ITU Telecommunication Standardization Sector
STUDY GROUP 15

Document AVC-831
Original: English

Yokosuka, 24 - 27 October 1995

Question(s): 2/15

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

Title: Usability of Today's LAN Multicast Environments for ITU-T Teleconferencing
TUB/TEL-02-B

Purpose: Information / Discussion

1. Introduction

ITU-T Q2/15 is currently developing a set of recommendations for LAN-based videophone / conferencing terminals and a *gateway* that interfaces the LAN to the WAN environment. When it comes to supporting multiple nodes on a LAN interacting in a teleconference rather than a point-to-point call the use of multicasting reduces the load of the LAN. Multicasting basically can be deployed in two ways:

a) The gateway also acts as an MCU and is responsible for distributing the information streams to the LAN terminals or

b) A multipoint teleconference interconnects only terminals on the same LAN so that no MCU might be needed.

The latter case is deferred for further study for the time being; therefore, we assume the existence of a gateway/MCU whenever multicasting is to be used.[1]

The discussion during the Q2/15 meeting in Haninge whether or not to consider multicasting brought up several issues that were considered to potentially hinder the use of multicasting and need further investigation. This contribution addresses these and some further issues concerning the use of multicasting. It indicates long-term solutions to some of these issues and, in some areas, makes proposals how to deal with them in the short term.

[1] Note, that gateway and MCU might also be separate entities on the LAN, i. e. when we are talking about LAN-only teleconferences, obviously no gateway needs to be involved. We make use of the term *gateway/MCU* to indicate both possibilities.

2. Networking scenario

Figure 1 depicts a possible network scenario that is to be dealt with. There may be any number of gateway/MCU hosts that interface the LAN to one or more WANs, and there may be other routers connecting the LAN to other (inter)networks outside the enterprise this LAN belongs to. Also, many different networking protocols may be used within this LAN.

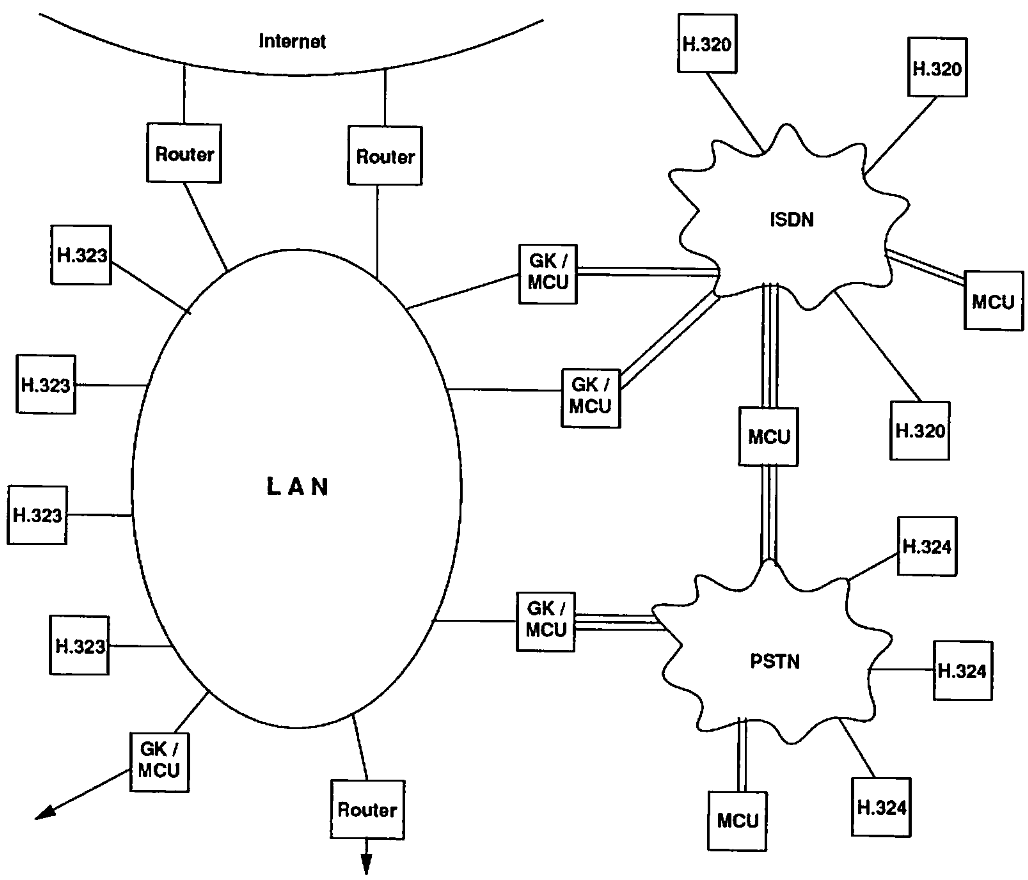

Figure 1: Possible scenario how a LAN may be interconnected

A LAN may range from single segment to large internetwork. In the latter case this internetwork may be inhomogeneous in many respects:

