MULTIMEDIA PERFORMANCE EVALUATION OF ETHERNET AND TOKEN RING MEDIA ACCESS PROTOCOLS: A NETWORK COMPARISON

by

Daniel A. LaHue

A Master's paper submitted to the faculty
of the School of Information and Library Science
of the University of North Carolina at Chapel Hill
in partial fulfillment of the requirements
for the degree of Master of Science
in Library Science.

Chapel Hill, North Carolina

June, 1999

Approved by:

____

Advisor

I. INTRODUCTION

I-1. Problem

Communication networks, especially those used in the infrastructure of digital libraries, are currently challenged by the latest technological developments to provide timely and coherent transmission of bandwidth intensive multimedia applications involving video conferencing, video/audio telephony or real-time video and audio products. Unlike asynchronous data, such as text, numerics, graphics and still images, these digital sources have interactive and/or time-dependent requirements that are of an isochronous nature, indicating that the data must be received by a network terminal at the same rate wherein the data was recorded, with negligible variation in the transmission rate. Consequently isochronous video and audio demand that the data be transmitted within performance bounds--normally 30 frames per second for video and 44,100 16-bit samples per second for high-fidelity audio--and that this transmission occur uniformly to facilitate synchronization and permit intelligible viewing or listening at the receiving end. Video recorded at 30 frames per second must be continuously transmitted at the same rate of 30 fps with minimal, consistent delays, since either sporadic lowering or raising the number of frames per second will cause disruption of the contents when played back.[1]

To ensure constancy in the transmission rate during playback is no easy task, illustrated by the bandwidth requirements of 30 frame per second video which, when recorded at 640 x 480 pixels for 24 bit resolution (low to moderate quality), requires 550-Kbps bandwidth at MPEG compression and 27-Mbps when transmitted uncompressed. Considering that common local area networks operate at 10 to 16-Mps, user requests for the provision of single let alone multiple isochronous streams of data can be a

[1]Marwan Krunz, "Bandwidth Allocation Strategies for Transporting Variable-Bit-Rate Video Traffic," IEEE Communications Magazine, (January 1999), p. 41-42.

tremendous network challenge.[2]

The nature of isochronous video and audio processing is particularly challenging for networks utilizing time-division multiplexing (TDM), a connectionless technique that allocates bandwidth throughout the network by fragmenting signals into packets and transmitting packets for individual signals at different times over a single connection. When the connection is used to transmit a packet, all the other stations waiting to transmit must wait their turn to be able to transmit the next packet, thus introducing some degree of delay for subsequent packets. TDM is commonly found in local area networks (LANs) and is in contrast to other methods of signal modulation used for network connectivity, such as connection-oriented, circuit-switch multiplexing, which uses a central switch to establish point-to-point connections for separate signals, e.g. telephone networks, or connection-oriented, frequency-division multiplexing used in cable television, whereby a single connection is divided into multiple, point-to-point frequencies and the different signals separated by frequency. TDM-based networks are popular and widely applied because of TDM's economical and efficient transfer of asynchronous data,[3] but the dependency upon time-dependent signaling instead of circuit or frequency modulated point-to-point connections causes variable delays as utilization increases, delays that become problematic for networks such as digital libraries where a substantial requirement is the fulfillment of isochronous multimedia requests.[4]

I-2. Research Question

---

[3]The exception in the TDM family is Asynchronous Transfer Mode (ATM), an expensive, complex and connection-oriented approach that operates at gigabit speeds and is highly effective for isochronous multimedia due to the fragmenting of its signals into miniature, fixed-size packets called cells. Notwithstanding its name, however, ATM is not cost-effective for time-insensitive asynchronous data and thus has limited applicability to the general environment of LANs, being more suited for wide area network (WAN) connections.

Although time-division multiplexing operates with an inherent utilization/delay ratio, there are variations in the TDM paradigm at the data link layer, known as the media access protocol, that can optimize performance involving isochronous multimedia. The most common protocol variants used in network architectures are the statistical-based contention of Carrier Sense Multiple Access (Ethernet) and the deterministic-based token-passing of token ring. For both media access protocols the primary network performance parameters with respect to isochronous video and audio are (1) the mean transmission rate as a function of the playback rate or throughput, with the mean throughput rate being equal to or less than the playback rate; (2) the delay in transmitting each packet and the variance in this delay, e.g. latency and jitter, respectively, with acceptable values for latency ranging from 20 to 400 milliseconds and that of jitter being less than 80 milliseconds; and (3) the rate of lost packets due to buffer overflow or media protocol factors, with acceptable values ranging from 0.01--0.001.[5] Based upon these performance parameters of throughput, message delay as defined by latency and jitter, and packet loss, an excellent research question to explore is which type of TDM media access protocol, Ethernet or token ring, most closely meets the above defined requirements needed to effectively process the isochronous multimedia loads demanded of the LAN infrastructures in digital libraries?

I-3. Outline

To answer the above research question it will be necessary to compare the strengths and weaknesses of both Ethernet and token ring. This comparison will first be treated in a theoretical fashion, providing (1) a literature review covering network evaluation, multimedia requirements in networks, and LANs applications in electronic

---

[5] The 3-part schema of throughput, delay, and loss rate are specified in Krunz, "Bandwidth Allocation Strategies", pp. 41-42; the performance ranges for each part are partially modified for greater stringency by the data in Ismail Dalgic, William Chien, and Fouad A. Tobagi, "Evaluation of 10Base-T and 100Base-T Ethernets Carrying Video, Audio and Data Traffic," IEEE Proceedings: 19th Conference on Local Computer Networks, 1994, p. 1094-1096.

and digital libraries; and then (2) developing a primer of local area network terminology and concepts, such as cabling, topologies, and media access protocols in general; afterwards (3) focusing on the pros and cons of Ethernet and token ring individually and (4) in relation to each other as regards isochronous multimedia transmission. Preliminary comparative conclusions about the two media access protocols will be propounded, including recommendations for overcoming the weaknesses of each access protocol.

I-4. Project

Complementing the theory will be a practical component of tests comprised of a computer program coded in C++, designed around an integrated set of network performance analytical tools, that permits the comparison of Ethernet and token ring networks at an optimized state to best determine which is most appropriate for isochronous multimedia. The structure of the application is as follows:

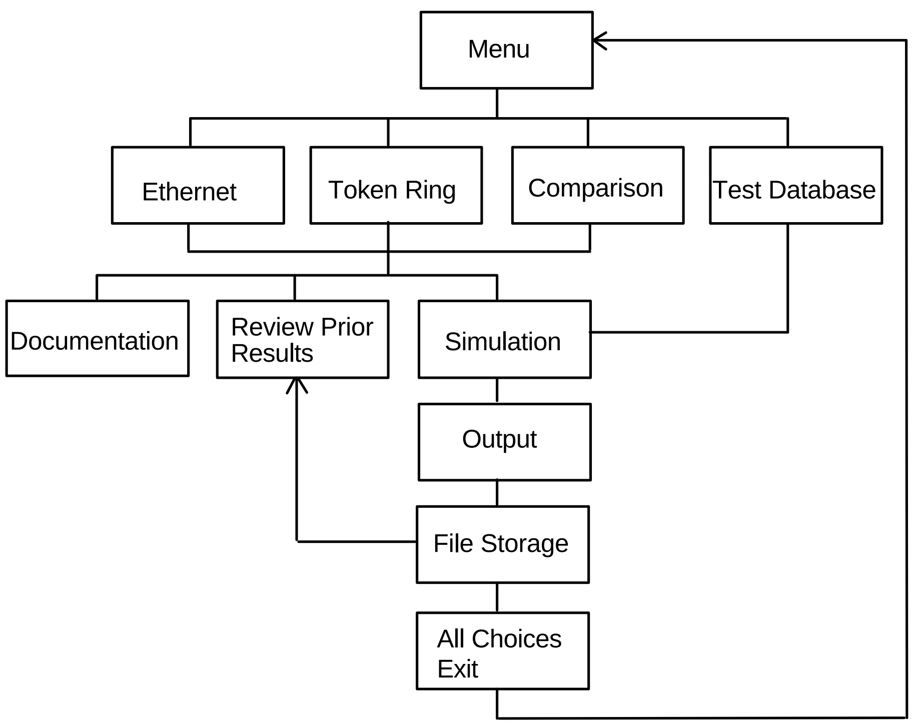

The performance parameters for the network evaluation software are:

1. The maximum achievable data rate of each node under heavy conditions, with heavy conditions defined as all nodes attempting to transmit simultaneously.
2. The message arrival rate per node to measure when network saturation occurs (with subsequent loss of packets).
3. The average message queuing delay.
4. The average number of data bits in the output buffer of a node.
5. The minimum, mean, and maximum expected latency and jitter.
6. The estimated packet loss rate, as a probabilistic function of the collision detection mechanism and buffer overflow in Ethernet or merely buffer overflow in token ring.
7. The maximum number of nodes the network can support when the network is at a certain level, determined as a factor of the average message arrival rate.

The algorithms used to define and calculate these parameters are described in the source code contained in Appendix A and stored with the attached software.

This evaluation will be accomplished in an optimized mode, which is the maximum attainable packet rate achieved when all network nodes simultaneously attempt to transmit one packet. This optimized mode assumes the following conditions:

a) The network is homogeneous, with all nodes having the same properties and average message load.
b) All nodes have exponentially distributed message interarrival times.
c) Time of message arrival to individual nodes on the network are statistically independent.
d) All packets are the same size.
e) A node can only transmit one packet each time it has access to the network.
f) All models assume steady-state operation.
g) All nodes have equal (non-prioritized) access to the network.
h) There is no idle time on the medium.
i) All transient events such as lost token, network reconfiguration, transmission error and recovery, and excluded from the analysis.

The results of the optimized mode will be stored in a file for comparison, with these findings and the preliminary, theoretical-based conclusions forming the framework for elucidating a final determination about the suitability of either Ethernet or token ring over its counterpart for isochronous multimedia purposes.

I-5. The Digital Library

The digital library's ascendancy was foreshadowed in the 1970s by a significant evolutionary development in data communication, the expansion of the decade-old concept of networking through the use of Wide Area Networks (WANs). Previously networking, i.e. the linkage of two or more computers in some fashion to facilitate communication, had been primarily restricted to LANs serving a limited clientele. A need to interconnect these individual LANs, however, led to the development of networked networks, such as metropolitan area networks (MANs) and WANs, covering regional and even national distances. At first comprised of a single long-range communication 'backbone' WAN with select users, such as the military and scientific researchers, within a decade the potential of extended networking had been identified by commercial interests and WANs proliferated. With multiple isolated sites connected electronically to each other regardless of local, regional or even global distances, a world-wide information network emerged and in its wake brought about many new innovations. One such innovation was the precursor of the digital library, the electronic library.[6]

The electronic library was the first introduction of networks in the library setting, and manifested itself either as a LAN interconnecting the departments of an individual library for enhanced data and resource sharing, such as administrative tasks, or as public access catalog to the bibliographic citations available on an online version of the card catalog, e.g. Online Public Access Catalogs (OPACs), which often were connected to a larger state or regionally-operated WAN supporting a consortium of cooperating libraries. The availability of online catalogs to the public caused a surge in user demand, since the medium involved permitted greater user access than physically possible under traditional card catalogs. As computer technology continued to improve in the areas of processing/storage capabilities, particularly with the introduction of the personal computer in the 1980s and common acceptance of the Internet in the 1990s, libraries began to respond to this increase in demand by extending the benefits of speed, organizational capability, and access afforded by the use of networks beyond mere administrative support or bibliographical citations into digitalization projects for their collections. The convergence of a variety of innovations began to foster yet another technological development, the digital library.[7]

All digital libraries are electronic libraries; not all electronic libraries are digital ones.[8] Unlike its predecessor, the mature digital library is a repository of collections stored entirely in digitized format and composed of all types of media, whether instances or combinations of text, images, video, or audio; and these diverse formats are accessed entirely through information transfer via networking, whether on-site or at a remote location. By design digital libraries are uniquely situated to provide access to not only conventional information of data or text, but also information previously considered to be difficult to access because of lack of geographic distance, format or rarity, such as historical collections, regional depositories, maps, architectural drawings, or rare books; or information with a high rate of change or currency-dependent value, such as scientific journals and news articles.[9] There has even been an extension to non-traditional, ephemeral collections, such as those generated from wire services and list-serves. Moreover the networked environment whereby the digital library subsists and interconnects with other digital libraries permits its users to work cooperatively with each other even when separated by extreme geographical distances.[10] In all cases the digital library's electronic collection requires immense storage, computational, and transmission capabilities, tasks accomplished by the network infrastructure..

I-6. The Art of Digital Library Design

Notwithstanding a digital library's utilization of precise, high-technology tools in acquiring, managing, and distributing its electronic collection, such as optical scanning of documents, hyper-linked object-oriented indexing, stratified database design, and

---

[7]Jack Kessler, The International Dimension, Artech House Computer Science Series, (Boston, MA : Artech House, 1996), p. 20-24.

[9]Nabil R. Adam, Milton Halem, and Shamim Naqvi, "Promising Research Directions In Digital Libraries," in Digital Libraries: Current Issues: Digital Libraries Workshop DL '94, eds., Bharat K. Bhargava and Yelena Yesha, Lecture Notes in Computer Science Series, no. 916 (Newark, NJ : Springer, 1994), p. 22.

extremely fast connectivity between workstations, the process of designing the network for a digital library is much less exact--indeed, in many ways it is an art. For the mere fact of technology's scientific provenance does not ensure that a particular technology can in practice be applied scientifically, since the extensive range of available equipment and applications, as well as the uncertainty introduced by new products and services, may often make it impossible to precisely determine which technology and in what form is best. Although the science behind the technology should enable one to design an 'optimal' network for each particular application, in practice, the analytical tools for networks in practice provide at best an estimate of performance and tend to be so theoretical in nature as to be of little use to most information professionals involved in selecting a network for a digital library.[11] Thus the field of digital libraries is one still very much in flux, making the usefulness of any technological standards limited in effectiveness to mere guidelines.

