Assessment of Traffic Prioritization in Switched Local Area Networks Carrying Multimedia Traffic

F. A. Tobagi, C. J. Fraleigh, M. J. Karam, W. Noureddine

Computer Systems Laboratory

Department of Electrical Engineering

Stanford University

Stanford, CA 94305

1.0 Introduction

Multimedia applications are those which involve multiple media types including audio, video, and still images, in addition to the more conventional data types. Many such applications are distributed in nature and involve networking and communications. Examples of such applications are: video-on-demand (VoD) for training, education and entertainment; audio/video broadcasting for news and information dissemination, and entertainment; audio/video conferencing for communication among groups of people; and computer supported collaborative work. There is a great interest in supporting these distributed multimedia applications over existing network infrastructures (which have originally been designed to support conventional data applications, such as data base access, electronic mail, file transfer, etc.), since changing to new technologies would be costly and would most likely delay the introduction of these applications.

To satisfy the requirements of multimedia applications, IEEE 802.1p introduces two new functionalities into bridges: (i) the provision of expedited traffic capabilities to support the transmission of time-critical information in a LAN environment; and (ii) the provision of filtering services that support the dynamic use of Group MAC addresses in a LAN environment [1].

This report evaluates the use of expedited traffic classes to prioritize multimedia and other time-critical data over non time-critical traffic. The evaluation of GARP/GMRP introduced to provide dynamic multicast filtering within Local Area Networks (LANs) is the subject of a separate report [2].

2.0 Multimedia Application Requirements

The characteristics of multimedia applications and their underlying traffic differ substantially from the more traditional data applications. These differences pertain to traffic patterns, data rates, delay and reliability requirements, modes of communication, and the need to integrate services for all media types in the same network. In turn these differences in traffic characteristics place new requirements on the networks that are to support multimedia applications.

2.1 Streaming Requirement

Traditional data applications produce traffic that is bursty. The time at which a demand may be placed by a source on the network is unknown, and the volume of data associated with that demand is random and may be highly variable. Furthermore, there is no preconceived idea of a required data rate associated with a source or a session; although it is known that the average data rate associated with a data source is typically on the order of tens, at most hundreds, of Kbps; (indeed, it is very common to see tens, and even hundreds, of data application users share the same 10 Mbps Ethernet LAN segment.) Accordingly, the design of networks aimed at supporting data applications has been based on the principle of bandwidth sharing and statistical time multiplexing.

In contrast with data traffic, audio/video traffic is stream-oriented; i.e., it consists of a more-or-less continuous flow of information of a given data rate, and thus requires an amount of network bandwidth equal to the underlying rate to be available on a somewhat continuous basis. By the same token, in order to support multiple concurrent audio/video streams, the network must have the necessary aggregate bandwidth required by all the streams.

The data rate associated with a stream can vary considerably depending on the type of media (audio or video), the compression scheme used, and the desired quality level. For digitized voice, for example, using Pulse Code Modulation (PCM) the data rate is 64 Kbps; with differential encoding, lower data rates are achievable. For video traffic, the data rates can range from tens of Kbps to tens of Mbps.

2.2 End-to-end Latency Requirements

Traditional data applications place no specific delay requirement on the network. Although a fast response time is always desirable, delay is merely an inconvenience; it results in a slow-down in the execution of the applications, but otherwise is harmless. Multimedia applications, on the other hand, introduce a maximum end-to-end delay requirement. End-to-end delay is defined as the time from when the information (e.g., audio sample, video macroblock) is generated at one end until the time it is available at the other end. Depending on the degree of interactiveness of the application, the maximum delay requirement may be anywhere from 100 ms (as is the case with audio/video conferencing) to upwards of one second (as may be the case with some VoD applications and one-way broadcasting). For example, with audio/video conferencing, to achieve interactive communication similar to that possible with the telephone network, it is essential that the end-to-end delay for both the audio and video signals be on the order of 100 to 150 ms; this delay comprises encoding and decoding delays, packet formation time, and delay incurred in the network in the form of queueing, medium access time, and packet transmission time. As conferencing most typically involves users that are geographically widely separated, the network in question comprises local area networks (one at each of the conferencing sites), and a wide area backbone connecting them; as a result the delay budget for one given local area network may be as low as say 20 ms. With VoD on the other hand, a delay of 500 ms to 1 s from the time when a user presses the playback key until the display of video at the user's computer begins is quite acceptable, especially if the duration of the video is not so short. Similarly, with video broadcasting, a delay of 500 ms to 1 s is, in most cases, also acceptable. In some cases, however, a much lower delay is required; this is the case, for example, with the broadcasting of news on a foreign exchange trading floor where a delay of a small fraction of 1 s may make a considerable difference in the outcome of a trade. In summary, when considering a single extended LAN, three delay ranges are of interest: (i) on the order of 15 to 20 ms, (ii) on the order of 100 to 150 ms, and (iii) on the order of 500 ms to 1 s.

2.3 Reliability requirements

Traditional data applications always require full reliability. Given that these applications have no explicit latency requirement, full reliability is achieved at the transport layer by means of error detection and retransmissions. With audio/video based applications, on the other hand, a certain amount of audio/video information loss is acceptable. This loss may be due to three factors: (i) buffer overflow at the sending station, in the switches and routers, and at the receiving station; (ii) loss at the MAC layer in the form of excessive collisions; and (iii) loss due to end-to-end delay exceeding the maximum latency requirement; (audio/video information that is received at the destination past the maximum allowed delay cannot be played-back/displayed; i.e., is considered lost). Information loss at the destination leads to quality degradation in the played-back audio and displayed video in the form of discontinuities referred to as glitches. The amount of information loss that can be tolerated is dependent on the application and the media involved (as well as the compression scheme and compression rate). In the case of video, for example, one glitch every minute may appear excessive, while one glitch every ten minutes may be acceptable, especially in video conferencing.

2.4 Multipoint Communications Requirements

Traditionally in data applications, communications has been primarily point-to-point; multicast traffic represented a relatively low volume, and was dealt with either by sending one copy to each destination individually, or by broadcasting the data to all hosts, and let the hosts filter the data as appropriate. (In LANs, all multicast traffic is just treated as broadcast traffic.) Many multimedia applications, on the other hand, involve multipoint communications. These may be point-to-multipoint as is the case with TV broadcasting, or multipoint-to-multipoint as is the case with conferencing and computer-supported collaborative work. Given that the data rates involved in multimedia applications can be considerably higher than those seen in data applications, the methods currently used for data traffic are highly inefficient and thus inadequate. In particular, broadcasting all multicast traffic would place a limit on the total volume of such traffic that low enough to be accommodated by the slowest ling in the topology and leave sufficient bandwidth to carry other traffic. Take for example a switching hub with 124 10Mb/s ports and an internal backplane bus capacity of 1.24 Gb/s. Assume that the multicast traffic consists of 1.5 Mb/s video channels. If one were to broadcast all multicast traffic on all segments, only 6 video channels would be sufficient to use up the 10 Mb/s bandwidth on all segments leaving no bandwidth for other applications, while the backplane bus capacity can handle over 800 such channels, and generally speaking, not all users would be interested in exactly the same channels at the same time. If one were to selectively forward on each port only the multicast traffic that is desired by the users attached to that port, then it would be possible to handle a much larger volume of multicast traffic and give the users the flexibility of choosing dynamically which multicast traffic to receive. This requires a protocol which allows the users to express dynamically their interest in receiving specific multicast traffic. GARP/GMRP described in IEEE 802.1 is such a protocol.

2.5 Integrated Services Requirement

All data applications currently supported by networks produce traffic of similar characteristics, and thus no differentiation among these applications within the network needs to be made. Multimedia applications, on the other hand, by definition involve different information types with different traffic characteristics and requirements. This implies that a network which is to support multimedia applications must be able to differentiate between the different types of traffic and provide the appropriate services required by each. For example, a network supporting a multimedia application which involves stream-oriented audio/video traffic with real time constraints, and data applications traffic which is bursty and which has no real-time constraint, must be capable of giving the audio/video traffic the bandwidth it needs, and throttle the bursty data traffic accordingly. Similarly, when serving both audio and video streams simultaneously, it is important to favor audio traffic over video traffic, given that the human ear is more sensitive to distortion in audio than the human eye is to degradation in image quality. Finally, given that different multimedia traffic are subject to different end-to-end delay constraints, it is important to favor traffic with lower delay budget over traffic with higher delay budget.

3.0 Expedited Traffic Classes

One method to reduce information loss (due to buffer overflow, MAC, and end-to-end delay) experienced by multimedia traffic is through the use of multiple priority classes within the network. Assigning higher priority to multimedia traffic throughout the network will prevent latency tolerant data traffic from delaying time critical multimedia streams. Priority classes will also improve performance in systems whose peak data rate exceeds network capacity. When there are large bursts of data traffic, there will be severe congestion on the network. The end-to-end delay experienced by the multimedia streams will increase resulting in poor performance. If the multimedia traffic is given priority, it will not be affected by these large bursts of data. This will be true as long as the average bandwidth requirements do not exceed the capacity of the network. Once the network capacity becomes insufficient to handle the average requirements, the only solutions are to increase capacity or impose restrictions on the transmitting stations.

