# A Comparison of High Speed LANs

**William J. Cronin, Jerry D. Hutchison, K.K. Ramakrishnan, Henry Yang**
Digital Equipment Corporation
550 King Street, Littleton, Ma. 01460
hutchison@npss.lkg.enet.dec.com

**Abstract**

*Ethernet at 10Mb/s and other legacy LANs such as 4 and 16 Mb/s Token Ring have matured and proliferated widely over the last decade. Now, applications have started to exceed the capabilities of these low speed LANs. A new generation of high speed networks are being deployed to meet with increasing expectations of users. This paper compares four such high speed LANs which operate at or above 100 Mb/s: 100BASE-T, 100VG-AnyLAN, FDDI and ATM. The attributes and performance of these four LANs are described and contrasted. The implications of migration to these new technologies are discussed.*

## Introduction

The past few years have seen the development of increasingly powerful desktop computers, servers, and the applications which can use this available computing power. As these applications become networked, users are demanding cost effective higher speed Local Area Network technologies that can keep pace with their distributed computing needs. Fortunately, or, unfortunately, depending on your point of view, the last few years have seen the development of many new technologies capable of providing high bandwidth services directly to users' desktops over widely installed Unshielded Twisted Pair (UTP) cabling. Four of these technologies, which are serious contenders for widespread deployment, are discussed in this paper. They are 100BASE-T (Fast Ethernet), 100VG-AnyLAN (Demand Priority), Fiber Distributed Data Interface (FDDI) and Asynchronous Transfer Mode (ATM).

In many cases users are adequately served by existing LAN technologies such as Ethernet and Token Ring. The development of switching hubs has allowed many of these bandwidth hungry networks to be upgraded, simply by connecting existing desktop network equipment to a switch port. This provides a bandwidth increase when the low speed adapter is adequate for the communication needs of the station. As the workload increases (e.g. larger file transfers become more common ) the transmission speed itself becomes the bottleneck, and the performance of switched low speed technologies becomes inadequate.

In this paper we will present a side by side comparison of important details of these high speed protocols. We will discusss the differences between shared and switched LANs, the operation and applicability for use in both the workgroup and the backbone environments, and the results of performance analysis of each protocol.

## Shared and Switched LANs

In a LAN, stations can share a single communication channel using a Media Access Protocol (MAC) or they may attach to a "Switch". The first mode of operation is a "shared bandwidth" LAN while a switch forms a sort of "extended LAN" for the attached stations. The data throughput in the switched LAN is greater than a shared LAN. Each concept is discussed below for use in later comparisons.

A shared LAN was the original LAN configuration. Stations connect to a hub, possibly a repeater or concentrator, which forms a broadcast channel for the stations. All stations receive the transmissions of any other station. Station operation is half duplex: there may be only one station originating data on the LAN at one time, determined by the MAC protocol. MAC protocols provide an arbitration scheme which controls station access to the channel. The aggregate of all the stations is limited to that provided by the LAN bit rate.

The same stations described above may instead connect to a switch. There may be one active originator on each switch port at any single instant. The switch copies packets from an input port to the specific output port(s) required while other ports may send or receive a different packet. For a switch with N ports of the same bit rate, the aggregate bandwidth approaches

$$BW = \frac{N * \text{(media-bit-rate)}}{2}$$

When connected directly to a switch,a LAN station may operate in full duplex mode, access control is not necessary. If all the stations use full duplex mode, the aggregate switch bandwidth approaches,

$$BW = N * \text{(media-bit-rate)}$$

Given the existence of multicast packets and contention for individual ports, the actual switched LAN will approach but not actually reach these estimated limits. Internal to the switch, contention by multiple input ports for a particular output port causes congestion.

The station used with a switch or hub may be the same type of station. So, a shared LAN may be upgraded in small increments by replacing hubs with switch ports as growth dictates. Switch ports may be shared by a group of stations with lower requirements while other stations may be assigned to a single port on a switch. This strategy allows for network growth over time, as needed.

