Multimedia on Local Area Networks

Amr Elsaadany, Mukesh Singhal and Ming T. Liu
Department of Computer and Information Science
The Ohio State University
Columbus, OH 43210

Abstract

In this paper we address the issues related to the delivery of multimedia streams on local area networks. Multimedia integrates voice and video along with text and images into existing systems. Local area networks were designed to handle regular data traffic which are bursty in nature and for which variable delay is acceptable. Multimedia traffic, on the other hand, requires constant delay in addition to fast (or real time) delivery. Multimedia traffic also requires large amount of bandwidth compared to regular traffic. Since local area networks are widely in use, their modification to integrate multimedia streams is very important. In our work we consider some possible solutions to the problems associated with multimedia on local area networks as well as studying the performance aspects of these solutions.

Key words: Multimedia, local area networks, synchronization, priority, switching hubs.

1 Introduction

Multimedia is a media that integrates the transmission of voice and video along with text and images into existing systems. Digital multimedia transmits multimedia information digitally and allows them to be manipulated and stored on a computer system. It also allows for interactive human interface. Interactive multimedia allows the user to interface with the system in real time which is important for education, business, entertainment, and communications.

There are two main problems associated with the design of multimedia systems: the synchronization problem and the media transmission problem. The former problem deals with synchronizing multiple media streams at the presentation site. The latter deals with the transmission of multimedia streams between sites using some kind of communication network.

Multimedia synchronization is the task of coordinating the order and timing of multiple streams for effective communication and high quality presentations. The idea is that no stream should fall behind another causing gap that the user can perceive. Lip synchronization is one example where video needs to be within about 0.1 second of accompanying audio for a good quality motion picture.

The synchronization requirements for the system is basically one of keeping multiple streams (arriving at a presentation site) in sync. In other words, any two related streams arriving at a presentation site should arrive within reasonable time limits in order to avoid gap during their presentation to the user. This may require some work at the sending site as well as the receiving site. At the sending site, streams can be sent ahead of their requested time. This would require some buffering at the receiving site. At the receiving site, however, more work is needed to assure the synchronization of the streams (for example, video frame must match the audio frame in a motion picture). Issues related to multimedia synchronization are reported in [1,2,3,4,5].

A large number of commercial organizations as well as universities use some type of local area network environment and they have two options for adding multimedia capabilities. One is to attempt to utilize their existing networks by adding some devices and/or software in order to transmit multimedia streams. The other option is use an ATM type network (see section 5.2). ATM protocol is designed such that voice and video can be transmitted effectively along with data packets. ATM switches and interfaces will be needed which would add a lot of cost to these organizations.

Existing local area networks have two main problems with regard to the integration of multimedia streams. First, the existing bandwidth capability for a standard LAN does not support multimedia transmission requirements. Second, there is no explicit priority mechanism for transmitting data/stream on LANs. Multimedia streams must be received at the user/presentation site at a time that there would be no interruption (or gap) from user's point of view (consider a moving picture presentation here). In other words, multimedia streams are isochronous, that is, they need to arrive in correct order and with constant delay between packets (and in real time).

Our work focuses on the problems associated with supporting multimedia on local network environments. The protocols of the local area networks are not suitable for multimedia traffic. They were designed to handle regular data traffic which are bursty in nature and for which variable delay is acceptable. Local area networks have to be modified to suit multimedia applications. We introduce some possible solutions for allowing timely delivery multimedia traffic on local area networks and we compare the performance aspects of the possible alternatives.

2 The System Description

The system under consideration is composed of a group of stations connected via local network in some fashion. The simplest example of such interconnection is a local area bus. Another example of such interconnection is shown in Figure 1 where stations are connected on sub-networks which in turn are connected together via some device such as a router/hub. This general network interconnection environment is commonly found in many departments or organizations both in industry and academia.

Multimedia applications are becoming popular and therefore is a need to transmit multimedia streams between stations using existing network interconnections. These streams could be the voice and video of pre-recorded sessions. The system must be able to deliver and present these multimedia streams in a timely fashion. In other words, the system should guarantee the delivery and synchronization of video and audio over these existing stations and their local area network infrastructure.