Furthermore the inability to adequately anticipate the response of the people for whom the system is built, whether in terms of current or future demand for a collection, forces the designer of a digital library to again estimate rather than precisely select a technology application.[12] There is also the variability in resource-demand exhibited by the actual collection, which being comprised of multimedia components requires a support system sufficiently powerful enough to efficiently process the most complex and resource-hungry media, such as isochronous video and audio. The difficulty here is that any improvement in these media is accompanied by a counter-balancing demand for greater quality, forcing the high-end threshold of technological performance required of networks ever-upward.[13] Consequently network media access protocols that were

---

[11]Dennis S. Mok, "A Software Tool for Performance Analysis and Evaluation of Local Area Networks," IEEE Proceedings: 11th Conference on Local Computer Networks, 1986, p. 77.

[13]Kessler, The International Dimension, p. 27-28; Michael Lesk, "Which Way to the Future? The Control of Scholarly Publication," in Digital Libraries: Current Issues, p.

previously deemed superior for most applications are relegated to being merely acceptable, even inferior, when considered for significant multimedia purposes such as digital libraries. For these reasons, the selection of a network in a digital library is combination of art and science, with the latter's analytical tools facilitating an approximation of performance but the experiential-based wisdom of the former ultimately dictating the final decision.

I-7. The Scope of This Paper

The focus for this paper's comparison of Ethernet versus token ring is local area networks in general and specifically when incorporated into the network design of a digital library. A comparative question of the 'best' media access protocol is especially important to digital libraries since, due to their intensive use of a mix of asynchronous and isochronous medias, the quality of a digital library's collection is dependent on not only the original source but also upon the local network infrastructure providing electronic access to its users with delays and lost packet rates consistently held within acceptable bounds. Moreover the favorable identification of one media access protocol over another would be a long-term benefit in terms of standardization, since standardized technologies tend to aid the diffusion of the technology and widespread adoption leads to lowering of per unit costs and ensures future support by vendors in the marketplace. Lastly, any effort to standardize network technologies through comparative evaluation will aid the reduce the proportion that digital library design is dependent upon a nebulous "art". In providing a software application, data and conclusions contrasting one media access protocol over another, this paper assists future designers of digital libraries, as well as multi-purpose networks carrying video or audio traffic, in making intelligent selection decisions that effectively match the capabilities of a media access protocol to its setting.

36.

II. LITERATURE REVIEW

II-1. Terminology

Before perusing the selected articles in the literature review, it would be helpful for the sake of comprehension to define some key terms, such as:

Bandwidth: the maximum capacity of the network media as expressed in megabits per second, e.g. a 10-Mbps Ethernet network has a bandwidth equal to 10,000,000 bits or 1,250,000 bytes per second.

Jitter: variable delays in the transmittal of packets caused by fluctuations in the data flow, jitter can be caused by increased utilization, signal distortion, collisions, errors, loss of packets, or intermittent "bursts" of activity that consume bandwidth. The intensity of the jitter measure is relative to the latency value from which it is derived at any particular point in time, e.g. if the jitter is a large value, such as 150ms, it doesn't matter whether that the average latency is only 20ms or 400ms, the jitter of 150ms will disturb isochronous transmissions; for this reason, it is better to have a constant moderate level of latency with low variance than delays that are in flux. Thus for multimedia applications both latency and jitter must be within acceptable bounds.[14]

Latency: also called response time, fixed delays, or average message transmission delay, latency is the time it takes, on the average, for the first bit of a message packet to travel from the sending to the receiving node, with the network's queuing model contributing in part to the latency measure. Latency can be described in various ways, such as (1) total system time, (2) total network service time, which arises from the signal propagation speed and media access protocol, (3) and the difference between the two, service waiting time, which is comprised of collision rates, error rates, queuing models, and network topology. Out of the three, service waiting time contributes most to overall delay since service capacity, dependent upon random and/or external factors, is the most variable.[15] For multimedia data, the latency experienced by transmitted packets must be within minimum and maximum bounds to prevent, or at least discourage, loss of fidelity.

Signal Propagation: the speed at which data is transmitted from node-to-node along a network physical layer. Electrical signals travel at the speed of light in a vacuum but their progress is slowed by the characteristics of the physical transmission media, whether it be copper or fiber cable, or the air. The actual speed of the electrical signal when considered in relation to its transmission media is the propagation speed.

Topology: the physical layout of the connections from node to node comprising the network. Common approaches are the linear bus, the centralized star, and the circular ring.

Throughput: also called channel utilization, throughput is defined as the amount of useful data that is transmitted in a duration of time, expressed in megabits per second (Mbps). The throughput for a given network at a particular time is dependent upon many factors, including maximum available bandwidth or channel capacity, for example 10-Mbps Ethernet or 16-Mbps token ring, number of active nodes, and size of the message packets. Overhead is also restrictive, since such factors as control information, packet header size, and retransmission due to error or collisions (Ethernet only) can rapidly detract from total throughput. To avoid significant distortion of video, audio and/or conferencing data, the network throughput must on average be of sufficient capacity to enable the data stream to arrive at the receiving node at essentially the same rate wherein the media was originally recorded.

Utilization: percentage of the bandwidth that is in use at any particular time or for a duration of time, expressed as a measure of the difference between maximum bandwidth capacity and the total load on the network.

II-2. Network evaluation

The compilation below represents only a small sample of the theoretical studies pertinent to network comparison and real-time or isochronous transmission requirements. The articles, arranged by publication date, were selected to be representative of the comparative essentials of CSMA-CD protocols such as Ethernet and token rings, in general and specifically addressing isochronous applications.

[15]For this paper, latency is considered to be the combination of network service time and waiting time. See Martin Nemzow, The Ethernet Management Guide, 3rd. ed., McGraw-Hill Series on Computer Communications, (New York : McGraw-Hill, 1995), p. 439.

Bux, 1981

In an informative study of Ethernet and token ring,[16] Bux provides a comparison of these media access protocol variants based upon the measure of average packet latency. Defined as the average time required to transmit a message as determined by throughput, a low latency is emphasized in this study for its real-time necessity. The concern for minimal delay rates anticipates the introduction of isochronous multimedia applications and is in contrast to similar studies of the period, which stress bursty, bulk-transfer performance instead of the constant data streams particular to video and audio message packets. Judging by Bux's findings, the use of latency as measure of performance favors token ring to the detriment of Ethernet, although the specification of shorter packet lengths on a CSMA-CD network can alleviate this deficiency.

Liu, Hilal, and Groomes, 1982

This study by Liu, Hilal and Groomes,[17] evaluates the performance of a variety of media access protocols, from the statistical contention protocols of ALOHA,[18] CSMA, CSMA-CD, to demand-assignment or deterministic protocols of token passing ring, slotted and register-insertion.[19] Using analytical models of throughput, response time, i.e. latency, and reliability, a simulated performance of the various protocols determines (1) that the deterministic protocols, especially register-insertion and secondly token passing, surpass in terms of throughput their contention-based counterparts for real-time applications; (2) that the former are also superior in terms of maximum worst-case

---

[17]Ming T. Liu, Wael Hilal and Bernard H. Groomes, "Performance Evaluation of Channel Access Protocols for Local Computer Networks," Proceedings COMPCON, Fall '82, (September 20-23), p. 417-426.

[19]Token passing employs a unique packet frame, called the token, that circulates throughout the ring as each node passes the token on to its neighbor. When a node wants to transmit, the token is removed and a message packet placed upon the ring. Once received, the message packet is again transformed into a token. While similar, slotted and register-insertion media access protocols differ from token ring primarily in utilization of the ring bandwidth, with slotted ring offering access to multiple nodes but at fixed packet sizes, while the more efficient register-insertion allows for multiple access at variable packet sizes

transmission delay when operating at high loads; and (3) that CSMA/CSMA-CD with its bus topology provides greater fault tolerance or reliability than ring loops. Overall it is the opinion of the authors, supported by their findings, that the group of deterministic protocols in general provide higher performance in terms of latency at moderate to high utilization.

Mok, 1986

Mok[20] offers a pragmatic network performance analysis software program that compares CSMA-CD bus systems, e.g. Ethernet, to other bus and ring systems using token passing, such as Arcnet and Ringnet (token ring). Designed to aid selection of a network to fit an application, Mok utilizes a series of analytical tools that enable performance characteristics of each system to be identified and evaluated under similar conditions. While non-conclusive for the most part, being devoted to elucidating the performance measures and logical structure of the software, this paper does indicate that Ethernet has a higher data throughput level when using long packets of over 1000 bits, while Ringnet achieves this performance when the bits per packet are from 10-1000 bits. Since minimization of latency is related to shorter packet sizes, the evidence is again biased against Ethernet for multimedia purposes.

Suda and Bradley, 1987

This paper by Suda and Bradley,[21] examines the effect of data traffic on voice transmission performance by a token ring. First a general distinction is made between the types of data traffic, being either (1) interactive and of a "bursty nature", consisting of short messages that demand low latency values; or (2) bulk data, which is comprised of long messages that are more bandwidth intensive and do not require bounded delays. In both cases errors must be negligible to ensure accurate replication of the data source.

---

[21]Tatsuya Suda and Tracy T. Bradley, "Packetized-Voice/Data Integrated Transmission on a Token Passing Ring Local Area Network," IEEE Proceedings: 12th Conference on Local Computer Networks, 1987, p. 238-244.

Secondly, these qualities of data traffic are compared to the words, sentences, and pauses of voice traffic, where the messages vary in length, possibly consuming large quantities of bandwidth, and are highly interactive, with periods of bursts punctuated by silence between transmittals. Due to the interactive nature of voice traffic, the latency must be kept within reasonable bounds, typically estimated at 150-200 milliseconds, and the total packet loss rate is less strict, being less than 1%, since the intelligibility of voice messages are not affected by a small proportion of lost packets. Lastly, the simulation model used for the study is outlined and then a series of tests are performed, with the resulting data indicating that token passing in a ring network can simultaneously process both data and voice traffic within the necessary performance constraints. In summary, the authors state that token passing provides lower overall latency and jitter values than other network designs, especially for high utilization applications since the operational delays of token passing do not exhibit as much sensitivity to fluctuations in traffic loads.

Sato, Nakada, and Sato, 1988

An elucidation on a new coding method for voice traffic and the analytical tools needed to conduct network delay analysis,[22] Sato, Nakada, and Sato's work focus on ATM networks limits the applicability of their study but their outline of the components to network delay is of universal value. Total node-to-node latency consists of:

Coding and packetizing
Nodal processing due to computational time incurred in preparing packet
Queuing or reception at the nodal buffers
Propagation speed due to cable characteristics and topology
Delay smoothing and de-packetizing

The necessity of coding before packetizing voice data and, for that matter, video and audio, arises from the tremendous bandwidth requirements of uncompressed multimedia, a characteristic that has made the reduction of the bits needed to transmit such data not only an attractive goal but a widely implemented one as well. When coding the data, however, there is introduced an intrinsic delay of about 20ms, which must be accounted for in the total latency requirement for message transmittal; moreover there is additional latency, variable in nature, incurred by the decoding/delay smoothing process upon receipt of the message packets. Therefore, when evaluating network performance, the time involved in preparing isochronous data for transmission and reception must be considered if potential latency is to be accurately estimated.

Gonsalves and Tobagi, 1989

This comparative simulation by Gonsalves and Tobagi,[23] examine the performance of voice and data traffic on bus topologies using CSMA-CD, e.g. Ethernet, and token passing, e.g. Token Bus and a folded bus variant, Expressnet. Although the study is restricted to buses, the authors state that the findings for Token Bus and the ring-like Expressnet are indicative of the performance of ring topologies employing token passing, such as token ring. The parameters set forth for effective voice transmission are an end-to-end latency of 10-20ms for telephone quality speech and upwards of 200ms for network conferencing, with acceptable packet loss values of 1-2 percent. Voice encoding delay is estimated at a maximum of 20ms. Based upon these requirements, the performance measures of throughput, latency, jitter, and packet loss are evaluated for each bus type, and in the comparison it is clear that, notwithstanding similar performance, Ethernet is less capable for voice traffic than Expressnet under the same conditions.

In regards to throughput, Expressnet achieves a higher total data throughput than either Ethernet or Token Bus regardless of the number of nodes, with Ethernet coming in a close second. Although the lowest total latency is produced by Ethernet, since the number of collision-less transmissions at low loads compensates for contention

---

[23] Timothy A. Gonsalves and Fouad A. Tobagi, "Comparative Performance of Voice/Data Local Area Networks," IEEE Proceedings: 14th Conference on Local Computer Networks, 1989, p. 657-669.

delays at higher loads, the variation in delay, i.e. jitter, favors Expressnet over Ethernet, which actually improves in constancy of response time as higher utilization occurs. Moreover while the packet loss value is within bounds for all bus types, the manner in which the losses sporadically occur in Ethernet, varying in the number of lost packets, is more detrimental for end-use performance than the regular, frequently occurring packet losses in the token passing protocols. For their conclusion, Gonsalves and Tobagi assert that Ethernet is preferred for networks with low signal propagation delays and handling low loads, since the overhead of token passing impairs its efficiency under these conditions. In the case of heavy utilization or poor signal performance caused by topology or cabling, token passing especially as configured in Expressnet is the preferred model.

Jeffay, Stone, Talley, and Smith, 1992

An innovative approach to the problem of providing high-quality isochronous data is supplied by Jeffay, et al,[24] wherein an existing interconnected network comprised of ethernets, token rings, and Fiber Distributed Data Interconnect (FDDI) rings, is evaluated for real-time conferencing using a new suite of software techniques designated MTP (Multimedia Transport Protocol). Designed to maximize performance and minimize latency even during congestion caused by high utilization, the transport protocol approach of MTP shifts the burden of multimedia transmission performance from the hardware to computational manipulation of the data at the sending/receiving processors. Central to MTP is the successful use of variable synchronization of audio/video streams as needed to compensate for latency and variations in this delay or jitter; furthermore large buffers are employed at the receiving nodes to modulate the flow of incoming data and thus smooth out any discrepancies in transmission rate.

Under simulated isochronous data transmission, the integration of MTP into the transport layer effectively reduces the differences between the 3 disparate media access protocols comprising the internetwork, resulting in a homogenized data flow achieved at only moderate cost. While not explicitly stated by the authors, an implicit conclusion from this study is the use of similar software approaches for an individual Ethernet or token ring to improve its capability to handle isochronous applications.