In order to implement priority classes in a network, a method of identifying the various priority levels is necessary. Some Media Access Control (MAC) protocols inherently support priority traffic. These include the IEEE 802.5 Token Ring protocol and FDDI. However, other MAC protocols, most notably the IEEE 802.3 CSMA-CD protocol, do not support multiple priority classes. The priority information must be encoded in some type of header information when using these protocols.

Once the priority information has been encoded in network traffic, a procedure for scheduling the priority traffic is required. There are many different options. The simplest is always transmitting high priority traffic first. This guarantees the high priority traffic will always have good performance. However, it is possible that lower priority traffic will never be serviced. To prevent this, a variety of fair queuing policies are also available. Starvation of lower priority traffic is also relieved by means of admission control exercised over higher priority traffic.

The IEEE 802.1p standard specifies a method for classifying traffic on bridged LANs into up to 8 different priority classes. To identify the classes of network traffic, 802.1p uses the inherent priority signalling mechanisms of FDDI and Token Ring, or the Tag Header field of the MAC frame when using Ethernet or other MAC types. It also requires that bridges implement a basic high priority first queuing mechanism. However, bridges are allowed to implement other queuing policies if they desire. This study assumes all devices on the network only use the high priority first policy.

4.0 Objectives

The purpose of the work presented here is to investigate the merit of adding prioritization of traffic at the MAC layer in multimedia networks as proposed in the IEEE 802.1p standard. The value of this new functionality is assessed by means of simulation of specifically selected scenarios of network topology and traffic patterns.

We chose to study these issues with Ethernet as the MAC technology, for several reasons. First, Ethernet technology is the most widely spread and used local area network technology. Second, Ethernet exhibits many unique and interesting characteristics; namely, highly variable delays as well as MAC loss due to contention and capture effects. We believe that these effects are worth careful study in the context of multimedia networking because of their strong impact on performance. Third, Ethernet does not intrinsically support priorities while other MAC technologies like Token Ring and Token Bus do. This means that Ethernet networks are among those that most adversely affect multimedia traffic, and may require special precautions to experience the benefit of prioritization.

However, the results obtained and the recommendations we make are equally applicable to other MAC schemes, provided that one takes into consideration the relevant issues and evaluates the effect of differences and common characteristics. For example, Token Ring networks implement a fair servicing discipline based on circulating a token around a ring of stations, each active station being able to transmit for a maximum amount of time (the token holding time, THT). This means that there is no contention for the transmission medium as in the Ethernet case, and loss at the MAC level is virtually inexistent. Also, the delay experienced by a station between two transmissions depends on the number of active stations over the ring and therefore should not be highly variable in practical situations and is bounded in any case. In addition, a traffic prioritization scheme exists and may be used to distinguish between time dependent traffic and other traffic types.

We try to depict scenarios that allow us to understand the interaction of video and data traffic inside stations and on Ethernet segments, and eventually to quantify the improvement gained by implementing prioritization as compared to unprioritized servicing of traffic. We also seek to identify potential sources of problems, and bottlenecks. Thereafter, we make recommendations as to the deployment and use of such networks.

5.0 Tools and Methodology

To evaluate the effects of priority classes and GARP, we use a network simulation developed at Stanford University by the Multimedia Networking Group. The simulation models accurately 10Base-T segments (the CSMA-CD scheme described in IEEE 802.3) as well as switches interconnecting such LAN segments. Switches are considered to be nonblocking and to have negligible processing delay, thus no input buffers are present at the ports. Output buffers at the ports, however, are modeled precisely, and the size of these buffers can be varied in the simulation. The simulator also models user stations which may generate either video traffic (video station), or data traffic (data station), or a mixture of both (multimedia station). The traffic generated at a station gets queued in the station's output buffer. If no prioritization is implemented in the station, then all traffic generated in the station gets queued into a single output buffer and submitted to the MAC layer according to a first-come-first-served service discipline. If prioritization is implemented at the station, then two separate output buffers are considered to exist, one for high priority traffic and the other for low priority traffic. The service discipline is first-come-first-serve within each queue, with the high priority queue being given priority over the low priority queue. The same considerations hold for the output buffers at the ports of a switch. In this study, when prioritization is assumed to exist, audio/video traffic is given priority over data traffic.

Video traffic is obtained from encoding a one minute video sequence of a popular motion picture (namely, Star Trek) using H.261. The sequence selected exhibits a variety of scenes, including both fast action scenes and slower dialogue scenes, thus constituting a good representation of most types of video traffic. The encoding is performed at 30 frames per second using CIF (352x240) resolution and a constant bit rate (CBR) encoding control scheme, producing a video stream with a data rate of 1.5 Mbps. The packetization of the produced video data for transmission over the network is done according to a variable size and rate packetization scheme (VSRP) whereby the amount of data that forms a packet is controlled by a maximum packet formation time referred to as T_f. Whenever a packet is formed, a timer with an initial value equal to T_f is started. From that point on, data is collected into a packet until either the maximum Ethernet packet size has been reached (1500 bytes) or T_f has expired (in which case a packet of size smaller than 1500 bytes is formed).

Data traffic is generated at a station according to the following model. Constant size bursts of M_s bytes each are generated randomly at intervals that are uniformly distributed, with the average interval between consecutive messages chosen so that the average load generated by the station is G_s. Messages that are larger than the maximum size packet are divided into maximum size packets which are placed in the queue all at once. The choice of a data traffic model as opposed to real traces allows us to easily experiment with different values of the parameters G_s and M_s and understand the effect of loading and burstiness on real-time traffic.

We focus on the video traffic and measure the rate of information loss in the network which may be due to three factors: (i) loss at the buffers, (ii) loss at the MAC layer, and finally (iii) loss due to the end-to-end delay exceeding the maximum tolerated by the application. Loss at a buffer occurs when a packet arrives to the buffer and finds the buffer full, and is function of the buffer size; there may be many buffers between the source and the destination where buffer loss may occur, the buffers at the sending and receiving stations as well as output buffers at the ports of switches through which the video packets flow. Loss at the MAC layer occurs when the number of collisions exceeds the maximum allowed in IEEE 802.3 (16 collisions) due to contention with other stations on the same segment; this occurs particularly when some stations have long queues of packets waiting to be transmitted thus capturing the medium.

End-to-end delay includes: (i) delay incurred at the encoder, (ii) delay incurred in the formation of the packets at the sending station, (iii) delay incurred in the network in the form of queueing delays in buffers, medium access delays, transmission and propagation, (iv) delay incurred in the playback buffer at the destination, and finally (v) delay incurred in the decoding and display process at the destination. In a highly streamlined system in which the steps of video digitization and encoding at the sending station are pipelined, and similarly the steps of decoding and display at the destination station are also pipelined, the sum of delays incurred in these steps is fairly low (on the order of 5 ms) and can thus be ignored. The dominant components in the end-to-end delay are then the packet formation delay and the network delay, the sum of which has to be smaller than the maximum delay requirement. When the systems are not streamlined, the digitization, encoding, decoding and display delays have to also be taken into account; since these delays are additive, it should be straightforward to apply the results shown here to such systems by considering the maximum delay requirement to be decreased by the sum of these delays.

The range of packet loss of interest must be determined in terms of the visual effects that viewers perceive due to loss of information. The following analysis intends to give an estimate on the range we need to consider. We first note that, due to the dependence that exists in the video encoded bit stream, the loss of information contained in a single packet has impact that extends beyond just that information; conversely, given that packet loss in a network is often bursty, the loss of several packets may account for the same perceived degradation as that due to the loss of a subset of these packets. Accordingly, it is important to consider the quality of degradation in relation to glitch statistics. From observation of simulation results, the number of packets lost contributing to a glitch is in most cases between one and ten packets and thus a rate of one glitch per minute translates to between one and ten packets lost per minute. For the 1.5 Mbps video stream divided in to packets of 1500 bytes (using $T_f = 8$ ms), we have a rate of 125 packets per second, or 7500 packets per minute. Thus the packet loss rate corresponding to one glitch per minute is between $1.3x10^{-4}$ and $1.3x10^{-3}$ and, equivalently, the packet loss rate corresponding to one glitch per 10 minutes is between $1.3x10^{-5}$ and $1.3x10^{-4}$. Assuming that one glitch per minute is the maximum acceptable glitch rate, the packet loss rate of interest should not be any higher than 10^{-3}. In the simulation results shown below, we focus on packet loss rates in the range of $(10^{-6}, 10^{-3})$. We guarantee the accuracy of the results obtained by simulation, by making use of the method of batch means to reach good confidence intervals around the quantities being estimated.

Note that although we do not consider data losses in the performance analysis that follows, we understand that excessive data losses or delay are unacceptable in a user network. Therefore, in a network that implements video traffic prioritization, a certain form of admission control should be included in order to prevent video traffic from monopolizing the offered bandwidth. Since admission control is an issue to be dealt with by upper layer protocols, it is outside of the scope of this study and we assume hereafter that proper measures are taken to correct this problem.

6.0 Scenarios and Numerical results

Figure 1 illustrates a typical network carrying multimedia traffic. At the lowest level it consists of a set of hosts transmitting data. These hosts fall into one of two categories. They are either traditional stations which only transmit data traffic, or they are multimedia stations which transmit a combination of video and data traffic. The multimedia stations could correspond to a host participating in a point-to-point videoconference, a video server transmitting a video stream to a set of users, or any other multimedia application.