- Different subnetwork types may use different link types and support a variety of QoS parameters (e.g. bandwidth, transit delay, guarantees, etc.).
- The various parts of the LAN are expected to have widely differing loads.
- The routers used in the LAN are supplied by different manufacturers and are expected to provide different functionality (routing protocols, reservation, etc.).
- The LAN components (as there are networks, terminals, and routers) may or may not be capable of handling multicast traffic.

3. Basic concept

In this section, we introduce the idea how multicasting can be applied for ITU-T teleconferencing in LANs. To show the basic concept, we make several assumptions that would not be valid in all environments but simplify the description and allow us focusing on what technically needs to be achieved in a simple environment. In the subsequent sections, we describe which problems will arise if these assumptions are not valid, and, afterwards, we point out some approaches to overcome these problems.

For the moment, we assume that we are concerned only with a single LAN segment all networking hardware of which permits multicasting. Furthermore, we have only a single gateway/MCU, and the LAN is dedicated to carry only H.323 traffic.

In this simple setting, the establishment of a conference call could work as described in the following:

1. Either a LAN terminal calls into the gateway/MCU or it is called by the gateway/MCU. As this is the first terminal on the LAN, it is not yet known whether or not additional LAN terminals will join later on. Therefore, the connection is handled as a point-to-point call and no further actions are taken.
2. However, during the capability exchange phase, the LAN terminal and the gateway/MCU negotiate whether they are capable of using multicasting for the various information streams: video, data, and audio (the latter of which has been declared for further study by Q2/15).[2] This negotiation process takes place for each LAN terminal that enters the teleconference through this particular gateway/MCU. Communication so far is done based on point-to-point connections — and this does not change unless an explicit switch to the use of multicasting has been performed regardless of how many LAN terminals are joined to the conference.
3. Switching to multicast may be done for all media types simultaneously or one by one (depending on some decision rules). Switching occurs for a certain media type as soon as a certain threshold number of LAN terminals supporting (a particular mode of) multicasting is reached.[3]
4. After the decision to make use of multicasting, some dedicated entity (the gateway, the gatekeeper, or some other agent) "allocates" an unused multicast address to the particular media type of a connection to be switched. Then a setup process for the new multicast connection is invoked:[4]
   a) If the underlying protocols require so, the multicast group is explicitly set up.
   b) All parties that are to participate in the newly created multicast group are informed through some protocol mechanism at least about the group address. Further information that need to be exchanged on a peer-to-peer basis may include authentication information,

[2] It might be not before capability exchange that the two communicating entities realize that they are connected to the same LAN.

[3] If LAN terminals drop out of the teleconference some (possibly different) threshold might be used to switch back to point-to-point connections.

[4] The following description applies to a multicast service that does not provide an explicit connection setup itself (such as IP). Other protocols (such as ST-II) would require different setup mechanisms.

session keys for encryption, the timing for switching and the like.

c) It is expected that the switching process takes some time to complete. Therefore, the data transmission is continued using the point-to-point connections until all invited parties have *join*ed the multicast group.

d) Then, the transition to the multicast group is initiated. Afterwards, information transfer is carried out only via the multicast group.

5. Depending on the media type and the protocols used, the point-to-point connection may then be released.

4. Problems identified so far

While the above concept would basically be a feasible approach towards the use of multicasting in isolated areas — although it is not yet verified that really all important issues have been addressed — several problems have to be overcome when applying this concept in one of today's "real" LANs.

1. As already stated above not all of the components to be found in a LAN may be able to support multicasting:

   a) networks may not support multicast in hardware or through specific software,

   b) terminal networking hardware and/or software may not support multicasting (e.g. Ethernet boards may not have registers to program multicast addresses into, device drivers may not support programming boards that do have this feature, operating systems may not have interfaces for this at all), and

   c) routers may not implement multicast routing protocols; in particular, standardization of multicast routing protocols within the IETF is in progress (well-advanced but not yet completed).

   Due to one or more of the above reasons, a node may not be reachable by multicast although it is by unicast packets.

