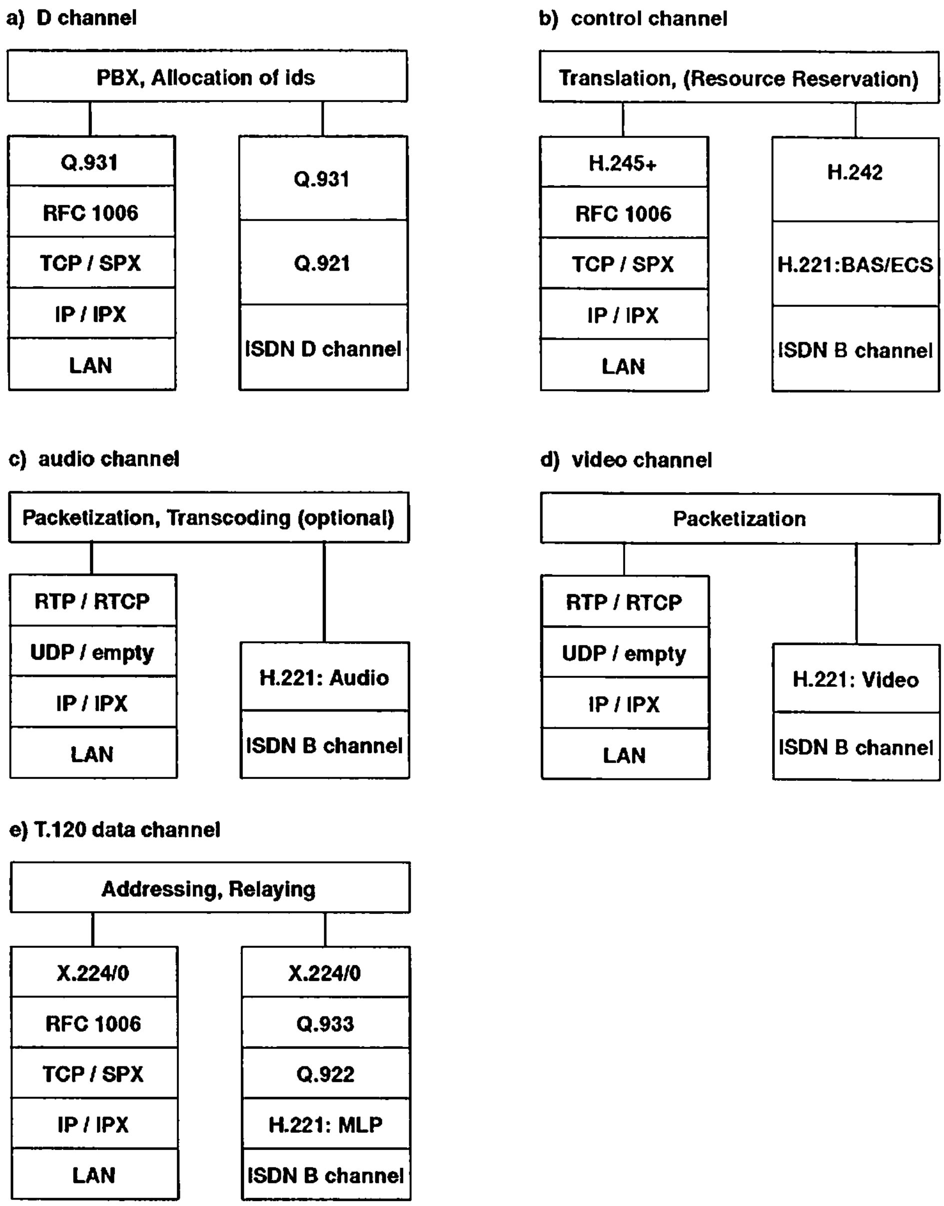

Figure 4: Adaptations of protocol stacks to be performed in the gateway

3.6. Phase E: Call release.

Similar to H.320. Recall that the LAN-D channel between two end points is not disconnected until the last "B channel" between those two entities is released.

3.7. Protocol stacks

Figure 4 shows the necessary protocol mappings and additional functionality to be realized in the gateway and the protocol layers at which these tasks are to be performed. In this figure, it is assumed that the gateway does not act as an MCU/T.120 node as well.