Two basic factors contribute to performance degradation for the video traffic. The first factor is interaction between video traffic and data traffic in shared buffers at the host and in the network switches. In a shared buffer, a video packet must wait until all packets queued in front of it are serviced; when data traffic is bursty, such delays may be excessive leading to video packets exceeding their end-to-end delay requirements. With the prioritization of video traffic over data traffic, such delays are eliminated. The second factor is the contention between video traffic and data traffic from different devices when accessing network segments. In the case that a station has a long burst of data packets and is contending on the same Ethernet segment with a station attempting to transmit a video packet, channel capture effects could occur causing the video packet to incur excessive delays. We note that this effect is not remedied by having prioritization within each of the devices; the only remedy here is to avoid the possibility of contention altogether; this can only be achieved by avoiding the sharing of Ethernet segments by many hosts (and thus have a single station connected to a switch port), and by having full duplex Ethernet links connecting the host to the switch port. We finally note that, since delays are additive, the number of hops in the path from source to destination becomes also a factor in performance degradation.

This study is to numerically assess the importance of these factors and to give a basis for the recommendations made for the deployment of networks supporting multimedia applications.

6.1 Buffer sharing

In order to assess the importance of buffer sharing in performance degradation we consider some relatively simple scenarios. The first is a host generating both a video stream and bursty data traffic to be communicated to a receiver over a dedicated 10 Mb/s Ethernet link. (See Figure 2.) This scenario essentially also represents the first hop in a path between two hosts in the absence of Ethernet contention (e.g., using a full duplex dedicated Ethernet connection between the sending host and the switch). We consider that the sending station generates and transmits a single video stream of bandwidth 1.5 Mb/s as well as bursty data traffic defined by the usual parameters G_s and M_s. We consider first the case where both types of traffic share the same buffer in the sending host, and examine separately the video packet loss ratio due to loss in the buffer and due to end-to-end delay, as a function of the buffer size in the station. The sum of these two loss ratios is the total packet loss ratio. Figure 3 corresponds to G_s=3 Mb/s and M_s=6000 bytes, and various end-to-end delay values ranging from 15 ms to 30 ms; higher end-to-end delay values were not considered in this figure because the loss ratio corresponding to these delays are extremely small and below the limit on loss ratio we deemed useful to consider. We note that buffer loss decreases with increasing buffer size, and that conversely, end-to-end loss increases with increasing buffer size (since more data packets may be queued in front of video packets) to reach a plateau that remains constant for higher buffer sizes. Although the total loss ratio exhibits a minimum at some optimum value for the buffer size, this minimum is not much smaller than the plateau reached for each of the end-to-end delay requirements, and thus is of little interest. This figure actually indicates that there is a minimum size buffer that is needed in order to achieve a small enough buffer loss ratio; and that beyond that minimum buffer size, video packet loss is dominated by end-to-end delay constraints. Depending on the tolerable loss ratio as required by the application, there is a minimum delay constraint that is achievable. For the example displayed in Figure 3, with a tolerable loss ratio of 10^{-3}, a minimum delay constraint of 20 ms is achievable, while with a tolerable loss ratio of 10^{-5}, the minimum delay constraint that is achievable is 30 ms. To understand the effect of data traffic load and burstiness, we show in Figure 4 similar plots for G_s=1, 3, and 6 Mbps, and M_s= 1500, 6000, and 12000 bytes. The packet loss ratio increases with increasing values of G_s and M_s, and it is apparent from the numerical results obtained that the loss ratio is more sensitive to M_s than to G_s. If prioritization of video traffic over data traffic is in place, the video packet loss rate drops to zero (for the single video stream considered in this scenario). Note that the results shown here are also applicable for the sharing of buffers in a switch at any single hop in a multihop path between a source and a destination (as long as the traffic characteristics used in this scenario are also applicable to such a hop).

We consider now the scenario depicted in Figure 5. Two multimedia stations are connected to a switch, each having its own dedicated port, and are transmitting both video and data traffic to receivers connected to a third port. In the absence of prioritization in the output port of the switch, video and data traffic of the two sending multimedia stations share the buffer at the output port to which the receivers are attached. With respect to prioritization of traffic, there are several cases to consider: (i) no prioritization anywhere, (ii) prioritization at the sending stations but not in the switch, (iii) prioritization in the switch but not at the sending stations, and (iv) prioritization both at the sending stations and in the switch. We show in Figures 6, 7, and 8 the buffer loss ratio and end-to-end delay loss ratio as a function of the buffer size (considered to be the same for stations and switch output port) for G_s=1.5 mb/s and M_s=6000 bytes, and for the cases (i), (iii) and (iv) respectively. From these figures, it is clear that the lack of prioritization either at the stations or in the switch have a very strong degradation effect, rendering ineffective the presence of prioritization at either of the two places; and that with prioritization anywhere a buffer is to be shared leads to virtually zero packet loss.

6.2 Contention on the Ethernet

To understand the effect of contention on the Ethernet we consider the simple scenario consisting of two multimedia stations connected to the same Ethernet segment communicating both video and data traffic with each other. (See Figure 9.) Two cases are considered: (i) no prioritization in either stations, and (ii) prioritization in both stations. In Figures 10 and 11, we plot the buffer loss ratio and end-to-end delay loss ratio as a function of the buffer size for cases (i) and (ii) respectively. It is clear from these results that prioritization in the stations in totally ineffective; indeed prioritization within a station does not shield it from the bursty traffic that the other station contends on the Ethernet with. Furthermore, the minimum end-to-end delay achievable for a given desirable loss ratio is much larger here than it was the case with buffer sharing. For example, for a tolerable loss ratio of 10^{-3}, the minimum achievable delay is 50 ms, and for a tolerable loss ratio of 10^{-5}, the achievable delay is on the order of 120 ms. This is explained by the capture effects resulting from when stations with bursty traffic contends on the Ethernet. (We note that the results shown in Figures 10 and 11 are for a burstiness of only 6000 bytes, and that the degradation is a lot worse for larger burst sizes.) This simple scenario represents what happens when a station and its dedicated port are connected by a half duplex Ethernet link, and when two switches are connected by a half-duplex Ethernet link. This result indicates the need to use full duplex links to prevent contention and the resulting capture effect.

6.3 The general multihop case

We now consider general multihop configurations. Two such configurations are examined; one in which half duplex Ethernet links connect the consecutive switches, and the other in which full duplex such links are used. (See Figures 12 and 13.) A video stream is generated at a sending station at one extreme of the topology and is sent to a receiving station at the other extreme of the topology. At each hop, bursty data traffic is injected from an auxiliary station and destined to another auxiliary station one hop away, in such a way that the data traffic shares the same link and thus output port buffer as the video traffic at that hop. In the half duplex configuration, such data traffic is created in both directions while in the full duplex configuration case, the background data traffic is injected in only the direction of interest (that of the video stream).

In Figure 14, we show the total loss ratio as a function of the number of hops for different end-to-end delay requirements for the half duplex configuration without prioritization and infinite buffer size, for G_s =0.5 Mb/s per station, and M_s=6000 bytes; (we note that with these parameters, the total data load on each Ethernet link is 2.5 Mbps). Figure 15 corresponds to the same scenario but with prioritization. As expected, given the issue of contention and capture, prioritization does not help.

In Figure 16, we show the loss ratio as a function of the number of hops for the full duplex scenario without prioritization, with G_s= 0.5 Mb/s and M_s = 6000 bytes; (note that the total data load on each Ethernet link is 2 Mbps). Figure 17 corresponds to the same scenario but with prioritization. It is clear from these two figures that prioritization in the full duplex case is helpful in bringing the packet loss ratio down and thus achieve lower delay constraints for the same tolerable loss ratio. Finally we note that in the case of full duplex configuration with prioritization, where contention has been totally eliminated, the only randomness that exists in the end-to-end delay of video packets is due to the randomness in the size of the video macroblocks generated by the encoders, and in the size of the resulting packets. If such randomness did not exist, then the curves shown in Figure 17 would have been of the threshold type, whereby the loss rate would have been literally 0 below a certain number of hops and 1 above that number of hops. This number of hops would then be such that the sum of packet transmission times exceeds the desired end-to-end delay.

7.0 Conclusions and Recommendations

Buffer sharing and contention on the Ethernet are the two basic effects that lead to performance degradation when real-time video and non-real-time data traffic are mixed. It has been shown that both are sufficiently important that no network deployment will provide satisfactory service to multimedia applications unless both issues are taken care of. The only solution to contention is to eliminate it completely, by means of switching hubs, dedicated ports for individual stations, and full duplex links connecting stations to their respective ports; the same for Ethernet links connecting adjacent switches in a multiswitch configuration. The solution to the second issue is to provide separate queues for the different types of traffic, and to prioritize access to network links, providing higher priority to traffic with real-time constraints. Finally, we note that, given a certain end-to-end delay requirement, there is a maximum number of hops that the path from source to destination could have.

8.0 References

[1] "P802.1p/D5 Standard for Local and Metropolitan Area Networks - Supplement to Media Access Control (MAC) Bridges: Traffic Class Expediting and Dynamic Multicast Filtering," Institute of Electrical and Electronics Engineers, 1996.