Amer, Christensen, Toher, 1993

This experimental study by Amer, Christensen, and Toher,[25] and corroborated by the later work Bisdikian, et. al.,[26] enumerates the advantages of carrying multimedia traffic over 16-Mbps token ring, primarily those of (1) a determined and low-variance latency due to the predictability of the token passing operation, and (2) the activation of a priority mechanism indigenous to token ring that permits higher priority of token access to be assigned to isochronous data, thus enabling bandwidth to be reserved for immediate and constant transmission.[27] Although the use of a priority mechanism is not restricted to token passing networks, the capacity to reserve bandwidth is an advantage that Ethernet's contention system lacks.[28] Simulating a token ring of 40 nodes covering 4 kilometers, and assuming multimedia traffic being generated by 12 stations at 1.2--Mbps each, Amer and his colleagues discovered that using priority schemes enabled effective multimedia transmission with delays for each message packet held to under 5ms at over 90% utilization. In contrast, without assigning priority the latency increased sharply to near 20ms at the same usage level. Furthermore the delays experienced by

---

[25]Khaled Amer, Ken Christensen, and Tora Toher, "Experiments with Client/Server Multimedia on Token Ring," IEEE Proceedings: 18th Conference on Local Computer Networks, 1993, p. 2-6.

[27]It is noteworthy that, while the datalink layer of token ring and similar systems offers priority capability, the integration and activation of these priority schemes have more often than not been neglected at the upper network layers, specifically the applications driving the message transmissions. See William J. Cronin, Jerry D. Hutchison, K.K. Ramakrishnan, and Henry Yang, "A Comparison of High Speed LANs," IEEE Proceedings: 19th Conference on Local Computer Networks, 1994, p. 48.

the isochronous streams at high priority were unaffected by occurrences of "bursty" data, illustrating how priority mechanisms in effect isolate one data structure from another.

Cronin, Hutchison, Ramakrishnan, and Yang, 1994

Addressing the problems that low-speed LANs such as 10-Mbps Ethernet and 16-Mpbs token ring encounter in a network environment replete with multimedia applications,[29] this paper by Cronin, et. al., describes and contrasts four high speed network options that perform at 100-Mbps or greater. The technologies covered by the authors are all based upon common and inexpensive Unshielded Twisted Pair (UTP) cable, and include an Ethernet upgrade of Fast Ethernet (100BASE-T), the hybrid 100VF-AnyLAN,[30] a token ring upgrade of FDDI, and the gigabit, cell-based ATM.

Using the comparative factors of implementation cost and latency control, 100BASE-T and VG-AnyLAN are, due to their simpler media access protocols, more inexpensive to install than either FDDI or ATM; in regards to latency, however, 100BASE-T is at a disadvantage since it alone lacks a priority mechanism for allocating or reserving bandwidth. Using switches with 100Base-T can compensate for this lack, since switches connect/disconnect each node from the connecting media in response to demand, creating a series of point-to-point links with each node having access to the full bandwidth capacity. In contrast a shared network has the total bandwidth divided among all the nodes, reducing the available bandwidth at any particular time to a proportion of the total. All four network technologies excepting ATM allows for either shared or switched performance, with FDDI in the authors' estimation outperforming overall its counterparts in handling isochronous data. The size of the network is also a consideration, with networks of less than 200 meters being adequately served by Fast Ethernet and AnyLan, for distances of up to 2 kilometers FDDI's accelerated token passing protocol achieves the highest efficiency, and for larger, regional distances ATM

[29]Cronin, "A Comparison of High Speed LANs," p. 40-48.

is the best contender as a multimedia backbone network.

Dalgic, Chien, and Tobagi, 1994

In this simulated comparison of 10Base-T and 100Base-T Ethernets,[31] Dalgic, Chien and Tobagi investigate the performance of hub-based CSMA-CD topologies for supporting isochronous data streams. Implicit to the author's significant research is an interest in maintaining Ethernet as a viable network architecture even under increasingly demanding multimedia applications, since Ethernet status as the most popular and widespread type of network used today affords advantages in terms of market support and infrastructure dependence that its millions of users have come to rely upon, and that only costly expenditures in other technologies could replace. The efficient transmission of isochronous traffic, however, requires that the packets be delivered within user-acceptable delays, bounds that Ethernet cannot guarantee due to its contention protocol and lack of a priority mechanism. One approach explored by Dalgic and his colleagues to overcome the inherent limitations of CSMA-CD is to restrict the number of nodes that may at any one time transmit multimedia streams.

To evaluate the performance of Ethernet/Fast Ethernet, simultaneous transmissions of "bursty" traffic and isochronous stream traffic are simulated, assuming latency values stringently bounded at 20ms and at a relaxed 100ms,[32] delays suitable for all types of isochronous data including time-sensitive interactive conferencing applications. Packet loss rate values, caused by either exceeding the specified delay bounds or through the contention process, are selected at 0.01 and 0.001. The results indicate that 10-Mbps Ethernet, with a pre-existing data load of 500-kbps, 20ms delay bound, and 0.001 packet loss rate, can handle up to 11 video streams of 384-kbps each; at the more relaxed 100ms bound and loss rate of 0.01, a total of 18 video streams can

[31]Dalgic, "Evaluation of 10Base-T Ethernets Carrying Video, Audio and Data Traffic, p. 1094-1102.

be supported. As the data traffic becomes heavier or experiences "bursts", less isochronous streams are supported and their packet sizes must become smaller to meet the latency requirements. For 100-Mbps Ethernet, with a pre-existing data load of 15000-kbps (15-Mbps) under stringent conditions and transmitting video packets of an impressive 1536-kbps each, the number of supported streams is 27.[33] Under relaxed conditions and dedicated isochronous use a total of 52 video streams at 1536-kbps are possible.

Deriving their conclusions from these results, the authors hold that Ethernets are capable of handling multimedia-intensive applications within user-acceptable parameters as long as the number of isochronous nodes are kept within a minimum and the data traffic activity is regulated so that the total load and the size of the bursts does not impair the isochronous streams already in service. Moreover the use of streams that have variable packet sizes that fluctuate to correspond with the intensity of the data traffic facilitates achieving transmission within the required latency and loss rate bounds. Implemented under these conditions, Ethernet and especially Fast Ethernet remains a viable network architecture in a rapidly changing multimedia environment.

Elsaadany, Singhal, and Liu, 1995

Elsaadany, Singhal and Liu,[34] examine the effect of network configuration upon multimedia applications, especially the difference between a shared and switched network environment. Acknowledging that isochronous transmission requires extensive bandwidth resources and low latency, the authors evaluate the extent to which switching hubs can increase the effective bandwidth and minimize delays on an Ethernet network. In switching a single LAN is separated via the switches into individual network domains, either as multiple nodes to a switch to retain a shared bandwidth schema, or as a single

[33]In contrast, 1536-kbps video can be transmitted on 10-Mbps Ethernet at a maximum of only 5 streams, and without any other data load upon the network.

node per switch port creating a dedicated bandwidth hub-to-node connection. In the course of their simulation involving a network of 32 nodes, Elsaadany and his colleagues report that response time drops 50% from 8ms to 4ms with the introduction of only 4 switches, or 8 nodes per switch; and that effective bandwidth increases proportionally to the number of switches, so that data throughput achieved at 80 percent utilization is accomplished at only 20 percent utilization with 4 switches, providing a savings of 60% in available bandwidth. Because of these gains in throughput and reduction in latency, the use of switched over shared bandwidth is clearly a suitable technique to accommodate the intensive isochronous demands of multimedia.

III. LOCAL AREA NETWORK CONCEPTS

III-1. Cabling

In every network there exists at the physical layer some means of connecting media whereby workstations at the nodes are enabled to communicate with one another. Normally this media is physical cabling but it can be via radio or microwave frequencies to earth-bound or satellite receivers. Although it plays a humble role in the operation of a network, quality cabling appropriately chosen for a particular application is vital to overall performance of all network components. So critical is cabling for network function that the legacy networks, such as Ethernet, Arcnet, and token ring, were originally designed for particular types of cabling at unique specifications. While there still exists a need to use the right cabling, innovations in technology have expanded the once limited options for network cabling configuration. Today, factors considered in selecting network cabling include network topology, type and speed of data transmission, cost, ease of installation, and future ability to add more nodes, or scalability.

III-1-1. Twisted-Pair

Commonly found as the connecting medium in telephone systems, twisted-pair cabling has become popular for network connectivity due to its inexpensive and malleable properties, especially as concerns installation. Fabricated from two, insulated copper wires that are twisted together to reduce noise, or cross-talk, when an electrical current passes through, twisted-pair wire is an effective medium when network performance specifications are less exact and can deal with a modicum of signal degradation.

Figure 1. Unshielded Twisted Pair Wire

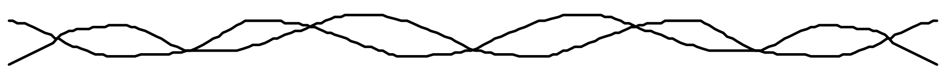

This first type of twisted-pair is unshielded twisted-pair (UTP), but there is a second type known as shielded twisted-pair (STP), which uses additional insulation to interference from external electromagnetic forces and thus is better suited for greater transmission distances. The concern for signal distortion arises out of the nature of copper cable itself and its limitations as a communication medium, such as:

<u>Attenuation</u>: an indication of a loss of signal strength, with signal intensity decreasing in proportion to the distance traveled and in relation to the propagation speed of the medium. To reduce attenuation, repeaters are used at periodic intervals to refresh and retransmit the signal.

<u>Capacitance</u>: the amount of residual energy that is retained by a cable, with longer cable lengths causing higher capacitance values. Capacitance can interfere with transmission by impeding processing of the signal.

<u>Impedance</u>: delay that occurs due to the resistance generated by the cable at certain signal frequencies.

<u>Noise</u>: distortion of the signal caused by external sources, such as electrical machinery or fluorescent light ballasts. A common example of noise is cross-talk, which is

experienced on the telephone as hearing multiple conversations and occurs because of electro-magnetic interference from the twisted-pair wires themselves.[35]

The primary problem with twisted-pair cable is cross-talk and a narrow bandwidth capability caused by poor signal performance, both areas being of concern for high-speed networks.[36] Despite its failings, the advantages of twisted-pair have led to its widespread adoption, including the development of five standardized types of UTP network cabling:

Category	Definition
1	Telephone wire
2	Four twisted-pairs that support speeds up to 4-Mbps.
3	Supports speeds up to 10-Mbps for Ethernet and 4-Mbps token ring.
4	Designed for 16-Mbps token ring.
5	Supports speeds up to 100-Mbps.[37]

III-1-2. Coaxial

In an attempt to avoid the inherent drawbacks to twisted-pair and improve transmission performance, coaxial cable was developed using two cylindrical copper conductors, one in the core and another externally, with an insulating dielectric separating the two and the entire unit covered by an protective jacket.

Figure 2. Coaxial Cable

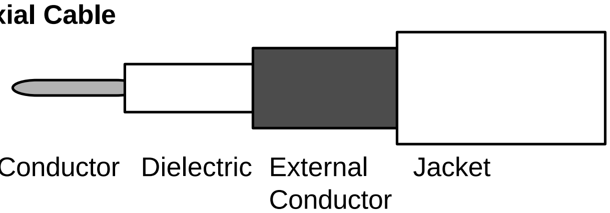

The result is a sturdy cable that is relatively immune to electromagnetic noise and provides excellent signal transmission, i.e. low cross-talk, attenuation, etc., across long communication spans. Consequently the signal propagation speed, or the rate at which data is transmitted, is extremely high, enabling wider bandwidths or broadband potential. Broadband is the multiple transmission of simultaneous messages, as compared to

---

[35]Joseph R. Matthews and Mark R. Parker, "Local Area Networks and Wide Area Networks For Libraries," Library Technology Reports vol. 31, no. 1 (January-February, 1995), p. 20.

[37]Matthews, "Local Area Networks," p. 19.

baseband media such as twisted-pair which handles single messages only. There are disadvantages to coaxial cable, such as its bulk and rigidity, which complicate installation, as well as its higher cost.[38] Furthermore a disconnected but still open tap in the cable can cause a length of coaxial cable to act as a mile-long radio receiver.[39] Despite these weaknesses it is due to coaxial cable's greater performance potential, especially for intensive data, video, and audio applications, that has made this cable the connecting medium most often used in networks.[40]

III-1-3. Fiber Optic

Abandoning the copper conducting base for concentrated laser beams, fiber optic cable is no light-weight connecting media. Fiber optic media is created around a slender glass or plastic core with a high refractive index, and surrounded by a glass sheath of a low refractive index. The core and its sheath are then covered with a protective jacket.

Figure 3. Fiber Optic Cable[41]

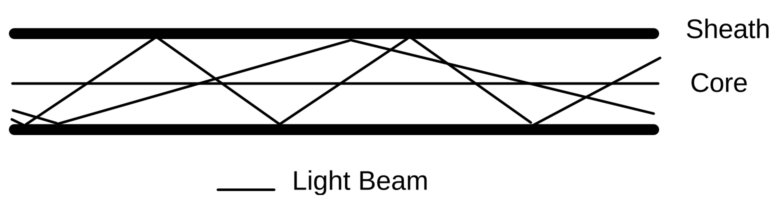

When there is only a single core, the optic cable is a monomode fiber; in contrast, when several of these cores are bundled together in their protective enclosure, it is a multimode fiber. Digital data is converted at the sending end into pulses of light by a laser or light-emitting diode, the pulses are transmitted at a slight angle against the boundary of the core and its sheath throughout the length of the optic core, and at the

---

[39]Cited from a class lecture. James Gogan, INLS 184-01, Network Protocols and Management, School of Information and Library Science, (Chapel Hill, NC : University of North Carolina at Chapel Hill, Spring 1999).

[41]Ferrero, The Evolving Ethernet, p. 46.

receiving end the pulses are translated once again into data.[42] While not strictly a cable in traditional sense, fiber optic cable fulfills all network cabling requirements: broadband capability, fantastic bandwidth permitting data transfer upwards of a billion bits per second, minimal signal interference, and security from unwanted access.[43] These gains are achieved at a cost, however, in terms of manufacturing expense, since the optic cable is engineered to high-quality specifications, and in regards to installation, which is complicated by the need to often use expert installers and sophisticated equipment. Continued reductions in cable cost per foot, optic components such as lasers, as well as recent innovations in installation, have alleviated some of the above disadvantages and subsequently fiber optic cable has enormous potential for future network applications, especially for high-speed networks operating at over 100-Mbps per second.[44]

Figure 4. Comparison of Twisted-Pair, Coaxial and Fiber Performance[45]

	Category 3 UTP	**STP**	**Coaxial (10Base2)**	**Multimode Fiber**
Impedance	100 ohms	150 ohms	50 ohms	----
Propagation Speed	0.4 c	0.6 c	0.65 c	0.67 c
Attentuation	< 70 dB/km	< 40 dB km	< 33 dB km	< 0.8 dB km
Noise	> 45 dB	> 70 dB km	----	----

III-2. Topologies

The logical arrangement of a network's hardware components is its essential structure, or topology. Forming part of the physical layer along with the connecting media, the topology defines the manner wherein the nodes access and communicate on the network. Since a given topology can be crucial to network performance, it is important to match the right topology to the anticipated demands that a network will experience once in operation. The two primary topological types are (1) logical

---

[43]Kim, "Local Area Networks," p. 16.