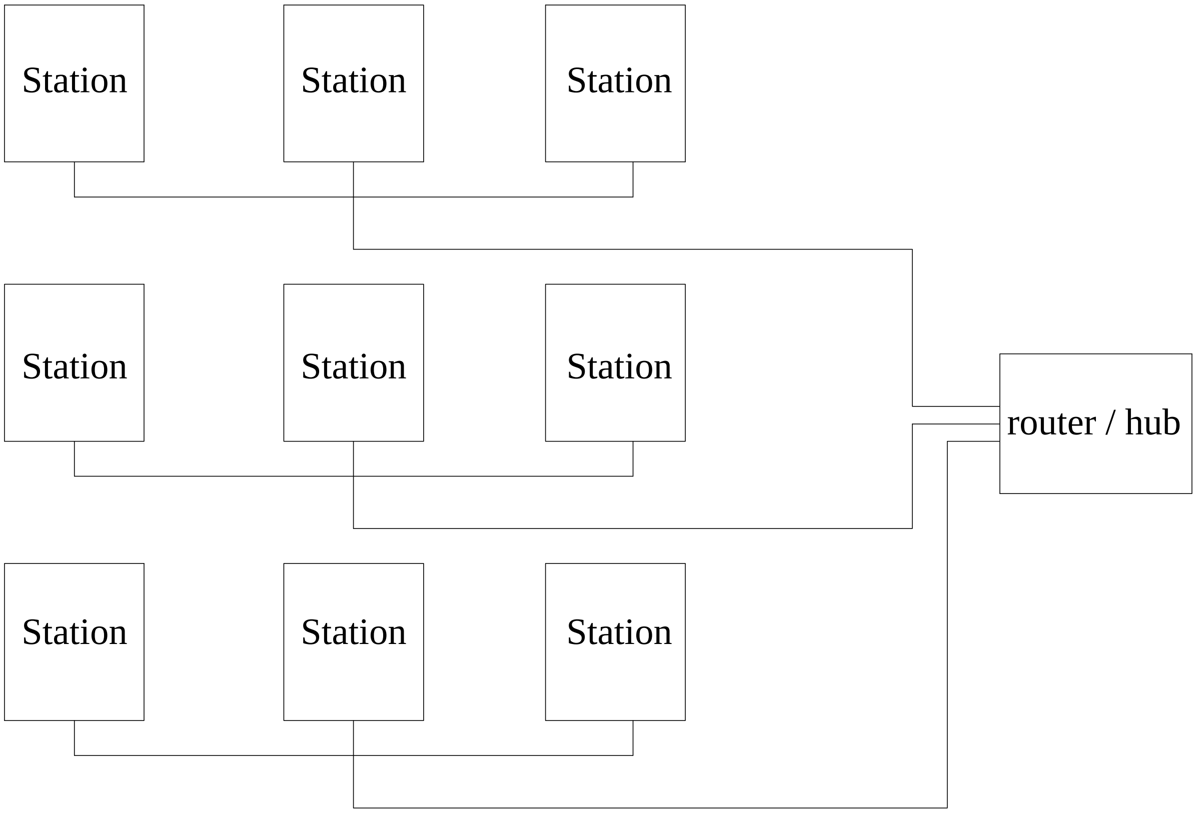

Figure 1. General Local Area Network

3 Motivations

By adding multimedia capabilities to the system described above, a large number of useful services become attainable to organizations. The following is a list of some of these services:

- Announcements/Broadcasts: In this case a person needs to get all members of the department/organization together for brief time to make some announcement. This can be done without having any member leaving his desk. Instead, on the screen, a window could be used to view (and hear) the announcement.

- One on one meeting: This is another example of service where two people can have face to face meeting without having to be physically being in the same room.

- Team meeting (many to many): This is just an expansion of the one to one meetings.