[2] Fouad Tobagi, Pablo Molinero Fernandez and Mansour Karam, "Study of IEEE 802.1p GARP/GMRP Timer Values", Submitted to Infocom' 98.

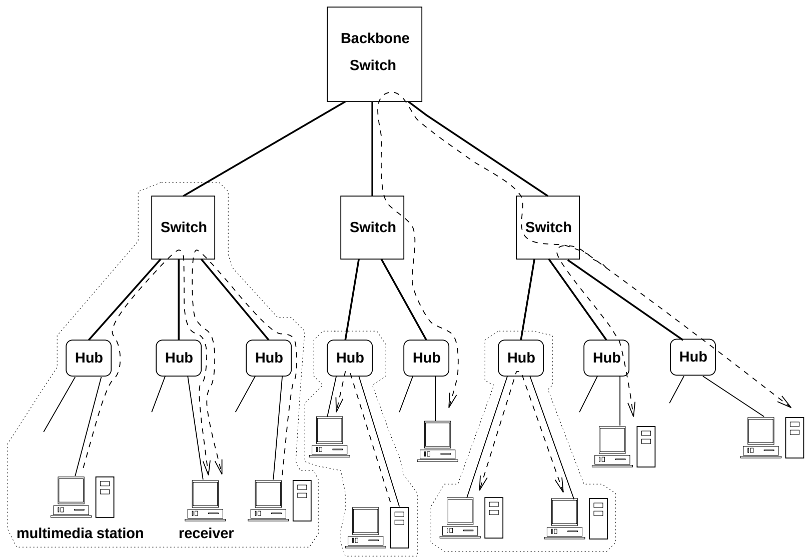

Figure 1. Typical network carrying multimedia traffic.

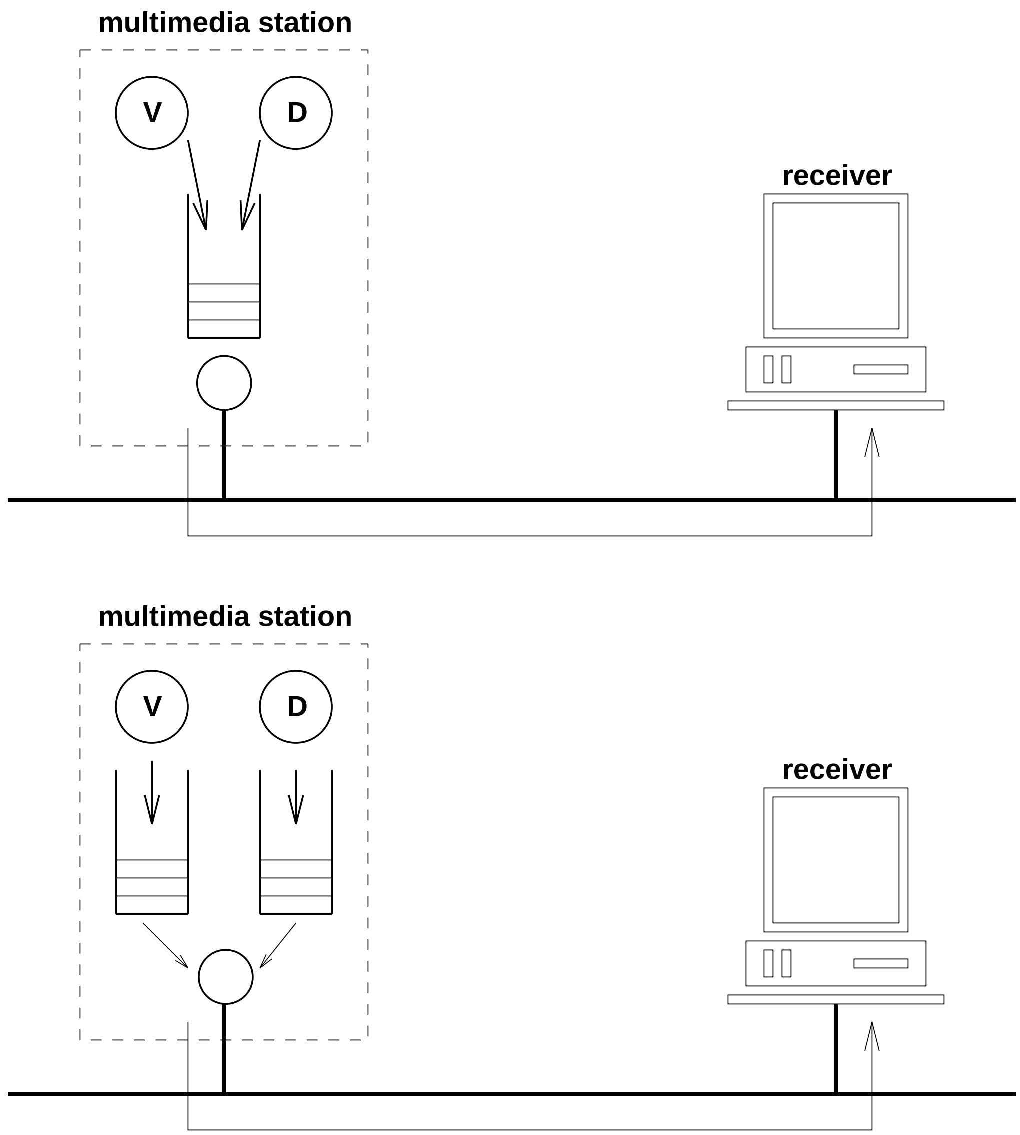

Figure 2. A host sending a video stream and data traffic to a receiver over a dedicated 10 Mbps Ethernet link (Top: buffer sharing at the station; Bottom: separate buffers and prioritization at the station).

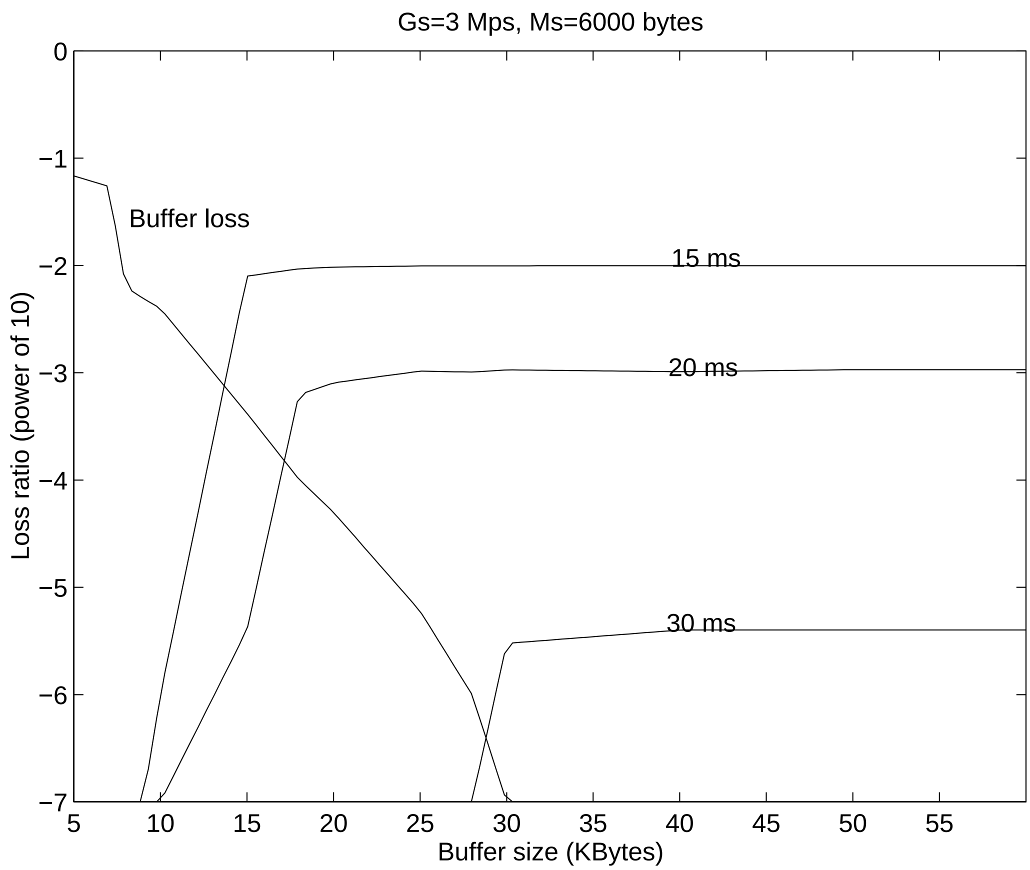

Figure 3. Buffer sharing at the station: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 3 Mbps and M_s = 6000 bytes.

