Telecommunication Standardization Sector
Study Group 15
Experts Group for ATM Video Coding
(Rapporteur's Group on Part of Q.2/15)

Document AVC-513
July 1993

SOURCE: Stephen Hall (Australian UVC Consortium)

TITLE: H.22z protocol for LANs

PURPOSE: Proposal

Abstract

This document describes an H.22z protocol for transporting real-time audiovisual/multimedia information over packet switched local area computer networks. The required functions are identified, and a corresponding packet structure is proposed.

1. Introduction

Document AVC-512 discusses the basic design of an H.32z terminal, which is intended for connection to packet switched local area networks (LANs). The present document describes an associated H.22z protocol, which provides the terminal with a means of conveying multimedia information across the LAN. The H.22z protocol assumes that the LAN is able to transport packets between two or more network nodes within suitable delay bounds, but with no guarantee of delivery (ie. using an "unacknowledged datagram" service). Such a service can be offered on suitably configured Ethernet, Token Ring, FDDI and ATM LANs, which incorporate basic resource allocation mechanisms in network switches and routers. Examples of unacknowledged datagram protocols include IEEE 802.2 LLC Type 1, and Internet IP. The H.22z protocol adds functions to enhance the basic datagram service as appropriate. For example, it implements an automatic retransmission scheme to correct errors in end-to-network signalling streams. A layered protocol model incorporating H.22z is shown in Fig. 1.

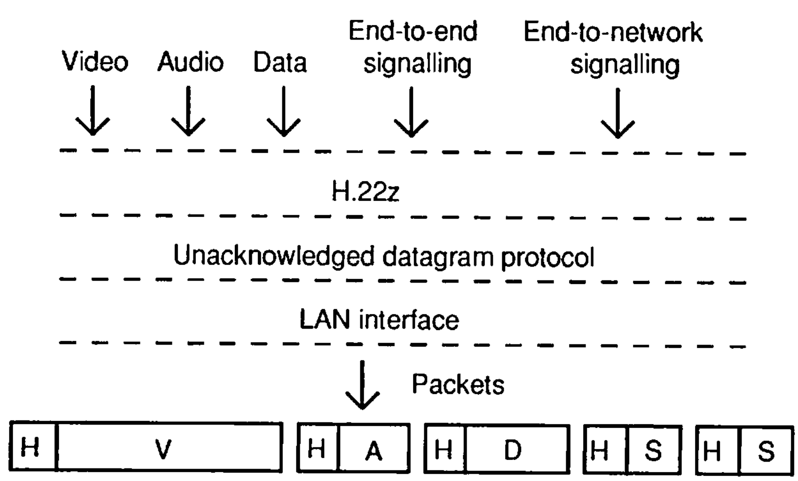

Fig. 1 Layered protocol model incorporating H.22z

2. Functions of H.22z

- Packetisation: The H.22z protocol must segment the various data streams into units which are suitably sized for transport on the LAN. The length of each LAN packet can be matched to the bit rate of the associated data stream, in order to minimise the packetisation delay. Where a natural data unit, or "message", is present in the stream, such as a video GOB or slice, then it may be advantageous to align the packets with these data units. For example, the effects of packet loss on video quality can be reduced by aligning GOBs/slices to packets. However, since a long message may span more than one packet, an indication of the last packet in the message is required, in order to minimise the reassembly delay. This can be provided by a "segment type" field in the H.22z packet header, which indicates that the packet contains the beginning of a message (BOM), the continuation of a message (COM), the end of a message (EOM), or a single segment message (SSM). In addition, it is possible that messages will not be aligned to *byte* boundaries in the packet payload. An example is where the payload is created by segmenting a continuous bit stream from a video codec. The provision of pointers in the H.22z packet header to the first and last bits in the payload avoids the need for time-consuming bit-shifting operations prior to packetisation.

- Multiplexing: As described in document AVC-512, the H.22z protocol is required to provide a multiplexing function for the various data streams from a single terminal. Furthermore, a packet multiplex approach, where each packet contains only one type of data, is shown in AVC-512 to be preferable to one based on the encapsulation of H.221 frames. This leads to the requirement for a "stream identification" field in the H.22z packet header. The field should be sufficiently large to allow the identification of streams from multiple terminals, in order to accommodate future multipoint configurations.

- Error handling: Packets may be lost on LANs due to bit errors or queue overflows. Bit error rates on LANs are very low, typically 10^{-9}, and queue overflows can be avoided through proper network design and management. Nevertheless, the losses which do occur must be detected by the H.22z protocol, and this can be done using a sequence number in the H.22z packet header. In addition, appropriate action must be taken once an error has been detected. In the case of lost audio or video information, the terminal must be informed of the occurrence and location of the loss, in order to allow it to employ its own error concealment or recovery techniques (eg. a "fast update request"). In contrast, lost end-to-network signalling information can simply be retransmitted. However, special measures must be taken to avoid the loss of end-to-end signalling and C&I information (eg. BAS), since this may be delay-sensitive. For example, packets containing such information can be duplicated prior to transmission.

- Timing recovery: Most LANs operate asynchronously, implying that there is no fixed timing relationship between successively delivered packets. Furthermore, there is often no master network clock which can be used as a timing reference by the terminals. The H.22z protocol must therefore provide a means of recovering the timing between packets, and synchronising the terminals. This can be done using a combination of buffering and sliding window flow control, leading to the requirement for a numbered acknowledgement field in the H.22z packet header.

3. Packet structure

A structure for H.22z packets, designed according to the above requirements, is shown in Fig. 2. Note that the H.22z packet header is only 3 bytes long, while the payload may be several hundred bytes in length, so that the packetisation overhead is very low.

SID	SEQ	ACK	ST	PF	PL	PAYLOAD

SID = Stream identifier (8 bits)

SEQ = Segment sequence no. (4 bits)

ACK = Acknowledgement no. (4 bits)

ST = Segment type = BOM, COM, EOM or SSM (2 bits)

PF = Pointer to first bit in payload (3 bits)

PL = Pointer to last bit in payload (3 bits)

Fig. 2 H.22z packet structure

4. Conclusion

A simple protocol for transporting multimedia information from H.32z terminals across packet switched LANs has been proposed. Further work is required to specify its operation in detail.