Latency in a switched LAN is generally lower than a shared LAN, especially if the network is heavily loaded. A shared LAN causes the latency between a pair of stations to be heavily influenced by the load on the network. Switches provide a form of isolation, but at a cost. For a network operating at low loads, however, the switch may not provide benefit in bandwidth or latency.

Some stations can operate in full duplex mode with a switch, others cannot. Other stations are of a design limited to use with a switch as they don't include a MAC protocol, the situation for stations used with ATM. For more than two ATM stations to connect together, ATM requires the stations connect to a switch. These concepts will be used in the subsequent sections.

## Overview of the Protocols

Each of the four protocols will be described in this section. Three of the four protocols can operate in a shared LAN while the last requires a switching environment. Details on the packet format and access protocol are given. Subsequent sections describe the topology of protocols and contrast the differences between them.

FAST ETHERNET

Fast Ethernet, also called 100BASE-T (the name we'll use herein), is a new protocol under development in the IEEE 802.3 working group. It is an extension of the Ethernet /802.3 MAC protocol [ISO], for operation at a 100 Mbps data rate. Operation with up to 1024 stations is supported on a variety of media types, described later.

The operation and architecture of 100BASE-T is nearly identical to that of its 10 Mbps ancestor. 100 BASE-T supports a broadcast channel made up of point to point media segments connected via repeaters, similar to the topology of 10BASE-T. Stations attach to the network without the need for the MAC to perform any initialization protocol. This simplifies the implementation by a significant amount.

Before initiating a transmission, a station listens to the channel to confirm that the network is not busy. The station then transmits, while monitoring the collision signal to ensure that the transmission has not experienced a collision. If no collision is detected, the station can initiate another transmission after an interframe gap interval of 960 ns. If a collision is detected, the station invokes the binary exponential backoff algorithm, which reschedules the transmission for a randomly selected time in the future.

The key parameter of the CSMA/CD protocol is the slot time, which is the period of time required for a station to be sure that it has not experienced a collision in a properly functioning LAN. This parameter, which determines the minimum packet size, is bounded by the end to end delay of the network as follows. The packet must "fill" the entire network before channel acquisition is ensured, and, if a collision occurs, the collision notification must have time to propagate back to the transmitter and be detected prior to the end of the slot time. Scaling the data rate by a factor of ten while not modifying the MAC's 512 bit slot time requires that the network diameter shrink by a factor of 10. This limits the station to station separation of the network to a maximum of approximately 210 meters.

VG-AnyLAN

VG-AnyLAN, also called 100VG-AnyLAN and Demand Priority, is a new 100 Mbps protocol under development in the IEEE 802.12 working group [AnyLAN]. It supports either the 802.3 or 802.5 frame formats, but the protocol itself is different than either of those protocols. While a network can operate with either frame format, it is configured to run either 802.3 or 802.5 frame format, not both at the same time. Multiple PMD options allow support of many different media types.

The Demand Priority protocol uses a centralized two priority round robin arbitration scheme which is controlled by a central hub. Stations and hubs fan out in a tree from this central arbiter. The central hub passes control of the right to transmit from hub port to hub port of this structure in a round robin manner. The protocol supports two types of transmission requests, referred to as Normal Priority and High Priority. High priority requests take precedence over Normal Priority requests. If a Normal priority request has been pending for longer than 200 to 300 ms it is promoted by the hub to High Priority status and serviced by the High Priority queue.

Each hub maintains a per port address table, which is used to store the address of the station connected to a port. When a packet is received, the hub buffers the packet long enough to determine its destination address, and forwards the packet to the port corresponding to that destination address. The hub also forwards the packet to the cascade port which is the port used to connect a hub to a higher level in the tree, and to all ports which were configured as promiscuous ports at initialization time. All other ports receive an idle signal, which provides a level of security from eavesdropping.