[45]Data for UTP, STP and coaxial cable courtesy of Heinz-Gerd Hegering and Alfred Lapple, Ethernet: Building a Communications Infrastructure, translated by Stephen S. Wilson, Data Communications and Networks Series, (Wokingham, England : Addison-Wesley, 1993), p. 59 and 113; data for multimode fiber from Ferrero, The Evolving Ethernet, p. 130.

structures based on point-to-point links, e.g. the ring and star, or (2) those permitting multiple access, e.g. the bus and the tree.

III-2-1. Star

A logical structure using a central controlling element with all the nodes radiating off from this location in centrifugal fashion is the star topology. Point-to-point links connect each node to the central element, or the hub, which controls access to the connecting medium and between all the nodes through a polling mechanism.

Figure 5. Star Topology

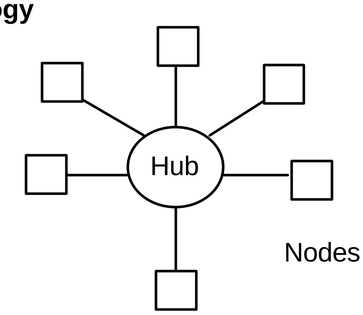

An asset to the star topology is those point-to-point links to the hub, permitting both the equal distribution among the nodes of shared medium systems or the allocation of maximum bandwidth on a per node basis through the use of switches. Moreover the operating standard for most topologies is half-duplex, indicating that only one type of signal, either receiving or transmitting, can occur on one channel at an instance of time. With point-to-point links in a switched hub environment, which uses a switch to establish at any instance dedicated channels between the hub and its nodes, simultaneous two-way reception and transmission, or full-duplex capability, is possible. In full-duplex the available bandwidth is doubled, so a 10-Mbps connection under full-duplex mode is actually operating at 20-Mbps; and, because dedicated channels exist, the need for mechanisms such as carrier sense and collision detection vanishes. Despite these benefits, the dependence upon the hub is also the star's greatest potential disadvantage, since reliance upon a central focus for distribution of the signal can be a source of bottleneck or, even worse, a failure of the central component leads to total network collapse.

III-2-2. Ring

Another method involving point-to-point links but with the nodes connected in series is the ring topology, a looped network. In a ring network, the nodes have immediate communication with their upstream/downstream neighbor and each node is responsible for repeating the network signal.

Figure 6. Ring Topology

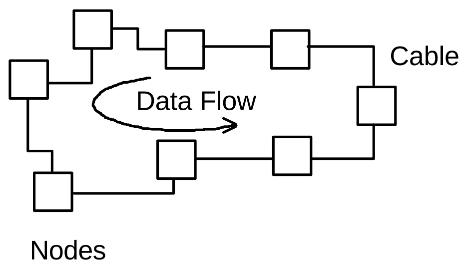

To accomplish the task of transmission, some means of regulating access is employed, normally via a token. Given the overall performance of the ring is dependent upon the individual performance of the nodes, a highly-tuned media access protocol must be used to permit efficient, distributed transfer of data. Benefits of the ring topology include handling near maximum utilization (> 90) without degradation of performance, consistent latency values at all load levels, managerial reporting based on system generation of statistics, and a high level of security. These benefits are attained at the price of complexity, in terms of software (the access scheme) and hardware (nodes double as repeaters), and introduce potential instability into the network, since each node is critical to the operation of the network.

III-2-3. Bus

The bus topology utilizes a central trunk connecting medium, with attached nodes lined up along the sides of the trunk like leaves, and a signal that propagates freely from one end of the cable to the other.

Figure 7. Bus Topology

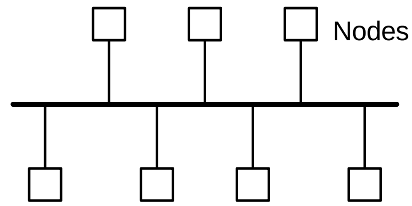

The interaction between the nodes and the connecting medium upon the transmission of a signal is transparent, as the nodes either passively receive messages addressed to them or wait for the medium to clear so another signal may be sent. Because one or any of the nodes may access the bus trunk, a media access protocol must regulate the manner in which access occurs so that signal collisions are minimal and to ensure fair use for all the nodes. The bus topology has advantages in that it permits a simple, efficient access method, the centralized trunk enables many nodes to connect to relatively short cable lengths, and nodal transparency disengages network operation from dependence upon any one node. Disadvantages of the bus, however, are difficulties accessing the network under heavy loads, resulting in disparities in distribution, and poor security due to the use of a shared medium.

III-2-4. Tree

Combining the other topologies, a tree topology is for larger network applications where the advantages of more than one topology is required or where pre-existing networks are being configured together into an intranet. Joined together by hubs and/or repeaters, which refresh the transmission signal, the tree topology allows for networks to cover greater distances than otherwise possible without discarding the benefits of a particular network type or incurring the cost of building an entirely new network. Normally the bus and star logical structures comprise a tree topology, but ring topologies can also be used.[46]

Figure 8. Tree Topology

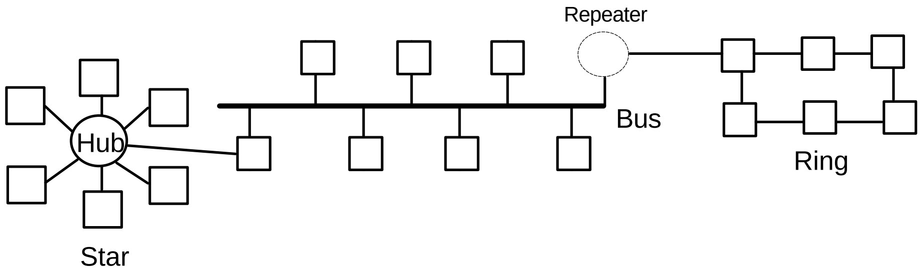

III-3. OSI and Media Access Protocols

The Open System Interconnection Reference Model (OSI) was developed by the International Standards Organization in 1977 as an attempt to introduce flexible standardization to the network design field. OSI provides a theoretical framework of network communication flow that divides the flow between various subcomponents, or layers, within the model with each subcomponent addressing a specific type of network activity. The 7 layers from top to bottom, with the lowest layer representing the physical transmission media and means of modulation, are

Layer	**Function**
Application	Higher level end-user and application purposes.
Presentation	Data transformation from bits to readable display.
Session	Oversees communication processes during session between two users.
Transport	Ensures for end-to-end integrity control to prevent loss and ensure quality of services.
Network	Handles the routing and switching of information, as well as packet assembly/disassembly.
Data Link	Transfers packets between the ends of a physical link through medium modulation.
Physical	The hardware rules that permit transparent transmission of packets as represented at the bit level.

The genius of the OSI model is that the 7 layers are viewed as an ordered set of subsystems through which a hierarchical communication protocol dominates, with the lower layers communicating to the higher layers via the intermediary layers, and vice/versa. Each layer is specialized in function and independent of the other sub-systems, and it is possible to build the OSI model without including some of the sub-systems.

The advantages of OSI are the specificity of function and independence applied to each layer, attributes that permit ease in testing the integrity of a sub-system layer, and that the hierarchical communication flow of the 7-layer open systems model provides a blueprint of network design while allowing for modifications of a particular layer. The result is that interoperability is ensured between other, dissimilar variations of sub-components or even entire networks due to variations caused by technical innovations or vendor modifications. Based upon OSI, it is conceivable to have one sub-system removed and replaced by another sub-system accomplishing the same task but in a different manner and yet not compromise network service.

Within the 7-layer framework of the OSI is the data link layer, which receives messages from the higher layers and prepares them for transmission by encapsulating the messages into frame packets. To interface with the connecting medium comprising the physical layer, the data link layer uses a controlling mechanism, the media access protocol. Media access protocols determine the manner in which network nodes obtain access to the medium so as to transmit and controls the length and relative importance of that transmission in relation to pending jobs. Essential to the integration of the physical layer medium and upper-layer network protocols, the type of media access protocol chosen for a particular network determines to a large degree the performance that will be achieved on that network. Since network access is either achieved through statistical-based contention or deterministic noncontention methods, the media access protocols can be categorized accordingly.

III-3-1. CSMA (Statistical)

The concept of Carrier Sense Multiple Access (CSMA) is derived from the Hawaiian ALOHA model of the early 1970s. One of the first networks to be developed, ALOHA was designed to connect the Hawaiian Islands through a connecting medium of radio carrier waves transmitted by local communications stations. In ALOHA, stations could transmit when the carrier was free for use and access by any particular station was governed by a probabilistic process of contention, which allocates access based on a statistical distribution. Relatively inefficient, with a maximum output of 37%,[47] ALOHA became the foundation for CSMA, which retained the idea of "carrier sense" even though for cable-linked networks there was no longer any true carrier medium.

CSMA relies upon the characteristic of each node on the network to listen to the medium, e.g. the carrier, and transmit when the channel is free. The media access protocol is "multiple access" because every node has the right to access the media at any point in time, as long as the channel is clear, with access controlled via contention. Once a node begins transmitting, the other nodes must wait for an idle period before transmitting their data. Due to contention's lack of assigned priorities or privileges, it is conceivable that two or more nodes will sense when the media's channel is free and all involved nodes will try to transmit, resulting in the collision and subsequent loss of the transmission signals. The proportion of these collisions increases in direct relation to utilization, thus impairing effective throughput at heavier loads.

Figure 8. CSMA Contention Cycle[48]

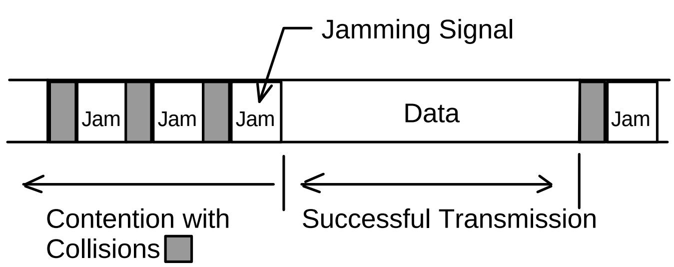

To compensate for the possibility of collisions on a CSMA network, two alternative algorithms were devised for non-ALOHA applications, such as Ethernet, that of Collision Detection (CD) and Collision Avoidance (CA). In CSMA-CD, a detection mechanism is employed that, upon a collision, clears the media by sending a jamming

---

[47]Bill Hancock, Designing and Implementing Ethernet Networks, 2nd. ed., (Wellesley, MA : QED Information Series, 1988), p. 12-13.

signal throughout the network, a signal that alerts all nodes to withhold transmission for a randomly generated holding period.[49] This holding period is controlled by binary exponential backoff, so that the first collision causes a wait of *n* seconds, the second collision *n**2, and so forth for a maximum of 16 attempts, at which point the node discards the packet. Collision Avoidance, on the other hand, has all the nodes monitoring the medium and refraining from transmitting during the moments a collision is most likely to occur. As Collision Avoidance involves greater delay periods between transmissions, its counterpart of CSMA-CD is more widely used. In both cases, however, the disadvantages of a contention media access protocol and frequency of collisions cannot be fully overcome.

In CSMA-CD, the average number of collisions in any network cycle (assuming optimized conditions) may be estimated using the following algorithm:

Example 1.[50]

1 - A / A

where,

$A = (1 - (1 / nodes))nodes^{-1}$

In C++, this formula is encoded in a subfunction as:

```
result = (1 - (1 / x)) * pow(x, -1);  // computes A as a factor of the number of
                                       // nodes (x)

result = (1 - result) / result;        // average number of collisions
return result;                         // returns the value to the "collisions" variable
```

The contention value is the product of collisions plus total detection time, i.e. the elapsed duration to detect a collision, and the delay introduced by the jamming signal transmitted at each collision, so that:

```
contention = collisions * (detect + jam);
```

Combining contention with the message transmission time, Ttran, which is the minimum time necessary to transmit a frame, the network cycle is established:

---

[49]Ibid.

cycle = (nodes * (Ttran + contention));

An important variable in the CSMA process is the size of the message packet, an administrative structure designated a frame, ranging from 64 to 1518 bytes in size and containing 18 bytes of control data at the fore of the frame that contains a variety of information, including the hardware addresses of the sending/destination nodes, the protocol type and/or length of the frame, and a frame check sequence (FCS) that alerts the receiving node of any transmission errors.

Figure 9. Frame Composition

		Protocol Type or Frame Length		
Destination	Source		Data (46-1500)	FCS

Excluding the control bytes, each frame may carry as little as 46 bytes of data to as much as 1500 bytes. For example a message 1 Mbyte in length would require either a high of 21,739 or a low of 667 frames, depending upon whether 64 or 1518 byte frame sizes were used. The overhead is 18 bytes per frame and itself can be a significant value, such as the above example where the 64 bytes frames generate 391,302 bytes in overhead and 1518 bytes frames create 12,006 in overhead. Frame size becomes an important issue when there are more than a few nodes contending for network access, since larger frames occupy the medium for longer durations than shorter frames, thus increasing the delay that other nodes experience in attempting to access the network. Moreover there are some applications, such as video and audio isochronous data, that benefit from a smaller frame size, due to the decreased latency between frame transmits, even though smaller frames require more frames per message and transfer less data per frame with more overhead. CSMA incorporates the advantages of variable frame sizes to suit the most efficient frame to the actual application.[51]

Although the overall approach of CSMA-CD/CA is simplicity at low cost and efficiency at low utilization, the contention nature of the access protocol or the selected

[51]Nemzow, The Ethernet Management Guide, p. 126-130.

frame size can result in less than ideal network performance at higher loads or for applications where minimal latency is a necessity.[52] For these reasons, the CSMA protocol is best suited for networks with sporadic transmissions, as this type of traffic volume incurs the lowest probability of collisions and produces the greatest throughput.[53]

III-3-2. Token Passing (Deterministic)

The concept of token passing can be traced to the early work of Newhall, Pierce, and others in the late 1960s, wherein a control signal or token is generated to continuously loop throughout a ring or bus network. The purpose of the token is to identify who has control over the connecting medium and to transfer this control at regulated times to ensure equal access. As the token travels along its path, it is subsequently retrieved, regenerated, and returned to the connecting medium by each and every node on the network. This process of repeating the signal is accomplished for all signals transmitted on the network, making the nodes a vital part of the communication process. Furthermore one node is selected as an active monitor and charged with ensuring the integrity of the token as it circulates through the network cycles. Upon monitor failure, the other nodes select a new monitor using a voting mechanism. To transmit a message, a node removes a circulating token, modifies it into a data packet, and replaces on the network for transmission to its destination. Once the recipient node receives the data packet, an acknowledgment is sent in the modified token to the sending node, which converts the packet into the original token and returns the token to the network. As each node can only hold a token for a specified duration, the repetition of this cycle results in a deterministic, predictable access to the medium for all node on the network.