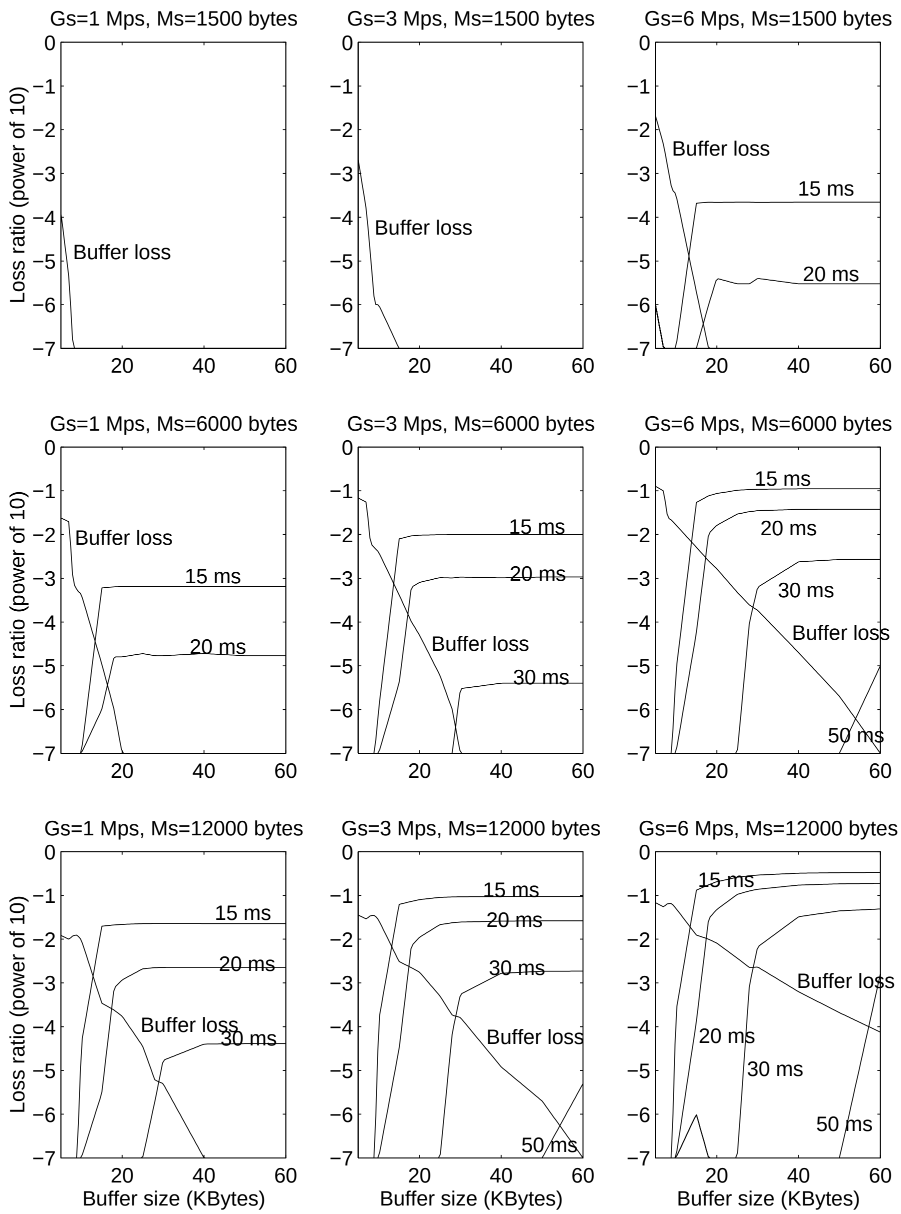

Figure 4. Buffer sharing at the station: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 1, 3 and 6 Mbps and M_s = 1500, 6000 and 12000 bytes.

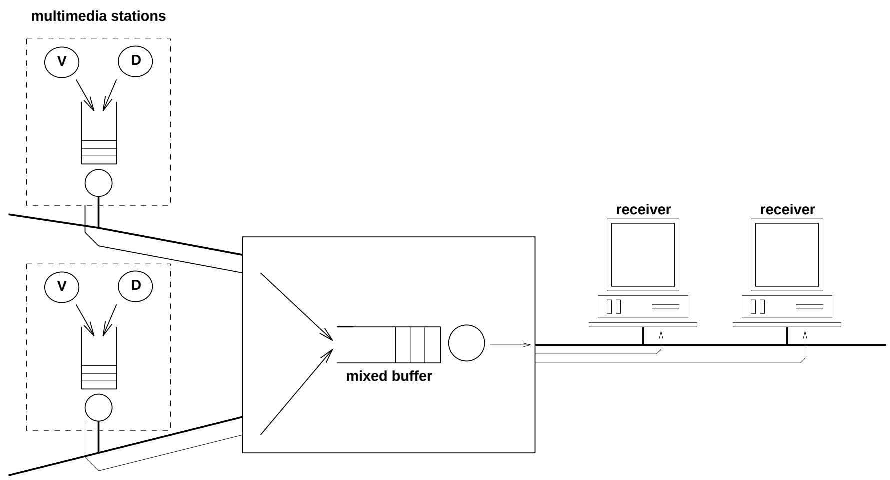

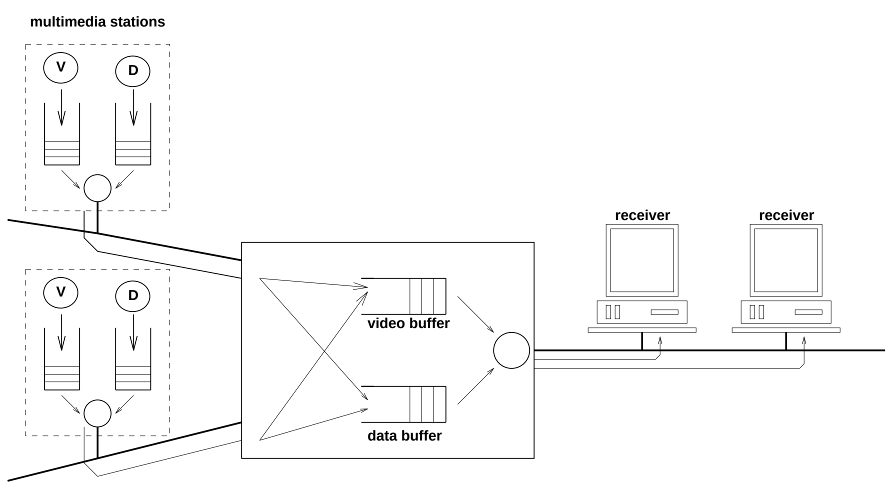

Figure 5. Two multimedia stations connected to a switch, each having its own dedicated port and transmitting both video and data traffic to a receiver connected to a third port (Top: buffer sharing at the stations and the output port; Bottom: separate buffers and prioritization at both the stations and the output port).

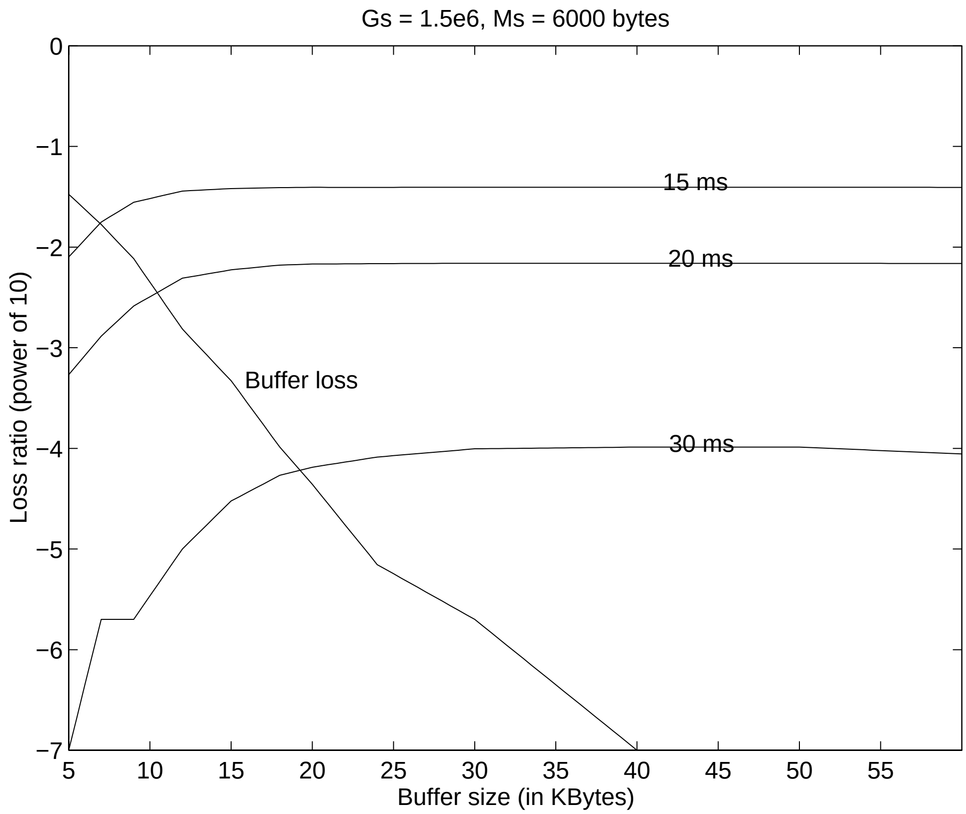

Figure 6. Two multimedia stations connected to a switch, each having its own dedicated port and transmitting both video and data traffic to a receiver connected to a third port with buffer sharing at the stations and the output port: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 1.5 Mbps per station and M_s=6000 bytes.