2. It is required to dynamically allocate multicast addresses for e. g. (connection, media type) peers. If the LAN is not isolated (or not limited to H.323 traffic), today, it cannot be guaranteed that there is no accidental overlap with other applications.

3. Using multicasting instead of unicasting requires changes — at least additions — to existing recommendations. For audio, video, and control (H.245) information, it is up to this group to specify the appropriate protocols. For T.120, it was found that this is the task of Q10/8.

4. For the transmission of data information, a reliable multicast transport protocol is needed as T.125 requires reliable transport connections below. However, currently there is no standardized multicast transport protocol available.

5. ...

5. Long term solutions

Many of the problems stated above will be overcome in the long-term, either by work done in Q2/15 or by other (standardization) bodies or enterprises.

1. The standardization process for multicast routing in the IETF will have produced stable standards in the near future. A trend can be recognized in the computer industry to support multicasting protocols: manufacturers of routers and networking hardware as well as operating system providers increasingly support multicasting in their products. For networks that do not support multicast in hardware it is expected that concepts to simulate multicast in the networking software are developed (as has been done with the multicast routing protocols for internetwork wide multicasting in the IETF).
2. A scheme or service for multicast address allocation (e. g. as envisaged in RFC 1458) could be provided. As an alternative, a more elaborate scoped multicast (as is to be developed for IPv6) could alleviate the problem by allowing enterprises to manage their own multicast address space.[5]
3. Some protocols will only be marginally affected by the use of multicast (such as H.245 to carry the respective capabilities and the switching commands and procedures), others may not be affected at all (such as the LAN-D channel as described in TUB/TEL-01). Protocols that are run over multicast "connections" are the ones that may need larger changes. Note, that RTP which is to be used for audio and video information is designed for use with multicasting. T.122/T.125 will require modifications that will be dealt with by Q10/8 in the not too distant future during the revision process of the T.120 series.[6]
4. A reliable multicast protocol — be it general purpose or T.120 specific — will eventually be defined probably within the IETF rather than within the ITU-T.
5. ...

6. Some short term solutions

The long term solutions mentioned in the previous section indicate that the usage of multicasting will be possible in the future — and thus justify considering multicasting at all. As Q2/15 needs to specify a revision of H.323/H.22Z now, in this section we will examine how the problems indicated above can be overcome *today*.

1. The fact that not all components of a LAN may be capable of doing multicast does not preclude those that are from using it. This leads to a mixed environment in a LAN where some terminals in a particular conference are connected through point-to-point connections while others are using multicast. The LAN still benefits from multicasting since this reduces the overall network load. What needs to be done is to support determining mutual reachability of a set of peers dynamically:[7]

---

[5] However, not even this can prevent (accidental) misusage of allocated addresses.
[6] Note, however, that this is currently a work item of lower priority.
[7] This might include reachability at a certain minimum QoS.

a) A terminal should know what kind of network it is attached to and whether this supports multicasting or not. If the network does, the terminal will advertise multicasting capabilities in its capability set that is exchanged during connection setup; otherwise it won't do so. Should the terminal not be able to find out about the multicasting capabilities of its network it should follow the procedure described in c).

b) The terminal should definitely know or be able to find out whether its operating system and its network hardware support multicasting. It should advertise its capabilities accordingly.

c) If both terminal and gateway/MCU have determined that they *should* be able to exchange multicast packets, they still need to find out whether or not the internetwork components in between them are capable of forwarding multicast packets. A possible mechanism for multicast reachability determination could work as follows:

   - the gateway/MCU forwards a testing multicast address to the terminal during connection setup;
   - then both gateway/MCU and terminal join this multicast address and start sending *ping* messages including some measurement information such as timestamps and sequence numbers;
   - this testing process is executed in parallel to the connection setup and may continue during the course of the conference;
   - after a (initially negotiated) period of time the testing process is completed and both parties calculate some statistics from their test;
   - the gateway/MCU uses this statistics to decide whether multicasting is basically possible and whether (and for which media) it might be useful for the QoS observed during testing;[8]
   - finally, the gateway/MCU announces to the terminal if multicasting may be used during the call.

   Note, that determining basic multicast reachability may be done between terminals and the gateway/MCU(s) permanently in the background independent from a particular conference. However, as the (multicast) network load varies and because multicast routers might become unavailable, it is suggested to have some determination phase for any connection setup — unless multicast router availability information may be retrieved from some management information base.

Finally, note that tunnels might be used even in comparably small LANs to simulate multicasting capabilities of the network and to simplify processing in the gateway/MCU.