---

[53]Gilbert Held, Token-Ring Networks: Characteristics, Operation, Construction and Management, Wiley Communications Technology, (Chichester, England : John Wiley & Sons, 1994), p. 32.

Figure 9. Token Passing[54]

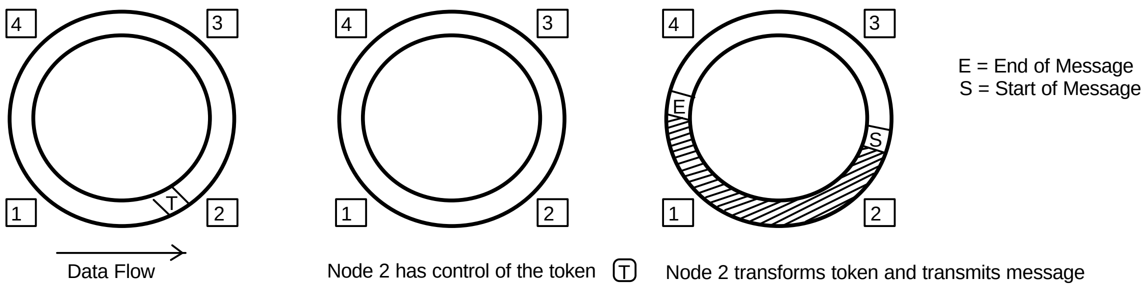

Because in token passing there is no contention or possibility of collision, the only delays introduced in a network cycle are those necessary expenditures incurred in transmitting the token, repeating the signal at each node, and copying of data at the receiving node. The algorithm for the token ring cycle is:

Example 2.[55]

Tpropnet = ((Lring / Vprop) + (nodes * interface_delay) + copy_delay + Btoken);

where,

Lring is the average signal distance (in meters),
Vprop is the signal propagation speed (in meters per second),
Interface delay is the nodal delay in repeating a signal,
Copy delay is the delay in receiving a message frame, and
Btoken is the time taken to transmit the token.

The combination of these variables, along with the time taken by a destination node to detect a token,

---

[55]See Appendix A, Tokenrun.cpp.

Ttoken = Btoken + (Lpass / Vprop) + interface_delay;

forms the basis of the network cycle:

cycle = (nodes * (Ttran + Ttoken));

where,

Ttran is the larger of the Tpropnet variable or the minimum message transmission time.

In token passing there are two primary types of transmission packets that may circulate on the network--the token and the data frame. The token consists of 3 bytes with the first byte indicating the token start, the second indicating control information such as priority access, reserved status, active monitor preference, etc., and the last byte indicating the end of the token. When seized by a node for message purposes, the token's 3 bytes are retained but converted into a frame format of a minimum 21 bytes. Included with this frame is the data,

Figure 10. Token and Frame Format.

Token

Start	Access	End

Frame

Start	Access	Frame Control	Destination	Source	Data (4500 - 18000)	FCS	End	Frame Status

which may range from 0 to a maximum of 4,500 bytes (4-Mbps token ring) or 18,000 bytes (16-Mbps token ring). A 1 Mbyte message transmitted on a token ring network would require either 222 or 56 frames, depending upon the speed, with generated overhead of only 4,662 or 1,176 bytes, respectively, a savings of at least 98% compared to the CSMA access protocol.[56]

Nonetheless the use of the token produces mixed results. On the negative side, the dependence upon the token can be a potential problem, since the loss or corruption of the token could bring down the entire network; a similar issue arises with the dependence upon the nodes as in-line repeaters. In fact, the use of a nodal repeater chain in token passing is a potential source of signal distortion, since each node's station generates a certain percentage of random noise that naturally introduces a probability of error in the refreshing/regeneration of the signal. As the network cycle progress, this distortion accumulates through the noise added by each node, a phenomena known as correlated jitter, and must be kept within a minimal level to ensure the integrity of transmissions. Furthermore, to avoid a network melt-down caused by the loss of a token or node, a complex and expensive system of checks and balances have to employed at the hardware and software levels, including automatic token monitoring and regeneration, node-to-node communication to ensure flow control, multiple connecting media paths for some systems, and self-healing mechanisms that permit a damaged node to remove itself from the network amidst continued operation. Often such complex internal management tools do not operate effectively, requiring time-consuming intervention on the part of the human operators.[57]

On the positive side, however, the advantages of token passing are predictable network access, with the elimination of contention disputes and variable delays caused by collisions as in CSMA-CD, and minimal degradation in performance as network utilization increases. Latency and jitter are consistently low, except at extreme loads, and the large frame sizes and low overhead make it possible for ring networks under certain configurations to develop throughput levels that exceed the actual bit data rate given in the available bandwidth, in otherwords values greater than 100% utilization.[58] The use of a token also permits larger frame sizes than allowed in CSMA and the

[57]Nemzow, The Ethernet Management Guide, p. 61; Held, Token-Ring Networks, p. 32-33.

incorporation of priority access, so that certain nodes or types of messages, such as isochronous data, are granted rights of access before other devices or messages. Lastly, the complexity of the hardware/software relationships enables automatic data-gathering and error reporting on the condition of the cabling and network nodes.

III-3-3. Polling (Deterministic)

Polling is a noncontention protocol that uses a central controller to deterministically allocate access of the connecting medium to all other nodes through a pre-arranged sequence. To achieve access, a node waits until the polling master generates its regularly-occurring query and then responds with an affirmative; otherwise, the transmit turn is given to the next node. The advantage to polling is predictable access and elimination of conflict between nodes, but these gains can be achieved with token passing with greater efficiency.[59]

[59] James S. Fritz, Charles F. Kaldenbach, and Louis M. Progar, Local Area Networks: Selection Guidelines, (Englewood Cliffs, NJ : Prentice-Hall, 1985), p. 12.

IV. ETHERNET

Originally developed under license from Xerox Corporation, the network architecture of Ethernet is based upon CSMA-CD and almost from its inception has been the inexpensive way to network. Ethernet is named for the hypothetical substance thought to fill all space, i.e., the *'ether'*, and was designed to provide low cost, high-speed 10-Mbps connectivity to a large number of potential users. Supporting up to 1024 nodes, Ethernet utilizes a shared network that allocates bandwidth equally among all the nodes via a bus trunk or star topology, with access to the network provided through contention. Due to the potential for collisions as utilization increases, the highest efficiency occurs when only a few nodes are accessing the medium at any one period of time, for under these conditions the proportion of bandwidth among those active nodes is generous and throughput is high. Since the primary task of networks in office environments is the handling of bursty data, e.g., intermittent text, numerical data, or bulk file transfers, the characteristics of CSMA have lent themselves well to such applications.[60] The combination of affordable, reliable performance networking has propelled Ethernet over its competitors in the marketplace.

Figure 11. Network Market Shares.[61]

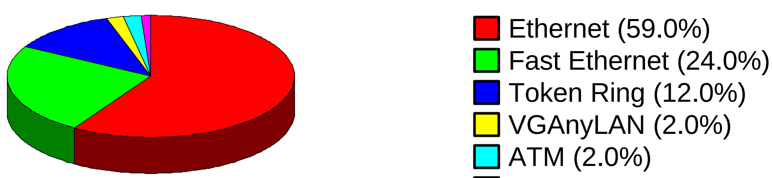

Ethernet is inexpensive and reliable because it uses a simple media access protocol and hardware components, the hardware consisting of a connecting medium, a transceiver, a repeater, and a network controller card for each workstation. The original

---

[61] 1997 projections adapted from data provided by International Data Corporation, in Robert Breyer and Sean Riley, PCWEEK: Switched and Fast Ethernet, 2nd. ed., (Emeryville, CA : Ziff-Davis Press, 1996), p. 51.

Ethernet, designated 10Base5,[62] uses a connecting medium of coaxial cable up to 500 meters in length per network segment, with a total of 5 interconnected segments permitted, or a total network span of 2500 meters. Along the length of the cable, every 2.5 meters, a marking indicated where a tap, an electronic interface known as the transceiver, could be installed. The transceiver is a combination transmitter and receiver that houses the electronics necessary to interact with or modulate the connecting medium, including the carrier sensing and collision detection mechanisms. In turn the transceiver was connected via a transceiver cable to the workstation's network controller card, which transmits and receives data to and from the network. Although the transceiver and to some extent the controller card perform complex tasks, the manufacturing of these devices is not cost-intensive. At the end of each network segment, with a maximum of 4 per total network span, a simple amplifying device, or repeater, is installed to refresh and retransmit the signal.[63]

Later versions of Ethernet retain the low cost and simplicity of the hardware/software but modify the operating parameters to attain different performance characteristics. Primarily the determining factor in the different versions is the type of connecting medium employed, with thin coaxial and twisted-pair cabling being the predominant options, although fiber optic solutions are available. Ethernet versions include:

- 10Base2, also known as "cheapernet", is a less expensive and more compact option that uses a cheaper coaxial cable with less distance capacity. In 10Base2, the maximum network segment is only 185 meters and permits a total of 30 potential nodes per segment, with the workstations attached not via a transceiver but directly to the cable segment. Expansion of the network span is allowed up to 925 meters or 5 network segments.

- 10Broad-36 uses a premium-quality coaxial cable capable of broadband, multiple channel transmission of distances spanning up to 3600 meters. Effective in performance, the higher costs associated with cable expense and the necessity of more intelligent circuitry in the transceivers have limited the application of 10Broad-

---

[63]Held, Ethernet Networks, p. 71-82.

36.[64]

- 10Base-T, on the other hand, achieves even lower costs through dependence upon a star topology and a connecting medium of unshielded twisted-pair wiring. Utilizing a hub and radiating nodes, 10Base-T offers point-to-point links and ease of installation since the UTP cable of lengths up to 100 meters is directly connected between workstations and the central multiport repeater in the hub. One benefit of point-to-point links is, under switched hub scenarios, the allocation of the full 10-Mbps bandwidth per switched port, thus modifying the shared network into a series of dedicated connections with greater throughput, lower latency values, and full-duplex capability. A 10Base-T network can be expanded to a maximum of 5, 100 meter segments connected by 4 hubs. While a 500 meter network span is significantly smaller than that offered by either 10Base5/2, the drastically reduced costs and availability of UTP have made 10Base-T a popular Ethernet version.[65]

- 100Base-T achieves a 10-fold increase in data rate and is known as Fast Ethernet, for this class offers 100-Mbps performance using category 5 UTP cabling. Built upon the star topology with switched hubs, 100Base-T offers both half and full-duplex options with total network spans ranging from a limited 250 meters with 2.5 100 meter network segments, to greater distances of 2000 meters with 100Base-FX, a fiber optic derivative of the 100Base class. Furthermore it is possible to upgrade an existing 10-Mbps Ethernet network into a 100Base model, a factor along with low cost and quality performance that makes Fast Ethernet a strong contender for the high-speed network market.[66]

- 1000Base class Ethernet is a gigabit-speed hybrid that utilizes the essential Ethernet media access protocol but with modifications to permit 1000-Mbps data rates. Although it is possible to operate Gigabit Ethernet as a 1000Base-T version over UTP, the total network span is restricted to 100 meters due to the poor performance of twisted-pair at these speeds. Better performance is obtained using one the fiber optic variants, 1000Base-X, which support network spans up to 5000 meters.[67]

IV-1. Advantages of Ethernet

Ethernet is popular because its CSMA-CD protocol on an easily installed bus topology offers equal access for all nodes accomplished at low managerial overhead and low hardware/software costs, with high performance in terms of throughput achieved at low to moderate utilization (< 50%), with the best performance obtained at less than 35% percent utilization. Frame sizes are variable with reasonable ratios of throughput to overhead and can be tailored to match the application data. Latency and jitter is minimal at light loads, and under switched environments continue to be reasonably low values

---

[65]Ibid., p. 95-103.

[67]Seifert, Gigabit Ethernet, p. 150-158.

under heavier data and/or isochronous loads.[68] Furthermore Ethernet is widely supported by market vendors for a variety of workstation interfaces and the manufacturing costs of Ethernet components continue to decrease as economies of scale are realized.[69] The expense of an Ethernet network controller card, for instance, can be as low as $15.00 per card,[70] and are readily available from a variety of vendors; in contrast, controller cards for token passing rings cost nearly 10 times as much, at $140.00 each.[71]

IV-2. Disadvantages of Ethernet

Although Ethernet offers 10-Mbps bandwidth capacity, due to the shared bandwidth and intrinsic delays caused by the collision-detection mechanism, the actual useable bandwidth is much less, being a proportion of the total bandwidth divided by the number of nodes minus the potential delay from collisions. Thus a 10-Mbps shared Ethernet network with 40 active nodes only offers each node a maximum of 250-Kbps in bandwidth. However a full-motion 30 fps slice of video, compressed using highly efficient algorithms such as MPEG, requires a network data rate ranging from 2 to 4 Mbps (2000 to 4000-Kbps), an amount that alone demands nearly 50% of the available 10-Mbps bandwidth on an Ethernet network.[72] Moreover the use of a bus topology and its shared medium with CSMA contention can result in a significant numbers of collisions, which, at higher utilization levels (> 60%), will drastically impede throughput and slow performance to interminable levels, with users waiting for seconds, even minutes, for processing. While a delay or latency this high is perhaps acceptable for text or numeric data files, any latency above a few milliseconds[73] becomes intolerable for real-time video or audio, where any surge in the data signal, or 'burstiness', will be

---

[69]Held, Ethernet Networks, p. 42.