- Review meetings (many to one): In this case one person is receiving comments/input from a group of reviewers/customers.

  All of the above services will save organizations a lot of personnel time and will increase productivity. People can make better use of their time because one could be doing some other work concurrently while attending a meeting. This, of course, is in addition to having access to all files and documentations.

- Training video library: A video library of classes or technical presentations can be stored on some disk and be accessible to user sites on the network. This is a very useful and cost effective way of training new employees. The trainee will have access to high quality video presentations or class materials in a very convenient matter (on a window at his workstation). This system gets ride of all the disadvantages of a video tape library with a VCR player where many conflicts can arise, let alone the quality.

4 Multimedia Data Requirements

All digital media requires digital end-to-end transmission of multimedia data streams. This means that the communication network must support the transmission of digitized voice, video and images. The network must also have a wide bandwidth to support these transmissions. Not all existing networks have sufficient bandwidth to accommodate the needs of all kinds of multimedia applications. For example, a LAN with 10 Mbps bandwidth may be able to provide limited support to transfer video and images for some multimedia applications.

Multimedia applications require fast processing and small visual response time. Moreover the effect of latency in transmitting streams between stations should be minimized or eliminated. In order to further understand the requirements for multimedia applications, we will review some of the properties of the multimedia data types/streams.

Graphical Images:

Graphic images to be displayed are stored in a display memory (or video memory) which usually has faster access times than regular memory. In a one bit display system, every pixel on the screen is represented by a bit in the display memory. The memory used to represent the entire screen in this case is called a bit plane. This system is capable of displaying images in a two color mode (black and white).

As the color resolution (or pixel depth) increases, more bits are required to encode the new colors. In this case the total of all the bit planes required to process and display an image is called bitmap. Most current workstations provide full colored screen capability using a 24-bit color resolution with 8 bits for each of the three basic colors (red, green and blue). This requires 24 bit planes for storing a display and can provide more than 16.7 million possible colors. However it requires a massive amounts of storage, for example, a 640 x 480 display takes about 900KB of memory. With the multimedia requirement of instantaneous display time, this system may not be adequate [6].

The amount of memory required can be reduced by using a color map which is a limited size table that contains an entry for each of the set of colors that can be displayed simultaneously. For a bit map consisting of eight bit planes, the color map size would be 256 long. Thus allowing for 256 of the possible colors to be used at any one given time. Each table entry (often called color cell) includes three color codes (8 bits each) corresponding to the three basic colors. Using this (indexed) color system the 640 x 480 would take about 300K of memory. The value of the pixel is used as the index to a color cell in the color map.

Some systems employ double buffering to avoid any glitches in the image presentation such as in real time animation. The size of the display memory in this case is enough to hold two screenful of data. Next frame is being rendered while the current one is being displayed.

Aside from the high storage requirements, bit maps have fixed pixel value and thus fixed image resolution which could present problems if image is displayed on devices with varying resolution. The work around this problem is to use structured graphics which describes the elements in the picture (using some commands) rather than the pixels. These commands are then used to render the picture into a bit map that matches the resolution of the displaying device. Postscript is an example of structured graphics systems.

Audio Technology:

Sound can be represented by either analog or digital technology. Analog sound exists in physical (waves) or electric formats. Physical sound waves is the natural format for human hearing (which can sense frequencies in the range of 20Hz - 20KHz). This format can be transduced into electrical current via microphones (and back into sound waves via speakers). The electrical format allows for higher amplifying ranges than that possible with the physical format (in addition to some manipulation capabilities).

Digital sound is transformed from analog sound via analog-to-digital converters. It can be manipulated, stored, transmitted or copied without degradation. The digital audio quality (and the storage requirement) depend on the sampling rate and the digitization resolution (number of bits used to to represent a sample). Most common audio CD uses 16-bit resolution and 44.1KHz sampling rate. One second of this high quality audio requires 176.4 kB of storage. A mono system which uses 8-bit audio at 22.01KHz requires 20.8 kB.