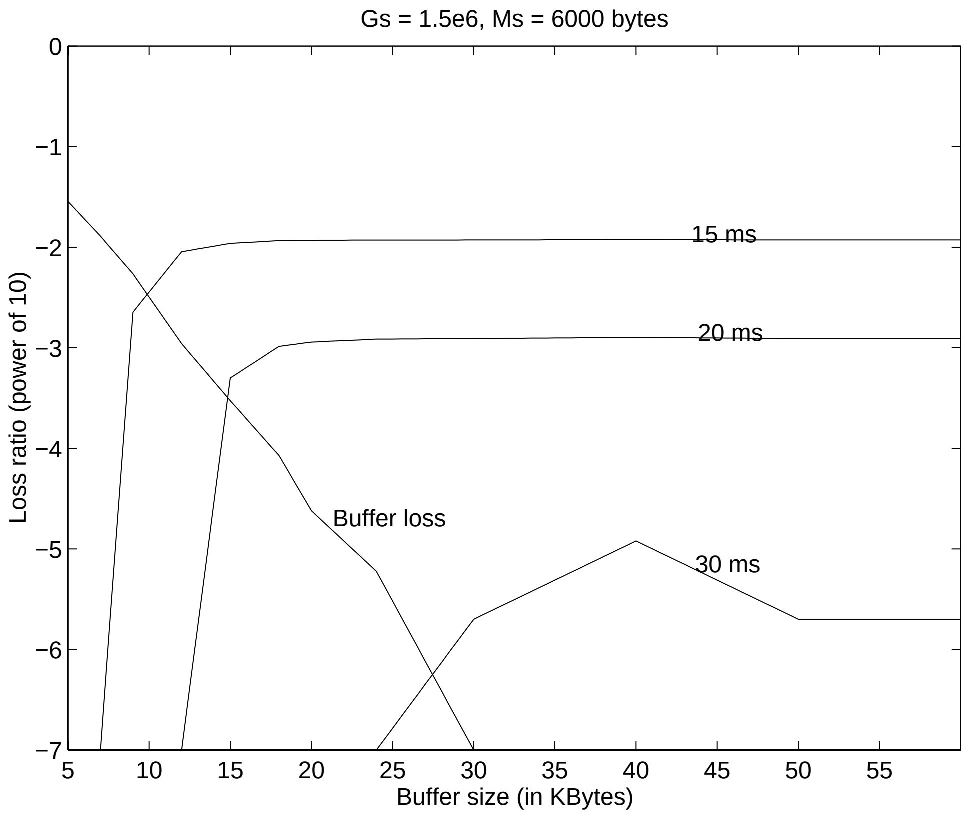

Figure 7. Two multimedia stations connected to a switch, each having its own dedicated port and transmitting both video and data traffic to a receiver connected to a third port with separate buffers and prioritization at the sending stations, and buffer sharing at the output port: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 1.5 Mbps per station and M_s=6000 bytes.

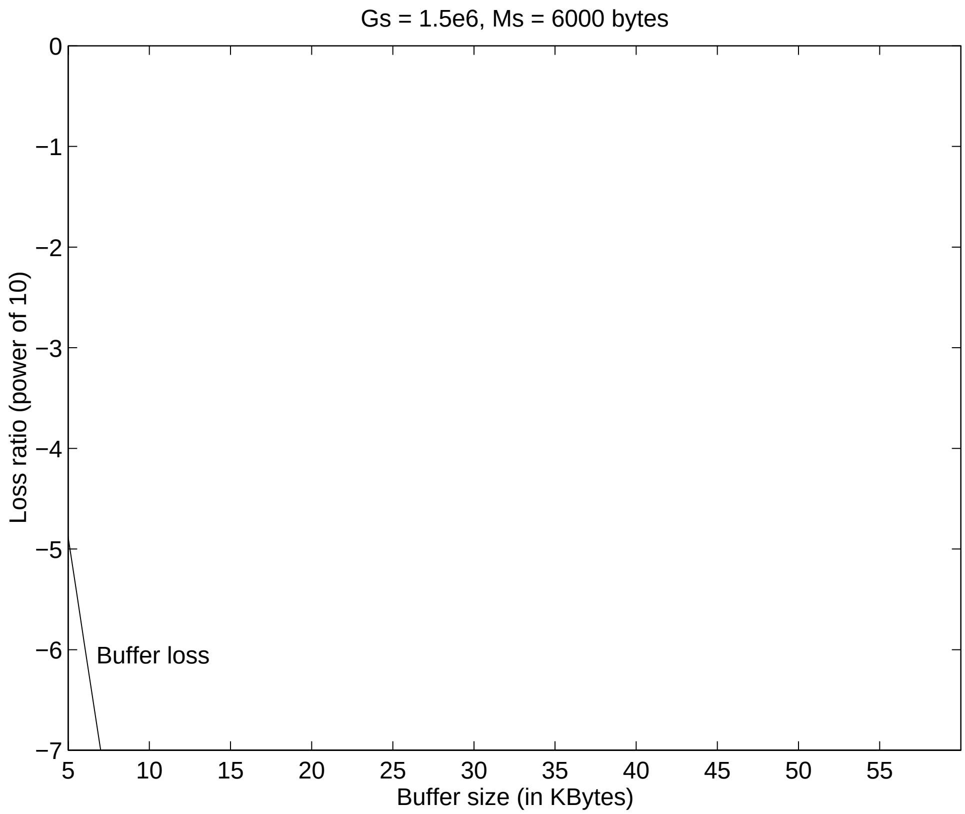

Figure 8. Two multimedia stations connected to a switch, each having its own dedicated port and transmitting both video and data traffic to a receiver connected to a third port with separate buffers and prioritization at the sending stations and switch: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 1.5 Mbps per station and M_s=6000 bytes.

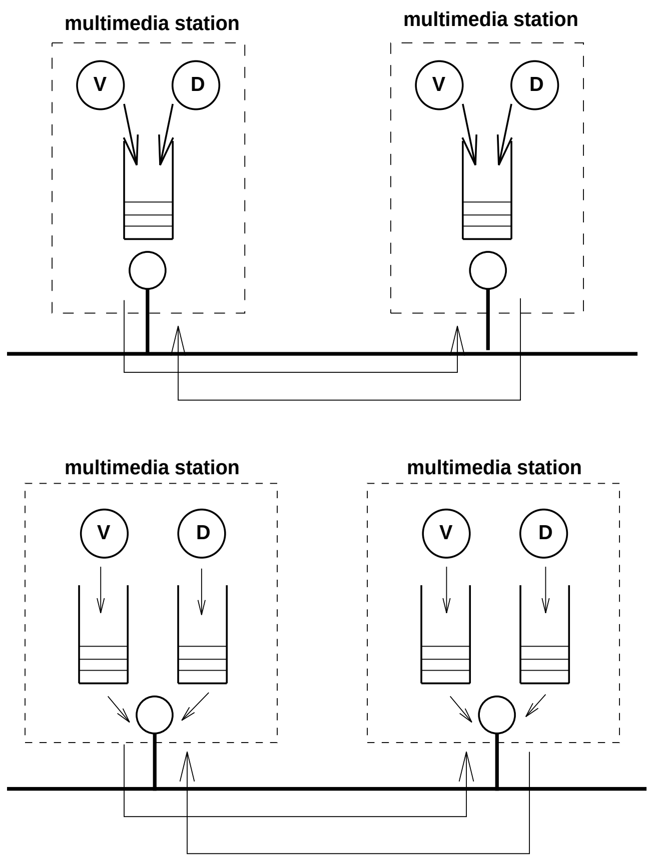

Figure 9. Two multimedia stations connected to the same Ethernet segment communicating both video and data traffic with each other (Top: buffer sharing at the multimedia stations; Bottom: separate buffers and prioritization at the multimedia stations)

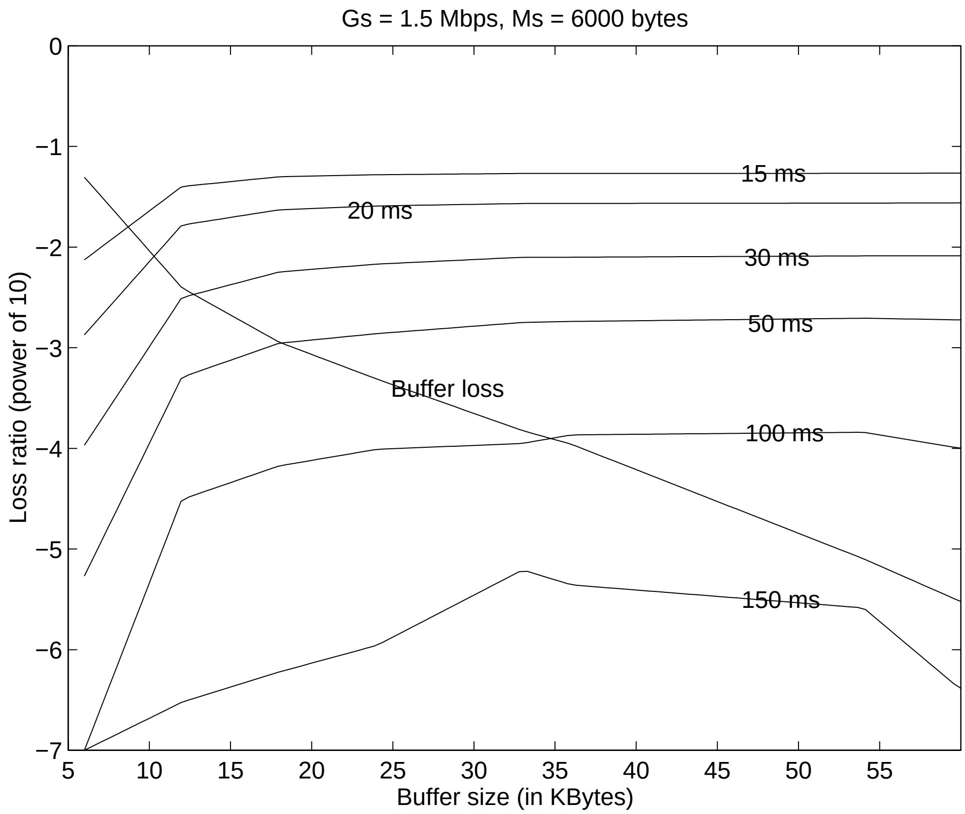

Figure 10. Two multimedia stations connected to the same Ethernet segment communicating both video and data traffic with each other with buffer sharing at the multimedia stations: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 1.5 Mbps and M_s = 6000 bytes.