[71]Current price quoted in CompUSA Direct, vol. B9.10, April 1999, p. 16.

[73]A commonly accepted range is from 20 to 100 ms.

extremely noticeable.[74] Lastly, due to the statistical nature of the media access protocol, the potential exists for a node to never obtain access to the network, or at least be considerably delayed in relation to its neighbors.

Because of these inherent limitations, Ethernet has had restricted applicability to environments involving high traffic or extensive transfers of isochronous data, e.g. video and audio, although attempts have been made to compensate for these weaknesses through limiting the amount of nodes, creating sub-networks via switched hubs, installing secondary dedicated media--such as IsoEthernet, which depends upon high-speed backbone links--or utilizing performance buffers at the nodes that capture and modulate incoming isochronous data to minimize latency effects. Therein lies the problem for Ethernet, as the strengths of CSMA-CD are also its source of weakness at higher utilization or involving isochronous data.[75]

With the introduction of Fast Ethernet, however, the traditional Ethernet model has been upgraded to 100-Mbps and even 1 Gbps, speeds that permit greater bandwidth per node and maintaining of low latency/jitter values despite an increase in traffic. Moreover the use of full-duplexing, switched environments can cause the 10 or 100 fold increase to be doubled once again in speed and free the network up from the restrictions inherent to CSMA. Thus a 100-Mbps Fast Ethernet with 40 nodes offers each node a maximum of 2.5-Mbps in bandwidth minus collision delay, but in full-duplex operation twice that bandwidth, or 5-Mbps, is made available for high-throughput central devices such as servers, storage media, and switches/routers. At these speeds and using dedicated switch-to-desktop connections with small frame sizes, the performance of Ethernet technology, especially Gigabit Ethernet, rivals the supposed champion of isochronous transmission, Asynchronous Transfer Mode (ATM), but without the disadvantages of cost, complexity and limited application that is incumbent upon ATM.

---

[75]Breyer, PCWEEK: Switched and Fast Ethernet, p. 40-43.

Consequently, with supercharged versions of Ethernet on the market, the inherent drawbacks to the CSMA-CD media access protocol in terms of isochronous transmission have to a large degree receded into the bilious ether.[76]

V. TOKEN RING

Adopted by IBM in the fall of 1985 as its baseband networking solution, token ring network architecture utilizes a logical ring topology and the concept of token passing, wherein a token is generated to circulate throughout a series of point-to-point links. Token ring as designed by IBM was, at face value, an impressive system, supporting (1) up to 260 nodes per network, (2) high-performance 4-Mbps or 16-Mbps speeds, (3) priority access mechanisms that provide a measure of quality-of-service guarantees, (4) enhanced fault-tolerance through automated management of physical connected medium and nodes, (5) a media access protocol that permits even, distributed access to all nodes with acknowledgment of packet receipt, and (6), could achieve and maintain over 90% utilization without significant impairment in throughput or incurring excessive delays. These features, along with the marketing muscle provided by a major manufacturer of the likes of IBM, caused industry analysts upon token ring's release to view the network as the successor to Ethernet technology in the LAN market.[77] These expectations, however, went unfulfilled for reasons that will be discussed in a moment.

Token ring hardware, while not consisting of more components than Ethernet, is required to accomplished tasks of vastly greater sophistication and intelligence. Connectivity is straightforward, as one could use coaxial cable or the IBM-provided, 7-tiered cabling system that standardized connecting media options, ranging from twisted-pair variants of differing quality to fiber optic cable. Selection of a particular cable was made based upon the desired network speed, number of nodes per ring, and distance to the connecting units, with total network spans as small as 100 meters to as large as 2 kilometers using fiber. Connecting nodes to the medium is more complicated, using a multistation access unit (MAU) to link together a group of nodes, which are, in turn, linked to a larger portion of the network.

Figure 12. 10 Port MAUs[78]

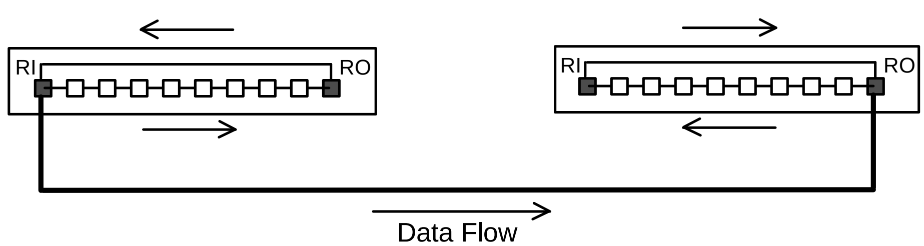

The MAU is a passive hub that permits 8 or more nodes per unit, with the Ring-In (RI) and Ring-Out (RO) connections used to combine a single MAUed group into a series of multiples, up to a maximum of 260 nodes and 33 MAUs. Note in Figure 12. that, in the

---

[77]Ibid., p. 6.

IBM-specified token ring's configuration, the network is more of a physical star structure than a ring, although the actual data flow is actually a logical ring topology. The purpose of the MAU is strictly that of establishing access to the connecting medium, with all the media access protocol handled by the complicated circuitry in the network controller cards and the software at the nodal workstations.[79]

The complexity of token ring is amply illustrated by the role of the Active Monitor, a workstation that is responsible for overseeing ring functions. All of the other nodes are considered to be Standby Stations and their primary role is that of repeating the signal and handling frames for transmission. As part of its duties, the Active Monitor's controller card must ensure that a free token is circulating and that there is never a prolonged occurrence of damaged, lost or multiple tokens. Moreover the Active Monitor is the timing center of the network, upon which all the other stations synchronize their clocks, and to maintain the proper flow control between all the nodes periodically (every 7 seconds) issues a Neighbor Notification Process that updates the tables for every Station's downstream and upstream neighbor relationship. Errors upon the network are handled either through a Purge process, which momentarily clears the network with a special purging frame, or, for more drastic faults, through a Beaconing process that forces a malfunctioning or suspected malfunctioning Station to leave the ring for a diagnostic check; in the latter process, if the defective Station is the Active Monitor, the Standby Stations enter a voting procedure based upon the highest hardware address and select a new Active Monitor. Remarkably, every one of the above functions is performed by the intelligent circuitry in the network controller card, as reliance upon software processing would impede the network data rate, hence the greater complexity and expense of token ring hardware.[80]

V-1. Advantages of Token Ring

---

[79]Held, Token-Ring Networks, p. 78-79.

As is evident, the intricate, finely-tuned media access protocol of token ring offers performance advantages that surpass other protocols in many ways.[81] These advantages include:

- Two data rate levels of 4-Mbps and 16-Mbps permit tailoring of the network to match bandwidth demands.
- The ability to assign priority to different types of messages or nodes enhances performance of high-priority data.
- Automated error reporting and generation of performance statistics facilitates network management.
- Larger supported frame sizes and minimal per frame overhead
- Distributed, predictable access to the network.
- High throughput even at >90% utilization levels.
- The benefit of centralized connecting links via a physical star topology.
- Consistently lower jitter values than contention-based networks

On the face of it, these features, along with IBM's support and a host of potential suppliers, should have ensured the success of token ring in its attempt to capture a portion of the network industry market, especially for isochronous data where token ring's performance is superior to Ethernet. For example, the same full-motion, MPEG compressed, 30 fps slice of video, that required 50% of Ethernet bandwidth (4-Mbps or 4000-Kbps), only requires 30% of a 16-Mbps token ring bandwidth; and this transmission can be prioritized over other traffic to ensure the best possible performance.

V-2. Disadvantages of Token Ring

Despite the excellent performance and added features offered by IBM's token-passing technology, upon its release token ring did not achieve the popularity and widespread adoption that had been anticipated. One problem at the outset was the insularity of the technology, being the sole proprietary design of IBM, from other manufacturers in the industry, which led to IBM's isolated position as the only promoter of token ring. In contrast, Ethernet had the support of a consortium of computer manufacturers, including Xerox, Hewlett-Packard and Sun Microsystems.

---

[81] Held, Token-Ring Networks, p. 32-33; Agnew, Distributed Multimedia, p. 150.

More importantly, however, were the complications arising out of the token ring technology itself. The use of complex media access algorithms forced implementation of hardware instead of software-based processing for network functions in order to maintain the necessary data rates. Reliance upon hardware demands that the network controller card and its workstation interface be of higher intelligence and capability than the simpler Ethernet design. Included in these capabilities is a dependency upon a central controller station, i.e. the Active Monitor, an essential system managerial property that requires Active Monitor circuitry to be included in every controller card, as any Standby Station could potentially be an Active Monitor in case of the designated Monitor's failure; this requirement along with hardware-based processing increases the complexity and cost for all the cards. And greater costs, regardless of performance advantages, tend to impede sales.[82]

Nonetheless higher costs are only part of the larger picture, since some of token ring's features are less beneficial than might be supposed. A central controller facilitates management and aids isolating errors, but the downside to centralization is the potential for loss of a segment or the entire network if the central component fails. While the ability to assign priorities is a practical idea, the effective use of token ring's priority mechanisms is a complicated scenario that causes disproportionate usage, and in reality most network applications of token ring fail to implement any other priority than "high-priority".[83] Another helpful concept, in theory, is the inclusion of automated error and performance data gathering; in practice, however, this feature went underutilized due to the inadequacy of upper-layer network applications to interface effectively with the information generated. Lastly, while a token ring is capable of high throughput and low latency, and thus is able to adequately process isochronous data originating from a few nodes, a function of the media access protocol is that as the number of nodes increases

---

[83]Ibid., 371-372; Cronin, "A Comparison of High Speed LANs," p. 48.

the overall bandwidth is deterministically divided among these additional nodes even though the resources cannot support this demand without introducing unwanted jitter.[84]

Solutions to enhance token ring for multimedia applications are thus somewhat limited. Priority access can be implemented, but a consequence is poorer performance for the lower priority data, imposing inefficiency for other types of traffic. The entire network could be devoted to isochronous use, but the expense of the components dedicated solely for multimedia becomes an issue. An upgrade of the token-passing concept can be adopted, such as the attempt to retain the performance edge while minimizing the drawbacks of token ring developed in Fiber Distributed Digital Interface (FDDI)

FDDI is a 100-Mbps local area network that utilizes a dual-ring fiber or copper connecting medium for increased fault tolerance with a span of up to 2 kilometers. Attaining in practice higher bandwidth utilization (typically 97.5%), than that of token ring through an enhanced media access protocol, FDDI performs well in terms of throughput, low delays, and stability. However FDDI is also extremely expensive and still lacks the ability to provide bounded delays necessary for isochronous data transfers, although there is a modification, synchronous FDDI, that segregates the traffic into asynchronous versus isochronous and then allocates a portion of the bandwidth to each division. As the same problems of prioritization discrimination arises in synchronous FDDI, only at higher speeds, the use of FDDI for multimedia purposes is debatable.

VI. COMPARISON

Having examined the various components and structures essential to local area networking, as well as the operating characteristics and strengths/weaknesses of the two dominant commercially-produced network technologies, that of Ethernet and token ring, it is now possible to consider these two technologies in relation to each other to determine which offers the best isochronous performance. This comparison will begin with the first of the performance parameters, throughput, then latency and jitter, and finally ending with error rates.

VI.-1. Throughput

Due to the distributed access of token passing, deterministic protocols such as token ring provide superior data throughput at almost all utilization levels as compared to stock, shared-bandwidth Ethernet networks. The only exception is at light utilization, where contention-based networks have an advantage in immediacy of bandwidth that excels over the slight delays inherent to token passing system, since the token circulates at a constant speed and a sending node must wait until the token is acknowledged and returned before being able to transmit again. A further advantage of token ring is that it is available in a 16-Mbit version, a significant bandwidth increase over Ethernet that facilitates reservation of network resources for multimedia. The combination of superior and expanded throughput capability enables token ring performance to provide the best match between the mean throughput rate and playback rate of isochronous media.

VI.-2. Latency

In regards to message delays, token ring's deterministic protocol provides transmissions with tighter latency bounds and lower jitter, although configuring Ethernet to use shorter packet lengths can compensate for this deficiency, since delay and the variation in delay is reduced with more frequent transmissions of small frames. However, since Ethernet has a higher throughput level when using long packets of over 1000 bits, resorting to short packets actually reduces effective data transmission. It is true that the overall lowest latency is produced by Ethernet, since the number of collision-less transmissions at low loads compensates for contention delays at higher loads, but latency values are more important under serious loading occurring at greater utilization levels, a performance area where token ring dominates given that jitter relative to Ethernet actually decreases at these levels. Furthermore, under token passing protocols, it is possible to implement priority schemes, mechanisms which can reduce latency over 75%, from 20ms to 5ms at over 90% utilization, and eliminate the jitter interjected into multimedia data streams from bursty data transmissions by isolating the isochronous streams from other types of traffic.

VI.-3. Error Rates

Error loss rates occur either through transient events, such as exceeding the maximum number of permitted collisions in Ethernet (16 per packet), and corruption of the data packet in token ring; or due to queuing dumps, i.e., the dropping of packets that occurs when the queuing buffers at the nodes overflow due to saturation. For both network technologies the error loss for transient events is minimal, definitely within the

0.01 - 0.001 target rate, with each media access protocol successfully handling and delivering an extremely high proportion of its packet frames. The difference, albeit slight, is in favor of token ring, because the contention and subsequent back-off algorithm evidenced in Ethernet leads to a greater number of damaged and/or lost frames. Furthermore the way in which loss occurs, the smooth, regular loss in token ring versus the sporadic nature of Ethernet's CSMA-CD, is again biased against Ethernet since consistent losses do not impair the quality of isochronous transmissions as much as randomly occurring ones. The queuing dump errors are more problematic, as the loss rate in this area depends upon the intensity of the network traffic and the size of the nodal buffers; nonetheless it is clear that there is an inverse relationship between the capacity of a buffer and packet loss. Thus, to minimize error loss at the receiving nodes, larger buffers should be employed.