Digital sound can then be stored on audio CD and would require a CD-ROM drive and speakers to produce the analog sound. Another option is to store sound in files known as waveform files. An audio card takes sound from an input device (e.g. microphone) and digitize, encode and store it (on the hard disk). At playback time this audio card converts the digital audio to analog and send it to the speakers.

While digital audio casts musical elements (notes) into bit streams, they can be described using MIDI (Musical Instrument Digital Interface). MIDI is a technology for storing, editing and transmitting music (using a MIDI sequencer). It allows the computer to accept input directly from a musical instrument. It records music by substituting each musical note with a command which describes the pitch, duration, and type of instrument playing the note. Later a MIDI synthesizer can recreates the music using these commands. In addition to using less storage space than digital audio, MIDI has the advantage of enabling the editing and manipulating the MIDI files in order to create an otherwise unachievable performance.

Video Technology:

Whatever applys to digital images and digital audio (combined) applys as well to digital video technology. The one additional factor is the frame rate which the number of picture frames that should be presented in one second (NTSC uses 30fps). One second of 30 fps video of size 640 x 480 pixels and 24 bit color resolution requires about 27MBps bandwidth. A 320 x 240 would require about 8MBps and a 1024 x 1280 requires about 120MBps. It is clear that the amount of data needed to represent video digitally adds a lot of requirements on storage, transmission, bandwidth as well as display capabilities of current systems.

Miscellaneous Items:

In addition to the storage and bandwidth requirements for multimedia data, there are some other problems associated with multimedia such as the suitability of existing network protocols. Multimedia applications and tools are being rapidly developed (by multiple vendors) and at the same time standards have not been fully developed. Standards should contribute to the growth of multimedia by resolving compatibility issues and thus allowing the development of multimedia applications across different systems and enable users to share multimedia technology and applications.

On the other hand, quite a lot progress has taken place in the area of multimedia data compression techniques. This would reduce the amount of bandwidth required for streams such as video. In addition, the amount of disk storage required for multimedia data can be substantially decreased.

The JPEG (Joint Photographic Experts Group) standard can achieve 25-50 times reduction in image size without sacrificing the quality of decompressed images. This technique works by compressing the variations in color and brightness. Another standard for full motion images is MPEG (Moving Pictures Experts Group). This technique works by storing only the variations between successive frames. The 27MBps transfer rate required for the 640 x 480 screen would be reduced to 550kBps using MPEG1. MPEG1 can also compress audio by a 5-10 factor [7].

5 Communication Network Background

5.1 Local Area Networks

Local area networks (LAN) exists in one of the following architectures; bus, tree, token ring and FDDI (Fiber Dual Data Interface) ring. The most widely used LAN is the Ethernet which is a bus based architecture. This type of network typically has 10 Mbps bandwidth as opposed to 100 Mbps for FDDI.

Local area network bus (non token) protocols are based on broadcasting. The protocol used for these networks is based on the idea of the CSMA (Carrier Sense Multiple Access) protocol in which each station listens to the channel before transmitting. If the channel is idle, then transmit and wait for acknowledgement otherwise back off and try again after some period of time. In case a collision happens (could be determined by lack of acknowledgement from destination), the station resends after some period of time [8,9].

The most dominating medium access protocol (MAC) in use is the CSMA/CD (Carrier Sense Multiple Access with Collision Detection) in which each station listens to the channel while transmitting. In case a collision is detected during transmission, the station ceases transmission and retries after some (random) period of time. This protocol requires that the message size be longer than twice the propagation time (in order to detect collision before transmission is complete). It also use binary exponential backoff technique, in which, after each repeated collision, the mean value of the random delay is doubled.

One other MAC that is widely used is the token bus in which stations on the bus form a logical ring with a token (control packet) to regulate access to the medium. When a station receives the token, it get control of the medium for a specified time. When the station is done, or time expires, it passes the token to the next station in the logical ring.

LAN topologies are based on sharing the bandwidth between stations. Data sent are assumed to be time-independent, they can be broken into packets without regard to their order. Also, when one station is using the network for transmission, no other station can and the entire bandwidth is assigned to that one station.