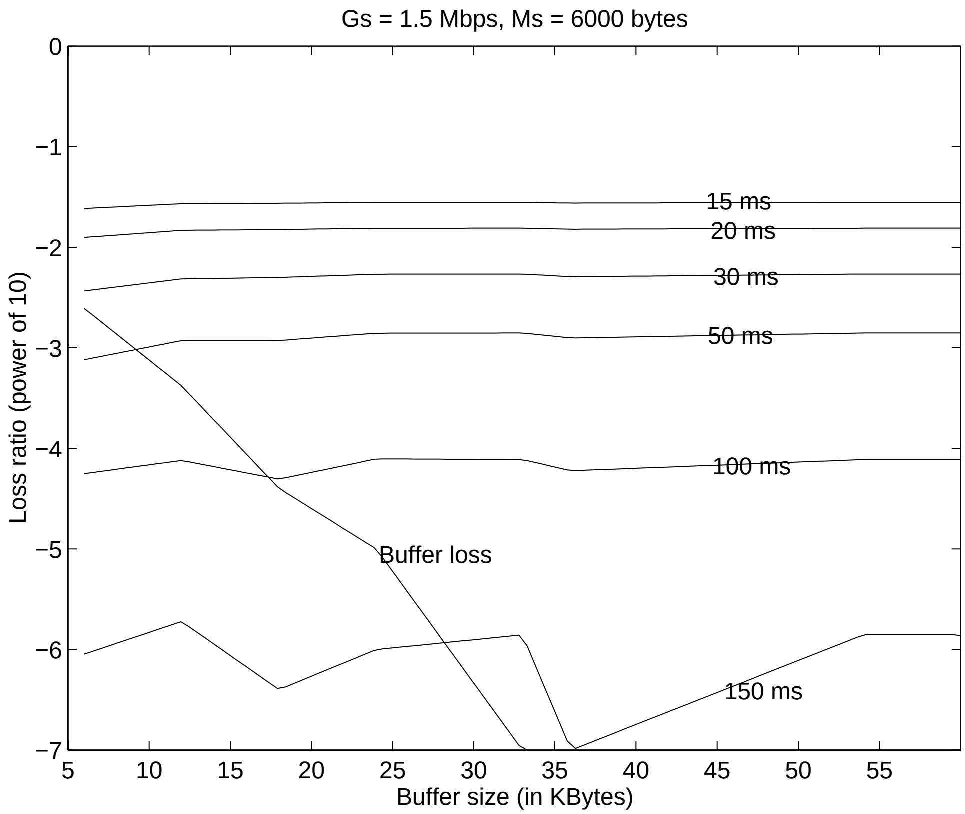

Figure 11. Two multimedia stations connected to the same Ethernet segment communicating both video and data traffic with each other with separate buffers and prioritization at the multimedia stations: buffer loss ratio and end-to-end delay loss ratio versus buffer size for G_s = 1.5 Mbps and M_s = 6000 bytes.

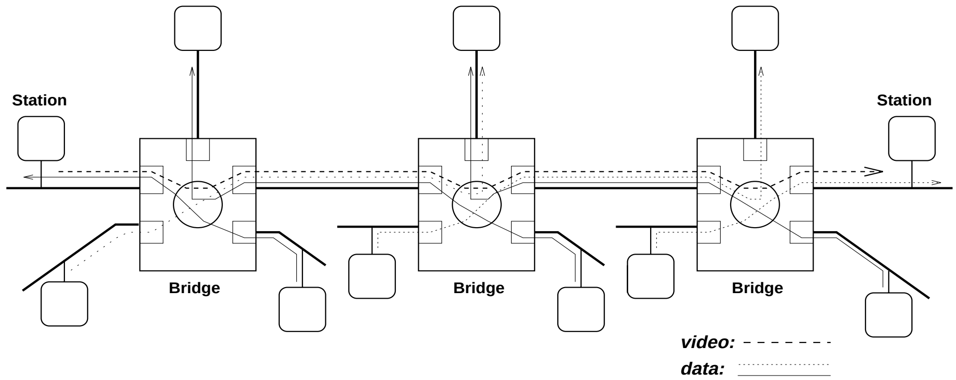

Figure 12. Multihop configuration with half duplex links: A video stream is generated at a sending station at one extreme of the topology and is sent to a receiving station at the other extreme of the topology. At each hop, bursty data traffic is injected from an auxiliary station and destined to another auxiliary station one hop away. Such background data traffic is created in both directions.

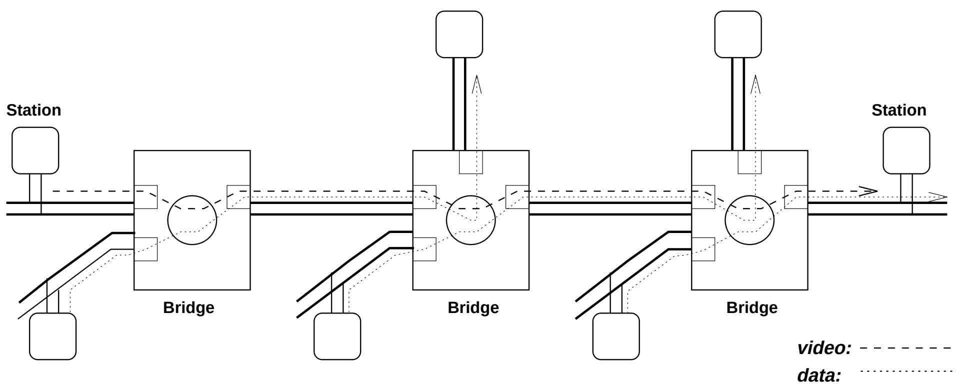

Figure 13. Multihop configuration with full duplex links: A video stream is generated at a sending station at one extreme of the topology and is sent to a receiving station at the other extreme of the topology. At each hop, bursty data traffic is injected from an auxiliary station and destined to another auxiliary station one hop away. This background data traffic is injected in only the direction of interest.

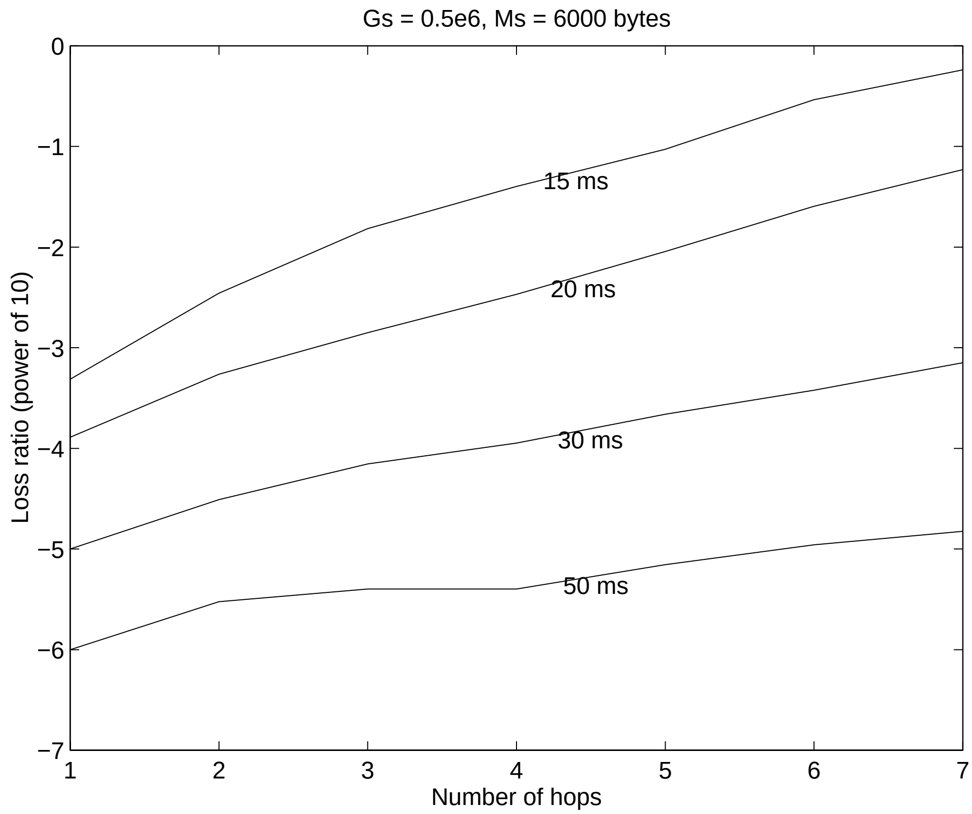

Figure 14. Multihop configuration with half duplex links and shared buffers in the switches: end-to-end delay loss ratio as a function of the number of hops for infinite buffer size, G_s = 0.5 Mbps and M_s = 6000 bytes.