VII. PRELIMINARY CONCLUSIONS

The research conducted and postulated in network literature, as well as the strictly theoretical foundation for topologies and media access protocols, indicate that both token ring and Ethernet are able to meet the specified performance parameters required in transmitting high-quality multimedia applications, parameters of a mean throughput rate being equal to or less than the playback rate; latency and jitter values of 20 to 400 milliseconds and < 80 milliseconds, respectively; and error loss rates between 0.01--0.001. The difference between the two network architectures is the degree of utilization level and number of supported nodes, wherein such performance bounds can be met, and the response of the isochronous streams to sudden loads of data traffic.

Ethernet as a platform for multimedia is best suited for network settings that will not experience on average greater than 50% utilization, where the imposition of data traffic upon the available bandwidth is controlled either through dedication of the network for isochronous data or by limiting the number of users that may simultaneously access.

The study by Dalgic, et al.,[85] conclusively demonstrates that the 10-Mbps, hub-based 10BaseT Ethernet variant can handle 5, 1536-Kbps (1.5-Mbps) isochronous streams at a peak 77% utilization, bounded by a 20 millisecond latency, and as many as 20, 384-Kbps streams at an equivalent utilization but relaxed to a delay of 100 milliseconds. These performance figures assume a non-bursty data load, with latency and jitter values quickly deteriorating when bursty data traffic is included, a normal occurrence for mixed media networks. Adjusting to allow smaller packet sizes alleviates this problem only slightly, since the gains are counterbalanced by increased contention arising from more opportunities to transmit.

In contrast, token ring networks are best suited for those applications experiencing constant utilization levels exceeding 50%, where there are numerous nodes connected to the medium, or where there is heavy data traffic causing bursts on the network medium. The study by Amer, et al,.[86] extols the advantages of using 16-Mbps token ring for isochronous data due to the low latency and jitter, as well as the option to employ priority mechanisms to reserve bandwidth for multimedia purposes. Moreover the deterministic property of token passing ensures that isochronous streams in transit are insulated by instances or bursts of asynchronous data, ensuring a large degree of quality of service in terms of delays expected at the receiving nodes.

---

[85]Dalgic, "Evaluation of 10Base-T Ethernets Carrying Video, Audio and Data Traffic," p. 1099.

VIII. PROJECT DATA

This project component involved a series of tests based upon varying frame sizes and number of active nodes to arrive at a variety of performance values for the optimized parameters. Central to the focus of this paper are concerns of throughput and delay, and the intent of the testing was to discover if the evidence in the data would confirm, contradict or at least highlight the performance differences between the two media access protocols. The algorithm used to develop the test is a FOR loop as follows:

Ethernet:

```
network [speed in Mbps] = 10000000;
for(nodes = 10; nodes <= 260; nodes += 20)
for(user_packet = 64; user_packet <= 1518; user_packet += 50)
```

Token Ring:

```
network [speed in Mbps] = 16000000;
for(nodes = 10; nodes <= 260; nodes += 20)
for(user_packet = 64; user_packet <= 18000; user_packet += 150)
```

Each subsequent pass through the loop increments the nodes value by 20 and the user packet value (frame size) by 50 for Ethernet; and 20 and 150 for Token Ring respectively, with 260 being the total number of nodes generated.[87]

VIII-1. Throughput

Ethernet's throughput across a range of frame sizes is consistent with the performance anticipated by the literature and experience. As is evident in Figure 13, the largest data volume is achieved with less nodes, as more nodes contending for access imposes a greater burden in terms of collisions and consequently impairs throughput.

Figure 13. Ethernet Throughput

In all cases, a scaling phenomena is exhibited whereby the greatest change occurs in those frame sizes from 64 to 714 bytes, and then gradually levels off throughout the remainder of the frame range, to a peak of nearly 1,200,000 bytes at the maximum 1518 frame, or a utilization of 97%. The ability of Ethernet to produce throughput at such high utilization is in contrast to the myth that Ethernet's CSMA-CD protocol will suffer a 'melt-down' or collapse beyond a certain level; in fact, the data rate is constant throughout the utilization range, although the contention process does limit the upper bound performance.

Throughput values for token ring are more impressive than Ethernet's, a result that is once again consistent with expectations, but the increased throughput of over 2,000,000 bytes is due primarily to the higher bandwidth capacity of 16-Mbps. When compared at 10-Mbps equally, token ring surpasses Ethernet by only a few 100,000 bytes, at 1,222,000 bytes. However, as Figure 14 illustrates, token ring throughput values begin at a higher level than Ethernet, at 750,000 bytes as compared to 400,000 bytes, and the higher throughput performance for token ring, albeit small at 10-Mbps and significant at 16-Mbps, is attained at all nodal levels, whether 10 or 110 (or greater),

[87]Upon analyzing the resulting data, it became evident that there was little change in performance beyond 110 and 260 nodes; for this reason, 110 nodes is used as the top nodal value in this comparison.

whereas Ethernet's throughput drops to under 900,000 bytes at 110 nodes.

Figure 14. Token Ring Throughput

In fact the only difference in token ring's throughput based upon nodal activity is evidenced at the lower frame sizes, up through 214 bytes. Because of this even distribution of bandwidth regardless of the number of nodes utilizing the medium, token ring is superior for larger networks.

A direct comparison of Ethernet and token ring throughput performance is provided in Figure 15, which has the first half of the graph illustrating 10 and 110 nodes Ethernet and the second part documenting token ring performance at the same levels.

Figure 15. Ethernet versus Token Ring Throughput

The noticeable disparity in starting and ending throughput levels is evidenced, with token ring providing greater throughput capacity at both lower and upper bounds essentially independent of attached nodes.

Figure 16, compares the minimum and maximum frame sizes for both Ethernet and token ring.[88] All frame sizes, with the exception of the maximum 18000 byte frame in token ring, are subject to declining throughput levels as the number of nodes increases, with the 1518 byte Ethernet frame providing the best throughput capacity among the three lesser sizes.

Figure 16. Comparison of Frame Sizes

It is noteworthy that the performance of the 64 byte token ring frame is nearly equal to that of the maximum Ethernet frame, again indicating the increased performance obtainable through token ring. Consequently, given that quality isochronous performance is dependent upon a throughput sufficient to enable a play back rate equal to the recording rate of the media, 16-Mbps token ring is best equipped to provide this throughput capacity, since its media access protocol permits consistently high data rates at nearly all nodal and frame size levels.

VIII-2. Latency and Jitter

Latency values for Ethernet follow the pattern expected for a CSMA-CD protocol, with smaller frame sizes generating lower latency than larger frames, and total average latency dropping as the number of nodes increases. Figure 17, indicates that latency is negligible, under 5 ms, for 514 byte and under frame sizes, and then quickly climbs as the frame becomes larger.

Figure 17. Ethernet Latency

Based upon the criteria of 20 to 400 microseconds for effective isochronous transmission, Ethernet's performance is excellent given that the maximum experienced latency is slightly over 40 ms, with minimal latency exhibited in the 1000 byte and under frame size range. Given that anywhere from 60% to over 90% of the data rate throughput, or upwards of 1,100,000 bytes, is attained at frame sizes near 714 bytes (see Figure 13), an Ethernet network utilizing multimedia applications should assign frame sizes in the 700-1000 byte range to maximize throughput while maintaining low latency values.

The test results also indicate that jitter is reasonable, with the maximum variance exhibited by 10 nodes of 0-40 ms for the frame size range of 64 to 512 bytes, and a minimum variance of 0-13 ms for 110 nodes. For the purposes of quality transmission, jitter must be held within the bounds of 80 ms to avoid distortion of the message signal, and Ethernet achieves this objective with a 200% variance cushion. Moreover, as

Figure 18, illustrates, the variance per individual frame is extremely low, an important factor since overall average values may not reflect the performance experienced by one or more items in a data series.

Figure 18. Variance Per Frame Latency

With jitter values per frame ranging from 0.08 through 1.2 ms, indicating that the variance in delay over the range of frame sizes will never deviate more or less than these values, Ethernet's timeliness and consistency in delivery under optimized conditions is clearly well-suited for isochronous transmissions. Despite these positive results, however, the final determination of Ethernet performance must consider how these latency and jitter values would be impacted by allowance for typical operating conditions, including bursty data traffic that introduces interfering sporadicity to the message transmission process.

Ethernet's latency and jitter values are excellent, being within the upper bounds by a generous margin; nonetheless, token ring's performance is superb. Figure 19, exhibits how average latency for a token passing ring never exceeds a maximum of 0.75 ms, a 98% reduction over Ethernet maximum latency value of 40 ms.

Figure 19. Token Ring Latency

Furthermore, and again unlike Ethernet, token ring's developed latency remains essentially stable over the range of packet sizes, enabling utilization of the data throughput delivery rates possible with larger frame sizes. There is also less disparity between average latency and number of active nodes, with the total variance between

10 and 110 nodes being at most only 0.70 ms, a jitter value that is extremely promising for isochronous applications.

It must be noted, however, that one disadvantage to token ring latency is that there is a higher initial latency, a minimum delay incurred through the token passing mechanism, that imposes a constant burden upon the system. While Ethernet's is more motile with a higher upper bound, the latency imposed upon initial packets begins at 0 and only then accelerates upward. The existence of an innate latency for token ring would be an inhibiting factor for immediate response applications, such as voice or telephony, if the latency values were of greater magnitude. Nevertheless, the extremely low, consistent latency values produced by token ring are strong indicators--asserted by the literature and by these results--that message delay will be essentially insular in nature even when the network is subjected to less than optimal conditions involving bursts of data traffic.

Jitter values of less than 1 ms on average are nearly ideal, and the variance per frame delay indicates that this performance will occur at minimal deviation, with Figure 20, illustrating how the average delay for each transmitted frame varies a maximum of only 0.14 ms, as compared to 1.2 ms for Ethernet, from the latency values for each instance of the nodal range. Moreover, the greatest deviation occurs in the 64-3000 byte frame size, tapering off to near 0 with larger frames, a capability that facilitates higher data rates without comprising quality of transmission through increased jitter due to individual variations in frame delay.

Figure 20. Variance Per Frame Latency

The varying capabilities of token ring and Ethernet for generating low latency values is dramatically illustrated in Figure 21, wherein token ring never exceeds 1 ms of delay over the entire range of frame sizes, while its counterpart experiences a quick climb to as high as 9 ms (10 Nodes) of delay for a frame size only as large as 714 bytes, or 47% of the maximum 1518 byte frame.[89] Extrapolating from the evidence, it is clear that Ethernet's CSMA-CD media access protocol is extremely sensitive to changes in frame size and/or the number of nodes

Figure 21. Token Ring versus Ethernet Latency

that are transmitting, and thus will be more impacted by irregularities in the traffic flow, especially in typical, less-than optimal conditions. In contrast, token ring is a predominately a stable-state system, with little to no variation in delay over the entire range of frame sizes--a true accomplishment given the range is from 64 to 17,914 bytes--and only minor increases due to additional nodes; consequently, a determination may be made that token ring's performance in typical, non-optimized applications will be at a high percentage, if not equal to, these impressive test results.

VIII-3. Error Rates

Lost frames arises from two conditions; transient events, such as lost tokens/frames due to hardware faults (token ring) or dropping of frames because of excessive collisions (Ethernet, > 16 collisions allowed per frame); and from saturation of the nodal buffers when the packet arrival exceeds the queuing capacity of those buffers. In practice, the occurrence of transient events in properly configured, installed and maintained networks of both protocols is minimal, with even such a likely source as Ethernet's collisions requiring retransmitting of dropped frames at rates that linger in the millionth of a second or smaller range. In contrast, however, lost frames due to buffer overload is more prevalent and can be a significant factor in impairing the quality of bandwidth intensive, time-sensitive isochronous applications. Acceptable error loss rates for multimedia are considered to be within 0.01-0.001 range, and the test results in

[89] A compensating factor, however, is that data throughput is already at the 60% level for this frame size.

Figure 22, indicate that both token ring and Ethernet are effective in handling their frame load, with the former's buffers never exceeding 4 hundredths of a bit and the latter's never exceeding 5 bits. Given that an Ethernet network generates 9,883,392 bits at the 1514 frame size for 10 nodes, an error rate greater than 0.001 (higher performance boundary), would require buffer values larger than 9883 bits or 1235 bytes, about the length of one maximum-sized frame. Therefore at a buffer capacity of 5 bits Ethernet's performance is substantially within the desired error rate, although there is an indication based upon the climbing values that Ethernet would be more prone to queuing loss than token ring.

Figure 22. Nodal Buffer Size

IX. FINAL CONCLUSIONS BASED ON VII. AND VIII.

The project data confirmed and illuminated network performance for the three measures of throughput, latency and jitter, and error rates that had been expected based upon the assertions in the literature. Although there were no surprises, the test results did highlight the essential operating difference between token ring and Ethernet, which may be summed as follows:

- token passing provides performance values that are, for the most part, isolated from frame sizes, level of nodal range, and external data traffic, while CSMA-CD performance is extremely dependent upon such variables.

Other than this important distinction, the performance exhibited by both network media access protocols is very similar, in that both token ring and Ethernet meet the demands of high-quality isochronous applications, i.e., a mean throughput rate being equal to or less than the playback rate; latency and jitter values of 20 to 400 milliseconds and < 80 milliseconds, respectively; and error loss rates between 0.01--0.001. The question, then, is not so much one of capability but of the anticipated requirements for the network application, such as number of users, the type of isochronous media and required message size, and presence and intensity of other data traffic.

If the network is to be devoted to isochronous use with a dedicated pool of network users, and transmission will occur amidst only a minimum of bursty data traffic, then Ethernet's acceptable performance, conjoined with its dependability and low cost, are assets testifying in its favor. Moreover the use of a switched environment instead of shared bandwidth, such as used in 10Base-T, as well as the switch-based and increased bandwidth upgrade of Fast Ethernet, are alternatives to the stock Ethernet model that have potential to further isolate the data traffic of one node from its neighbors and to reduce the latency times involved. Nonetheless it is evident from the test results that the Ethernet's statistical CSMA-CD contention protocol is not innately suited for the stringent transmission requirements of isochronous data, with potential scalability issues--especially in the area of latency and jitter--arising in the future as network use or multimedia demands increase, and Ethernet's advantages as the lower cost network should be carefully weighed in light of this fact.