To overcome the problems with the speed and/or order of transmission on LAN, several techniques are being developed. Fast Ethernet is one way to solve the bandwidth problem of the standard Ethernet and it is expected to be 10 times wider. It may retain the CSMA/CD MAC protocol to allow internetworking with standard Ethernet LANs.

Another way to overcome the bandwidth problem is to use switching hubs. Switched Ethernet establishes a single point-to-point LAN between each workstation and a high-capacity switching hub. This gives each workstation its own LAN with all of its bandwidth. (Some topologies may have more than one workstation per LAN.) However the switching hub may be a bottle neck at some point. That is why this network architecture impose limits on the number of workstations that could be connected to the hub and also on the number of simultaneous users.

Other options include using FDDI (running over twisted pair for a short distance) to increase the bandwidth to 100 Mbps. FDDI technology could be extended to support the timely delivery required by multimedia applications. Isochronous Ethernet is another option which provides isochronous channels to stations equipped with isoENET cards [10].

5.2 Switching Networks

Switching networks transfer information between source and destination via three functions: transmission, multiplexing, and switching. Transmission provides the point-to-point transfer of information. Multiplexing uses a single transmission facility (link) for several independent channels. Switching uses properties of telecommunications traffic to replace a fully meshed network interconnecting all pairs of users with a shared substructure providing on demand connections, that is, packets can be assinged to channels based on the address information in the header of the packets themselves.

Circuit switching networks use time division multiplexing (TDM) where bandwidth is multiplexed between several channels with assigned amounts of bandwidth. Time slots within a frame are allocated to the channel for the duration of the service. This type of networks is suitable for applications that require guaranteed bandwidth or low latency such as PCM voice and constant bit rate data traffic. The mode of transfer in this type of network is referred to as synchronous transfer mode (STM).

Packet switching networks use bandwidth efficiently to suit bursty applications such as file transfers. However, the delay in these networks is unpredictable and the latency tends to be high. Packet switching uses statistical multiplexing which dynamically allocates time slots to incoming channels based on demand. In contrast with synchronous TDM, the positional significance of the slots (in TDM) is lost and hence each packet has to carry addressing information in addition to data.

ATM (Asynchronous Transfer Mode) uses cell switching which uses bandwidth flexibly and efficiently for all sorts of applications. For constant rate applications, a guaranteed number of cells per second can be allocated and any unused cell slots can be allocated (on demand) to bursty applications. In addition, the number of virtual channels carrying these cells can vary as well as the rate of each channel.

In cell switching technology, a cell is a 53 byte fixed length packet with 5 bytes for header (containing routing and quality of service information) and 48 bytes for data (payload section) [11]. This fixed length cells are better suited for high speed switching because it enables simpler implementation of switches and provides predictable system behavior.

The asynchronous nature of ATM comes from the fact that cells are assigned asynchronously to any service based on its needs. This is facilitated by the fact that each cell is addressed independently to any destination. In this mode of transmission, there is no notion of owning a time slot on periodic basis.

Although existing wide networks (WAN) can be used by some multimedia applications, these WAN-based applications would suffer from existing problems such as the varying network delays. When it become widely available, ATM will be the most suitable technology that has the high bandwidth and low latency requirements for all kinds of multimedia applications.

6 Possible Solutions

As mentioned before, our work focuses on the problems associated with supporting multimedia services on local network environments. The protocols for regular data transfer on the local area networks are not suitable for multimedia traffic and they have to be modified to suit multimedia applications. They were designed to handle regular data traffic which are bursty in nature and for which variable delay is acceptable. We present some possible solutions for allowing timely delivery of multimedia traffic on local area networks. These solutions are either software solutions or a combination of hardware and software. A software solution is one who can be fully implemented without any additional hardware equipments. We focus our attention to the problems associated with supporting multimedia traffic on with local area bus (Ethernet) as opposed to token ring networks [12].