2. Currently, there is no mechanism to dynamically allocate multicast addresses[9] and (as stated above) even if there was this would not guarantee information exchange without any

[8] QoS observation should not be an issue if a LAN-wide reservation system is available.

[9] Note that static (administrative) assignment of multicast addresses is not an issue within the IETF. For acceptable purposes — e. g. for ITU-T protocols — it should be possible to obtain a set of multicast addresses.

interference. If a multicast address is not unique this could be for two reasons:

i) Some entity outside the LAN might have created a multicast session[10] using the same address and at least one entity within the LAN participates in this session. In this case, there is no possibility to avoid collisions except blocking incoming multicast packets for multicast addresses that are used for internal gateway/MCU to terminal communication — such an approach to *administratively scoped multicast* has been presented on the 30th IETF in the working group for *Inter Domain Multicast Routing (IDMR)* by Van Jacobson and Steve Deering. However, to do so the gateway(s) would have to inform all multicast routers which groups are not permitted (alternatively, one could assign a block of multicast addresses for internal usage and statically configure the multicast routers not to pass through packets destined for one of these addresses.

ii) An entity within the LAN has created a (local) multicast session. In this case, each entity could be required to query some LAN internal service for address allocation to prevent overlapping with the addresses used by the gateway/MCU(s). However, there may be other address assignment tools in use that need to be considered.[11]

A suitable approach for the ITU-T might be to request a block of multicast addresses for use by H.323 with one or two static addresses being dedicated to gateway/MCU and terminal communication (e. g. one *all-gateway/MCUs* and one *all H.323 terminals* (including the gateway/MCUs?) address. The remaining multicast address space should be used for LAN-internal multicast communication within H.323 conferencing calls. These addresses should be assigned dynamically by some LAN authority, and packets using such addresses as destination should never leave a particular LAN. The address block to be allocated for exclusive use by H.323 terminals and gateways would be expected in the order of 10^5 to be able to even work with companies that have large LANs and a single or a pair of addresses is used for each teleconference call.[12]

Nevertheless, address conflicts still may occur — e. g. due to non-knowledgeable or even malicious users or programs. Therefore, it must be guaranteed that the terminals can also operate in environments in which they receive multicast packets originated from different applications. The single mechanism needed is a way to unambiguously distinguish H.22Z PDUs from any other protocol.[13] Obviously, this can only partially resolved within the ITU-T: the design of the packet header for H.22Z should provide a clear means to identify each packet as being an H.22Z PDU and differentiate it from any other protocol *currently* in use. It is up to other standardization bodies to avoid such ambiguities for protocols that will be specified in the future.[14]

---

[10] The term *multicast session* refers to any usage of a multicast address by a set of application entities.

[11] Conflicts with *sd*, the Internet session directory, can simply be avoided because a block of multicast addresses has been allocated for exclusive use by *sd*.

[12] The total number of available multicast addresses is 2^{28} equivalent to approx. $2.7 \cdot 10^8$.

[13] If authentication is used along with encrypting all packets with a session key, the failure of the decryption indicates automatically that a misaddressed packet was received (that will then be dropped). Thus, the use of authentication solves the problem of identifying the packets belonging to the session — at a greater cost in terms of processing required, however.

[14] Note, that the RTP specification provides an algorithm to check incoming packets for validity based on many of the fields in the packet header.

3. The protocols for connection setup, any additional H.245 messages and procedures and the data protocols can be dealt with by this group when defining so that the short term solution is basically equal to the long term solution. However, due to time constraints some problems may be excluded from the first revision of the Q2/15 specifications — such as (T.120) data communication (see item 4.).
4. For the short term, T.120 implementations that make use of multicasting can only be developed based upon non-standard protocols as no standards for reliable multicast transports are available. The hooks that can be provided are certain non-standard capabilities that allow terminals to learn about compatible peers. Interoperability between multicasting T.120 implementations cannot be achieved entirely within the ITU-T but should be based on agreements in industry consortia.
5. ...

7. Conclusions

From this first analysis, multicasting seems applicable even in today's "unregulated" LANs. Possible mixtures of unicast and multicast system in the same LAN and the same teleconference session are a prerequisite for early and incremental deployment of this technology.

Concerning the control protocols to be used, they should provide hooks for negotiating multipoint capabilities for all information streams and allow for any kinds of scenarios. Signaling for switching between different protocols or protocol modes need to be provided (either as in-band signaling for each information stream or as out-band signaling within H.22Z).

Static allocation of a set of multicast addresses seems to be an important step towards making the internal use of multicasting manageable.