Token ring, on the other hand, is truly a higher performance network paradigm than Ethernet, an opinion argued for in the literature and proven by the test results, especially where concerns of regulated access and minimal, consistent delays are paramount, as in multimedia applications. Token ring's logical ring/physical star topology with deterministic access, insular transmission independent of other data traffic, centralized control, and expanded 16-Mbps bandwidth option are all key components to the obtained performance gains in throughput and decreased latency and jitter. Moreover there exists an option for priority access to ensure that a measure of quality of service is explicitly granted to isochronous data. The additional costs of these features and added performance as compared to Ethernet are not inconsequential, as well as the risk of malfunction introduced by hardware complexity, but for the right application the considerable performance advantage of token passing may predominate over such costs/risks. For token ring the question becomes one of balancing cost over performance; if one truly wants an effective, scaleable, and highly efficient network platform for isochronous purposes, than the cost of implementing a token ring network actually becomes irrelevant. Indeed it is better to expend more but achieve one's current and future goals then spend less now but in a few years incur costs through needed retrofitting or replacement.

Concerns of scalability and cost/performance ratios are of particular interest to digital libraries, institutions whose livelihood is dependent upon the network infrastructure into which resources have been invested. As stated, designing the network for a digital library is a combination of art and science, since the possible technological options in networking and uncertainty inherent to new technologies often precludes a precise determination as to which is the best form of technology to adopt. A purchaser may decide upon one network architecture over another but, depending upon the time of adoption, such as initial release, the market may not fully support the servicing of one's chosen options. Furthermore the complexity of network design, and the particularity of each specific library application, makes it difficult to estimate the selection of the best network architecture for a digital library. Lastly, multimedia is a substantial component of a digital collection and such media are remarkable for being subject to an ever-increasing demand for higher quality which pushes the high-end threshold of acceptable network performance continuously upward. Consequently, the artful selection of an appropriately capable network architecture is essential to a digital library's success, for the analytical tools are only a guide to potential performance rather than a concrete determination of that performance.

In conclusion, although it is not possible to recommend an 'optimal' network design, the results from the project data do strongly lean towards token ring as the best time-division multiplexing (TDM) network architecture for settings involving heavy multimedia applications, including digital libraries. For token ring offers sufficient bandwidth and delay criteria to efficiently process resource-hungry media, such as isochronous video and audio, and, unlike Ethernet, the intrinsic nature of these performance values permits future expansion as demand and/or multimedia requirements increases. While token ring has lost a significant percentage of its market share since initial release, its token passing media access protocol is still widely supported by manufacturers and will likely continue to be a viable network architecture for some years in the future. This recommendation of token ring recognizes the disproportionate cost difference between token ring and Ethernet but, given the premise of a network with a dedicated or large multimedia requirement, the superiority of token ring for transmitting isochronous data outweighs the additional expenditures. In identifying this TDM media access protocol over its counterpart, network designers of digital libraries and other settings will be aided in selecting either token ring or Ethernet as the network architecture solution for their multimedia purposes, and the repetition of this selection process will gradually incorporate more of the art inherent to digital library and network design into networking standards.

BIBLIOGRAPHY

Lecture

Gogan, James. INLS 184-01, Network Protocols and Management. School of Information and Library Science. Chapel Hill, NC : University of North Carolina at Chapel Hill, Spring 1999.

Monographs

Agnew, Palmer W. and Anne S. Kellerman. Distributed Multimedia: Technologies, Applications, and Opportunities in the Digital Information Industry, A Guide for Users and Providers. Reading, MA : Addison-Wesley Publishing, 1996.

Böllmann, Elizabeth and Engelbert Gößler, ed. 17th Library System Seminar: The Virtual Library: Karl-Franzens-Universität Graz, Meerscheinschlößl, 14-16 April 1993. European Library Automation Group. Universitätsbibliothek Graz, 1993.

Breyer, Robert and Sean Riley. Switched and Fast Ethernet. 2nd. ed. Emeryville, CA : Ziff-Davis Press, 1996.

Comer, Douglas E. The Internet Book: Everything You Need to Know About Computer Networking and How the Internet Works. 2nd ed. Upper Saddle River, NJ : Prentice-Hall, 1997.

Ferrero, Alexis. The Evolving Ethernet. Translated by Stephen S. Wilson. Data Communications and Networks Series. Harlow, England : Addison-Wesley, 1996.

Hancock, Bill. Designing and Implementing Ethernet Networks. 2nd. ed. Wellesley, MA : QED Information Series, 1988.

Hegering, Heinz-Gerd and Alfred Lapple. Ethernet: Building a Communications Infrastructure. Translated by Stephen S. Wilson. Data Communications and Networks Series. Wokingham, England : Addison-Wesley, 1993.

Held, Gilbert. Ethernet Networks: Design, Implementation, Operation, Management. 2nd. ed. New York : John Wiley & Sons, 1996.

Held, Gilbert. Token-Ring Networks: Characteristics, Operation, Construction and Management. Wiley Communications Technology. Chichester, England : John Wiley & Sons, 1994.

Fritz, James S., Charles F. Kaldenbach, and Louis M. Progar. Local Area Networks: Selection Guidelines. Englewood Cliffs, NJ : Prentice-Hall, 1985.

Kessler, Jack. The International Dimension. Artech House Computer Science Series. Boston, MA : Artech House, 1996.

Kim, Jeong-Jae. "Local Area Networks: A Methodology for Decision Making in Choosing Cabling, Topology, and Access Method." Master's Thesis. Chapel Hill, NC : University of North Carolina at Chapel Hill, July 1992.

Nemzow, Martin. The Ethernet Management Guide. 3rd. ed. McGraw-Hill Series on Computer Communications. New York : McGraw-Hill, 1995.

Nilausen, Jesper. Token Ring Networks. Translated by Walther Falk. New York, NY : Prentice Hall, 1992.

Seifert, Rich. Gigabit Ethernet: Technology and Applications for High-Speed LANs. Reading, MA : Addison-Wesley, 1998.

Proceedings

Amer, Khaled, Ken Christensen, and Tora Toher. "Experiments with Client/Server Multimedia on Token Ring." IEEE Proceedings: 18th Conference on Local Computer Networks. 1993.

Bisdikian, Chatschik C., Baiju Patel, Frank Schaffa, and Marc Willebeek-LeMair. "On the Effectiveness of Priorities in Token Ring for Multimedia Traffic." IEEE Proceedings: 18th Conference on Local Computer Networks. 1993.

Cronin, William J., Jerry D. Hutchison, K.K. Ramakrishnan, and Henry Yang. "A Comparison of High Speed LANs." IEEE Proceedings: 19th Conference on Local Computer Networks. 1994.

Dalgic, Ismail, William Chien, and Fouad A. Tobagi. "Evaluation of 10Base-T and 100Base-T Ethernets Carrying Video, Audio and Data Traffic." IEEE Proceedings: 19th Conference on Local Computer Networks. 1994.

Elsaadany, Amr, Mukesh Singhal, and Ming T. Liu. "Priority Communication Schemes on Local Area Networks for Multimedia Traffic." IEEE Proceedings: 19th Conference on Local Computer Networks. 1994.

________. "Alternative Configurations for Local Area Network Design." IEEE Proceedings: 20th Conference on Local Computer Networks. 1995.

Gonsalves, Timothy A. and Fouad A. Tobagi. "Comparative Performance of Voice/Data Local Area Networks." IEEE Proceedings: 14th Conference on Local Computer Networks. 1989.

Jeffay, K., D. L. Stone, T. Talley, and F.D. Smith. "Adaptive, Best-Effort Delivery of Digital Audio and Video Across Packet-Switched Networks." In Network and Operating System Support for Digital Audio and Video: Third International Workshop, La Jolla, California, USA, November 1992 Proceedings. Ed. P. Venkat Rangan, Lecture Notes in Computer Science Series, no. 712. Berlin : Springer-Verlag, 1993.

Lesk, Michael. "Which Way to the Future? The Control of Scholarly Publication." In Digital Libraries: Current Issues: Digital Libraries Workshop DL '94. Eds. Bharat K. Bhargava and Yelena Yesha, Lecture Notes in Computer Science Series, no.

916. Newark, NJ : Springer, 1994.

Liu, Ming T., Wael Hilal and Bernard H. Groomes. "Performance Evaluation of Channel Access Protocols for Local Computer Networks." Proceedings COMPCON, Fall '82. September 20-23.

Liu, M. T. and D. M. Rouse. "A Study of Ring Networks." In Ring Technology Local Area Networks: Proceedings of the IFIP WG6.4/University of Kent Workshop on Ring Technology based Local Area Networks, Kent, U.K., 28-30 September, 1983. Eds. I. N. Dallas and E. B. Spratt. Amsterdam, Holland : North-Holland, 1984.

Mok, Dennis S. "A Software Tool for Performance Analysis and Evaluation of Local Area Networks." IEEE Proceedings: 11th Conference on Local Computer Networks, 1986.

Sato, Ken-ichi, Hiroshi Nakada and Youichi Sato. "Variable Rate Speech Coding and Network Delay Analysis for Universal Transport Network." IEEE Proceedings: 13th Conference on Local Computer Networks. 1988.

Suda, Tatsuya and Tracy T. Bradley. "Packetized-Voice/Data Integrated Transmission on a Token Passing Ring Local Area Network." IEEE Proceedings: 12th Conference on Local Computer Networks. 1987.

Serials

Adam, Nabil R., Milton Halem, and Shamim Naqvi. "Promising Research Directions In Digital Libraries." In Digital Libraries: Current Issues: Digital Libraries Workshop DL '94. Eds., Bharat K. Bhargava and Yelena Yesha, Lecture Notes in Computer Science Series, no. 916. Newark, NJ : Springer, 1994.

Bux, Werner. "Local-Area Subnetworks: A Performance Comparision." IEEE Transactions on Communications COM-29. 10 October 1981.

Krunz, Marwan. "Bandwidth Allocation Strategies for Transporting Variable-Bit-Rate Video Traffic." IEEE Communications Magazine. January 1999.

Levy, David M., Catherine C. Marshall. "Going Digital: A Look at Assumptions Underlying Digital Libraries." Communications of the ACM. Vol. 38, no. 4. April 1995.

Matthews, Joseph R. and Mark R. Parker. "Local Area Networks and Wide Area Networks For Libraries." Library Technology Reports. Vol. 31, no. 1. January-February, 1995.

Trade Magazines

CompUSA Direct. Volume B9.10. April 1999.

Dartek.Com. May/June 1999.

APPENDIX A

C++ PROJECT COMPONENTS

I. Source Code

The source code for the research network comparison application is divided into 23 files primarily comprised of either functions or modules (collections of functions) to accomplish the tasks of management, selection and modification, operation, and review of test results. These files are listed, in alphabetical order, as follows:

Modules	Functions	Prototype File
Comp_sel.cpp	Average.cpp	Prototyp.h
Compethr.cpp	Averorig.cpp	
Comptoke.cpp	Evaluate.cpp	
Digital.cpp		
Eth_read.cpp		
Eth_sele.cpp		
Ethernet.cpp		
Etherun.cpp		
Filecrea.cpp		
Read.cpp		
Select.cpp		
Test_net.cpp		
Testbed.cpp		
Testrevw.cpp		
Tok_aver.cpp		
Tok_read.cpp		
Tok_sele.cpp		
Token.cpp		
Tokenrun.cpp		

The role of the modules and functions, and their respective importance to the overall analytical model, may be easily displayed using a slightly modified version of the flowchart from the Introduction.

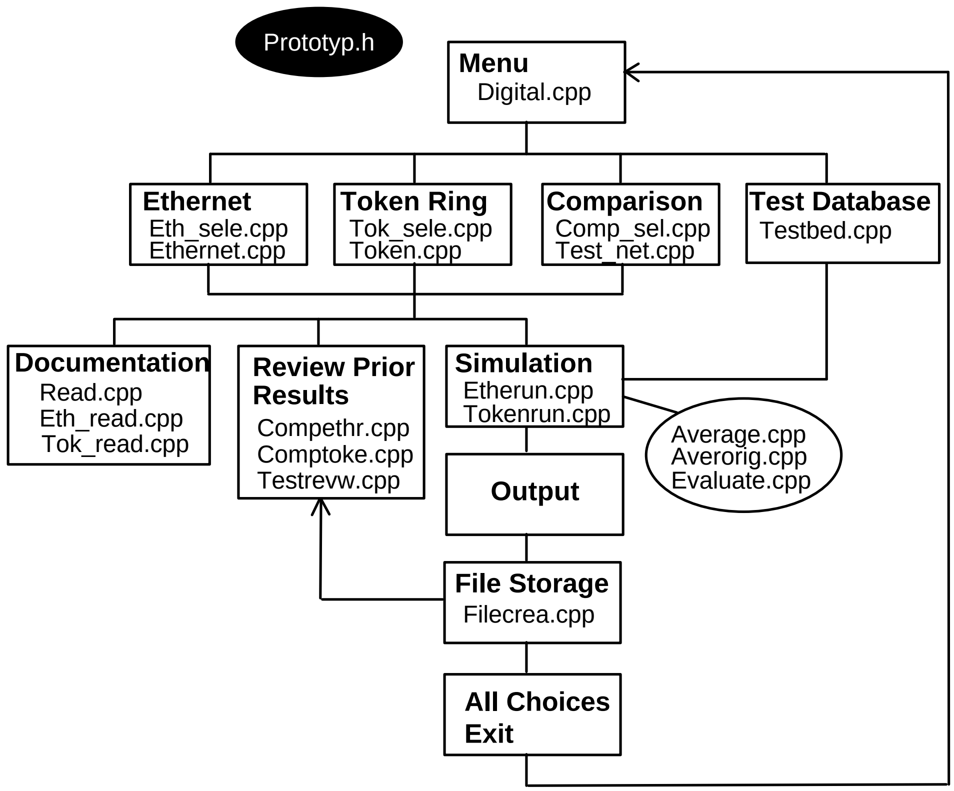

More detailed information about each module and function may be obtained from the print-outs included for the source code within this Appendix or by examining the electronic files stored on the attached diskette in directory "/source."

II. Object Code

The object code, resident in "digital.exe" in directory "/program" on the attached diskette, is a fully functional version of the network performance evaluation software that was designed and encoded in C++ for this project. Options are provided for individual, mutual, or bulk research comparison of the Ethernet and token ring network paradigms.

III. Research Data

To ensure future verifiability of the data's integrity and enable cross-referencing of the conclusions propounded by this project, the research data generated by the Test Database portion of the evaluation software is located, in raw and analyzed form (Works 3.0 format), in directory "/research" on the attached diskette.