6.1 A Two-Priority Scheme

The current medium access bus protocol does not provide priority for transfer of data. Stations on the bus use the CSMA/CD protocol to compete for bus bandwidth. This means packets can indiscriminately collide and get retransmitted. This medium access protocol is suitable for regular data traffic but not for multimedia traffic. We propose a two-priority scheme that add some sense of priority to data traffic on the bus. It supports two levels of priority: high priority (for multimedia or real-time traffic) and low priority (for regular data). It relies on message passing to support these two levels of priority.

Before a station sends a "high priority" traffic (for a multimedia session), it informs every other station of its intention so that they would abstain from sending "low priority" traffic for a period of time. This is accomplished by broadcasting a request message that includes the amount of time that the station needs (i.e., the session length). During this time period, only high priority traffic is allowed to use the bus.

Since the request message will compete for the bus with other traffic, the scheme has to ensure that this message get successfully transmitted before the high priority traffic period starts.

When stations receive a message indicating that high priority traffic need to be sent by a station, they stop sending low priority traffic at once. Stations records the time at which they will be allowed to send regular traffic. This time would be equivalent to the current time plus the length of the session (which is included in the request message). Stations also keep a count of the number of multimedia sessions currently in progress. Whenever a station receives a request message, it increment this count by one.

Within a high priority period, if another station wants to start a multimedia session, it can use the bus simultaneously since it knows that the bus is being used for high priority traffic. However, the station has to send a (request) message to other stations. This message will be used to update the time at which the high priority period will terminate. In addition, stations will increment the number of current multimedia sessions by one. During this time period, high priority traffic compete for the bus using the CSMA/CD medium access protocol (but they do not compete with low priority traffic).

After a station finishes sending its high priority traffic, it broadcast a completion message indicating that the number of multimedia sessions is decremented by one. Again the scheme has to ensure that the completion message get delivered even after it collide with other traffic.

Figure 2. Bus Time

Thus the bus is time multiplexed between high priority and low priority traffic, as illustrated by Figure 2, where Tm is time period used by multimedia (high priority) traffic and Tr is time period during which regular (low priority) traffic can be sent.

The period (Tm) during which the bus time is dedicated to multimedia traffic is measured from the start of the high priority traffic session(s). This period ends when the count of multimedia sessions goes down to zero or if Tm exceeds a pre-determined maximum limit. If Tm exceeds this limit, regular traffic will be allowed to use the bus. The ratio Tm/Tr need to be fixed (or be within some margin) to insure fairness to regular data traffic.

The number of simultaneous multimedia sessions has to be limited, otherwise, they would interfear with each other producing low quality in addition to increasing the regular data traffic delay. This limitation is dependent on the amount of available bandwidth and the nature of the multimedia applications.

This technique should result in better multimedia traffic performance over the non-prioritized bus protocol but regular data traffic will suffer some delay. This delay could be acceptable to a user sending regular data. On the other hand, a delay in video frames, for example, may result in an unacceptable presentation quality.

6.2 A Station Multiplexing Scheme

In the station multiplexing scheme, each station keeps status information about the transmission requests of all other stations. Using this status information, each station will be able to make decision to send packets on the bus. Each station is still required to send a request message before starting its multimedia session (i.e., high traffic session), but stations are not required to send a message after completing their sessions.

In this scheme, each station maintains a data structure that has an entry for each station. The entry for a station contains the time of the latest transmission request made by that station for high priority traffic. An entry for a station is updated when a (request) message is received indicating that it wants to start its multimedia session.

All stations use a fixed length quantum (Tq) for high priority traffic. A station uses the bus for this quantum and renews its request for another after the current quantum expires. A station wanting to start a session or renew its request for another quantum, compares the timestamps in its local structure. The station that corresponds to the smallest timestamp is the one that can use the bus next (after the current quantum expires).

A station can find if the current quantum has expired by comparing the difference between the latest (highest) timestamps in its local structure and the current time with the default session length quantum. If the current time minus the value of the latest timestamp is less than or equal the value of Tq, the station may proceed with its request immediately. Otherwise, the station tries whenever the current quantum expires.