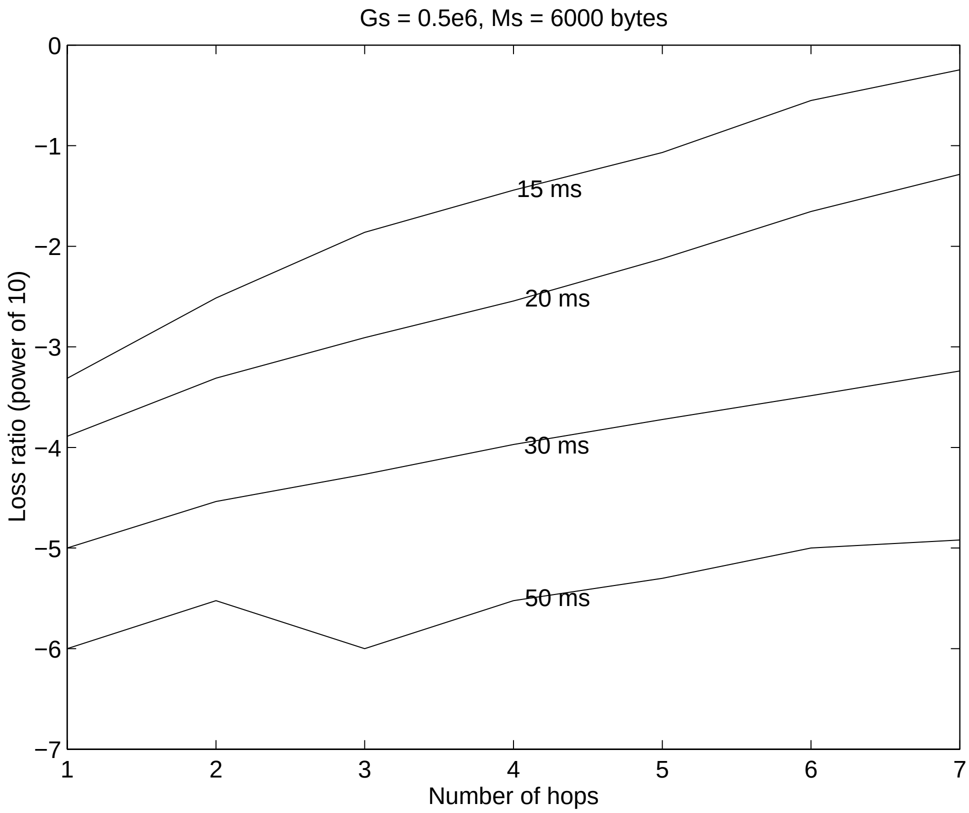

Figure15. Multihop configuration with half duplex links, separate buffers and prioritization in the switches: end-to-end delay loss ratio as a function of the number of hops for infinite buffer size, G_s = 0.5 Mbps and M_s = 6000 bytes.

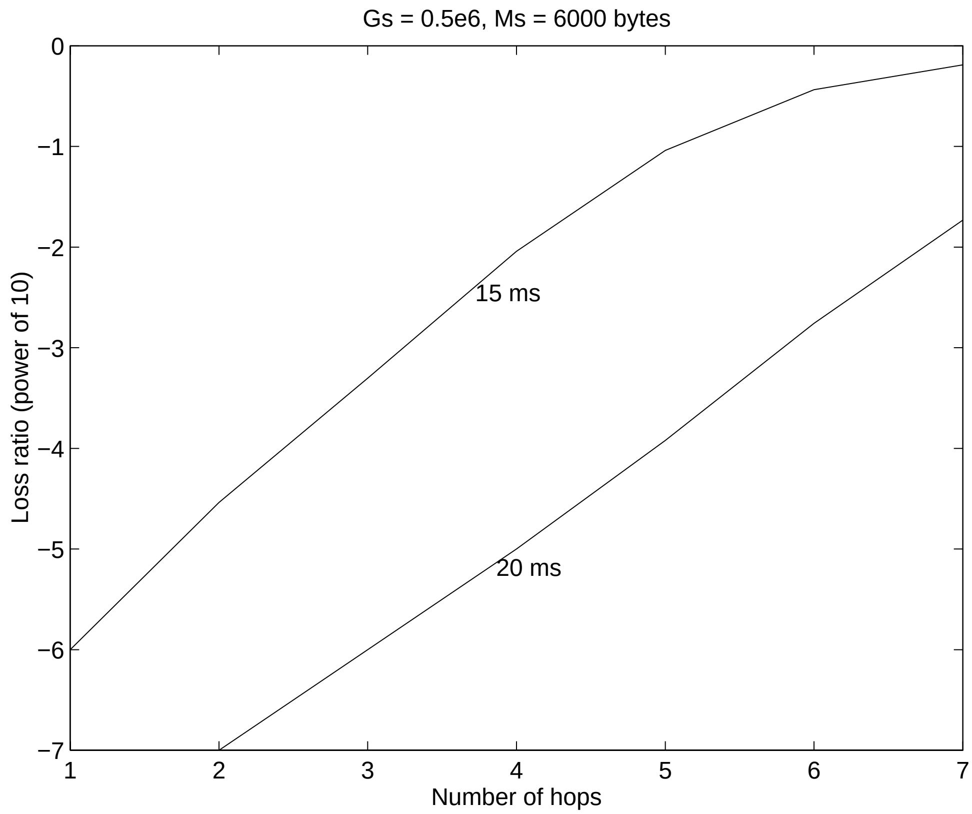

Figure 16. Multihop configuration with full duplex links and shared buffers in the switches: end-to-end delay loss ratio as a function of the number of hops for infinite buffer size, G_s = 0.5 Mbps and M_s = 6000 bytes.

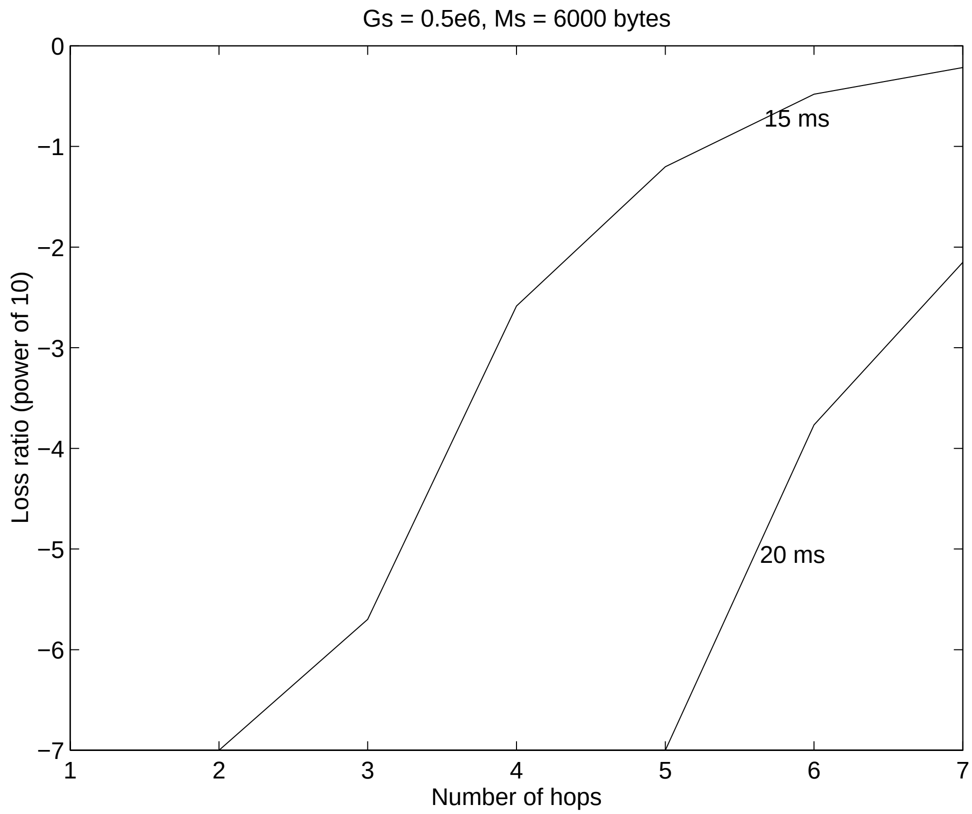

Figure17. Multihop configuration with full duplex links, separate buffers and prioritization in the switches: end-to-end delay loss ratio as a function of the number of hops for infinite buffer size, G_s = 0.5 Mbps and M_s = 6000 bytes.