In this scheme, the bus time is still multiplexed between high priority and low priority traffics, but in addition, the Tm period is further divided into small time quanta (of duration Tq) and each station uses a quantum (if needed) based on the latest timestamp criterion stated above. This would improve performance further since the competition between stations (during Tm) is regulated. In other words, this approach further reduces the possibility of collisions.

The basic idea is that stations that need to send high priority traffic will use the bus in a last recently used order. The station with the oldest (smallest) timestamp will be allowed to use the bus once the current Tq expires.

If two stations send their request messages at the same time, these messages will collide. In such cases, one station will backoff its transmission of request message and the other will proceed. (Station numbers can be used for tie breaking.)

As in the previous scheme, the period (Tm) during which the bus is dedicated to multimedia traffic has a maximum length from the start of the high priority traffic session(s). If Tm exceeds this limit, regular traffic will be allowed to use the bus. Again, the ratio Tm/Tr need to be fixed (or be within some margin) to ensure fairness to regular data traffic.

A variation of this scheme is a round robin scheme, where stations are served in a predefined order and for a fixed period of time. This scheme is similar to a token based scheme, where stations pass a token among themselves in a pre-defined order and only the station that has the token can use the bus for its high priority traffic for a predefined period. When timer expires, station passes the token to the next one. While this scheme reduces the competition between stations trying to send high priority traffic, it may result in an unacceptable response time for some users.

6.3 Switching Hubs

Switching hubs are used to connect LAN segments. A large LAN can be divided into smaller segments linked by a switching hub (see Figure 1). This would overcome some of the bandwidth limitations on the large LAN. Intra-segment communications are carried out on the local buses (may be using one of the software approaches introduced above). While inter-segment communications can be carried out through the hub.

The software approaches discussed above provide priority to multimedia traffic. But due to the limited bandwidth there will be some limitation imposed on multimedia applications. For instance, the number of multimedia sessions have to be limited. By splitting the stations between several buses which are connected via a switching hub, the number of simultaneous users can be increased.

However the switching hub may be a bottle neck at some point. That is why there may still be some limits on the number of stations that could be connected to (the segments on) the hub and also on the number of simultaneous users. In our work we will consider factors like the number of buses going through the hub and the number of stations on a single bus.

7 Discussion

From communication point of view, there are three areas that are currently being investigated for multimedia applications. The first uses cell switching in an ATM-type network which can provide the needed capacity for all multimedia applications. However this is a new technology and is not widely available. The second uses the existing packet switching protocol in WANs to transmit multimedia streams. As discussed earlier, varying delays is the major problem the normal packets switching protocol. The third uses the protocols of LANs to transmit multimedia data. LANs are widely available but they suffer from low bandwidth and they do not support priority traffic.

Our work focuses on supporting the timely delivery of multimedia traffic on local area networks. In particular, we study the protocols and topologies of bus-based local area networks. Other issues still need some design considerations [13]. These issues include file management system, storage systems, the design of routers and bridges, and the modification of the transport protocol to suite video traffic.

The medium access protocol for local area bus does not provide priority for regular data traffic. Stations along the bus compete for bus bandwidth. This mean that packets can indiscriminately collide and get retransmitted. This medium access protocol is suitable for data traffic but not for multimedia traffic. We introduced some techniques suitable for multimedia streams on LANs. These techniques adds some sense of priority to the traffic on the bus.

The problem with multimedia streams is that they are isochronous in nature and they need to be delivered in real time. Thus a mixture of (multimedia and regular) traffic on a LAN bus would not satisfy this requirement. Multimedia traffic need to be transmitted at higher priority than regular data traffic. One way to establish such priority is to have stations send messages to each other (and get acknowledgements) before they start their multimedia sessions (to send high priority traffic). Considering the broadcast nature of LAN buses, this approach is not suitable. Bandwidth would be wasted by the transmission of all these messages. We propose a two-priority scheme and a station multiplexing scheme.

The two-priority scheme supports two levels of priority (high priority and low priority) and it multiplexes the bus time between the two traffic types. It relies on sending a message to inform all stations that bus should be used for high priority traffic for some period of time during which low priority traffic are not transmitted. This technique should result in better multimedia traffic performance over the non-prioritized bus protocol but regular data traffic will suffer some delay. This delay could be acceptable to a user sending regular data. On the other hand, a delay in video frames, for example, may result in an unacceptable presentation quality.

In the station multiplexing scheme, the bus time is still multiplexed between high priority and low priority traffic, but in addition, the high-priority-period is further divided into small time quanta and each station uses a quantum (if needed) based on the last recently used criteria. This would improve performance further since the competition between stations during the high-priority-period is regulated which in effect reduce the number of collisions.

Another area under study is the use of switched hubs to overcome some of the bandwidth the limitations. Using a switching hub in a LAN would split the stations between several bus segments which are connected via the hub. Only inter-segment communications need to go through the hub. Factors like bottle neck at the hub, the number of buses going through the hub and the number of stations on a single bus are being considered.

Currently we are building a simulation model of the system. First, a simulation of the bus without any priority mechanism will be examined under a mixture of multimedia and regular traffics. Then simulations of the bus with the priority schemes introduced above will be made to study the improvement in multimedia traffic performance.

REFERENCES

1. S. Ramanathan and Venkat Rangan, "Feedback Techniques for Intra-Media Continuity and Inter-Media Synchronization in Distributed Multimedia Systems," The Computer Journal, Vol. 36, No. 1, 1993, 19-31.

2. T. D. C. Little and A Ghafoor, "Multimedia Synchronization Protocols for Broadband Integrated Services," IEEE Journal on Selected Areas in Communication, 1991, 1368-1482.

3. R. B. Dannenberg, "Remote Access to Interactive Media," SPIE, Vol. 1785,Enabling Technologies for High-Bandwidth Applications, 1992, 230-237.

4. R. Steinmetz, "Synchronization Properties in Multimedia Systems," IEEE Journal on Selected Areas in Communication, 1990, 401-412.

5. C. Nicolaou, "An Architecture for Real-Time multimedia Communication System

6. Jeff Burger, "The Desktop Multimedia Bible," Addison Wesley, 1993.

7. Bernard Cole, "The Technology Framework," IEEE Spectrum, Vol 30, No. 3, March 1993, 32-39.

8. William Stallings, "Data and Computer Communications," Macmillan Publishing, Second Edition, 1988.

9. Andrew S. Tanenbaum, "Computer Networks," Printce Hall, Second Edition, 1989

10. Daniel Minoli, "Isochronous Ethernet: Poised for Launch," Network Computing, August 1993, 156-162.

11. William Stallings, "ISDN and Broadband ISDN," Macmillan Publishing, Second Edition, 1992.

12. Khaled Amer, Ken Christensen and Tom Toher, "Experiments with Client/Server Multimedia on Token Ring," Proceedings of 18th Conference on Local Computer Networks, September 1993, 2-6.

13. Fouad A. Tobagi, "Multimedia: The Challenge Behind the Vision," Data Communications, Jan 21, 1993.

ACRONYMS

ATM	Asynchronous Transfer Mode
CSMA	Carrier Sense Multiple Access
FDDI	Fiber Distributed Data Interface
fps	Frames per second
GB	Giga Bytes
JPEG	Joint Photographic Experts Group
kB	Kilo Bytes
kBps	Kilo Bytes per second
kbps	Kilo bits per second
LAN	Local Area Network
MAC	Medium Access Protocol
MIDI	Musical Instrument Digital Interface
MPEG	Moving Pictures Experts Group
Mbps	Mega bits per second
NTSC	National Television Systems Committee
PCM	Pulse Code Modulation
QOS	Quality of Service
STM	Synchronous Transfer Mode
TDM	Time Division Multiplexing
VCR	Video Cassette Player
WAN	Wide Area Network