When a station or hub wishes to join the network it initiates a training sequence, consisting of an exchange of frames between itself and the port it attaches to. This training period lasts between 2 to 5 ms, and is used to determine whether the attaching device is a station or hub, whether it uses 802.3 or 802.5 frames, the MAC address of the connecting station, and whether the port will be allowed to operate as a promiscuous listener. While training is occurring, the network suspends its round robin operation for the duration of the training period. Based on the results of the training, a station is either admitted to the network, or informed of the reason why it will not be admitted.

Between packets, stations and hubs send an idle signal to each other indicating that the channel is available to make a transmission request. A station which wishes to transmit a packet signals a transmission request to the hub, indicating the priority of the transmission request. It then awaits a response which grants it the right to transmit one packet onto the network. If the hub that receives the request is the only hub in the network, it waits for the current transmission to complete, and then services the port which is next in the round robin service order. It then grants the right to transmit to that port. If the hub is connected to a hub higher in the tree, it waits for the current network transmission to complete and signals a transmission request to the higher level hub. In this manner, the transmission requests eventually reach the central arbiter, which determines which hub will control the granting of transmissions. When the right to control the network passes to a hub, it services requests in port order (i.e. port 1, then 2 up to port N) and then passes control back to the higher level hub. If a High Priority request is made anywhere on the network, the control passes to the hub which has that request, and returns to the original hub after all High Priority requests have been serviced throughout the network.

FDDI

FDDI is a 100Mb/s LAN which uses a token ring protocol to schedule station access. FDDI is the type of ring where each station removes its own transmission and which operates without a central controlling station. It supports up to 500 stations on a single LAN with a maximum of 100 km of duplex cable. The MAC protocol provides for ring initialization and two categories of service called, "Asynchronous" and "Synchronous" [FMAC]. The initialization and access algorithms are described in more detail below.

Ring initialization must occur prior to the operation and is described first. A station must first physically joint the ring using Station Management (SMT) protocols [SMT]. These protocols test and initialize the link then reconfigure the token path to include the new station. Next, Ring initialization is done using a distributed algorithm called "Claim Token". The claim token process is invoked when the absence of a token is detected. Then, stations send special packets called claim frames which are used to elect a single station to create the new token.

Once the token has been created, it is used to arbitrate shared access for the stations using a timed token protocol [Grow]. To originate information onto the ring, a station must first capture the token which is otherwise circling the closed loop of the ring. The standard provides two criteria for capturing a token which results in the two service priorities. Each station may use either or both priorities, as described below.

The first is synchronous service where each station is assigned a Synchronous Bandwidth Allocation (SBA) using a management protocol [SForum]. The station is allowed each time it sees a token, to capture the token and originate a number of packets related to its allocation. The allocation allows a network administrator to give priority to certain stations over other stations and provides low la-

| | FDDI | 100BASE-T | VG-AnyLAN | ATM |
|---|---|---|---|---|
| # Stations | 500 | 1024 | (unspecified) | implementation limits |
| Access Method | Token Passing | CSMA/CD | Round Robin | Full Duplex |
| packet size | 4500 bytes | 1500 bytes | 1500 or 4500 | 48 byte Cells |
| Extent | 100 km | 210 m | 2.5 km | unlimited |
| Complexity | medium | low | medium | high |

**Table 1: Summary of protocol attributes**

tency access at regular intervals compared to the other class of service, described next.

The asynchronous token protocol divides the bandwidth not consumed by the synchronous allocation among all stations equally. To achieve this, each station maintains a token rotation timer (TRT) which is reset each time the token is seen. When the station next receives the token, it compares the TRT to the value of the Target Token Rotation Time (TTRT), often set to 8 mS [JAIN]. If the current value of TRT is less than TTRT, the station is allowed to capture the token and transmit asynchronous data frames. When TRT >TTRT, the token is said to be "late" and stations limited to asynchronous priority must yield. When the token is late, stations with SBA may capture it giving them priority of access. If the load is high enough to saturate the ring but is mostly asynchronous traffic, maximum access delay for each of N stations is [Jain2]:

- synch delay = 2 * TTRT (nominally, 16 mS)
- asynch delay = N * TTRT

### ATM:

Asynchronous Transfer Mode, ATM, is different than the other protocols described. ATM is connection oriented, point to point, full duplex and uses a small, fixed packet size called a "cell". It is primarily a format for use by switches and includes no access arbitration protocol. Each port on a switch behaves in many ways like a station. The cell header information in a received cell is used to look up forwarding information needed to route the cell within the switch. Error checking is performed on the cell header, and errored cells are discarded. The cell header address information is changed at each switch to represent the route at the next switch. Addressing in the ATM cell is of local significance to a switch, in contrast to the MAC address which identifies individual users either with locally or globally unique values.

Applications which would communicate over the ATM LAN must first establish a Virtual Connection (VC). The VC is a path through one or more switches which provides an end to end pipe to carry the application data. VCs are established in two ways. First, a Permanent Virtual Circuit (PVC) can be manually configured by a network manager. Second, a Switched Virtual Circuit (SVC) may be established by call setup procedures at the time of need [UNI3]. Procedures which are interoperable between switch vendors for SVCs are being developed. These procedures will be sufficient to build ATM networks.

The problems of congestion and ‘VC-routing’ among multiple switches are still outstanding in the standards process. Congestion management is important because a small level of cell loss (e.g., 0.1%) gets magnified to a dramatically large frame loss (e.g., 20%). This is considered unacceptable, and several alternative congestion management policies are under active study and evaluation.

Additional protocols are needed to integrate ATM with legacy LANs such as Ethernet or token rings. IP Routers and bridges can be interconnected with PVCs using an ATM network [RFC1577] and [RFC1483]. Separately, LAN Emulation for ATM is being developed to allow end stations running existing applications to be adapted to ATM services [LANEM]. LAN emulation is necessary because of the large base of applications which require multicast facilities and a broadcast channel.

Other protocols, already in place for ATM, include the ATM Cell format, the ATM Adaptation layer, and the PHY layer. The ATM Cell is used to carry data transmitted between switches [UNI3]. A 48-byte segment of the user data payload is placed in a cell along with a 5 byte header forming the 53 byte ATM cell. The cell header carries the information necessary for switch operation. The cell header contains two address fields, the Virtual Path Identifier (VPI) and the Virtual Circuit Identifier (VCI) which together total 3.5 bytes and define the route of the cell at any particular switch. These fields are updated by each switch in the path.

User information is mapped into the cell payload according to the ATM Adaptation Layer (AAL) protocol. AAL protocols exist for several applications. For variable rate data characteristic of LANs, AAL-5 is often used [CCITT363][ANSIT1S1]. An 8 byte field including data length and an error detection checksum are appended to a frame (or block) of user information up to 64K bytes in length. The AAL-5 PDU is then separated into a stream of smaller PDU's by the Segmentation and Reassembly Sublayer (SAR). Each SAR-PDU is 48 bytes long, the correct size for the ATM Cell payload.

The PHY layer of ATM contains a Convergence Sublayer (CS) and a PMD sublayer. The convergence sublayer maps cells into a PMD service. The convergence protocol provides for delineation of the cells in the serial bit stream. CS sublayer protocols include SONET [Bellcore]. PMD options are discussed in a later section.

Table 1 summarizes the relevant parameters for the four protocols and points discussed in this section. We next discuss the topologies which can be built with these networks. Lastly, we'll compare the protocols based on these descriptions.

## Description of topologies

Until recently, fiber optic cable, and some types of high quality shielded cables were the only media available for delivery of high speed data services to the desktop. Users generally prefer to use twisted pair cabling to their desktops, and reserve fiber cabling for their backbones; this has been an impediment to the acceptance of high speed LANs in the workgroup. Installed UTP cabling typically has a maximum length of 100 meters, and high speed networks have chosen this distance as the maximum UTP segment length. In structured cable installations which adhere to the EIA/TIA 568 wiring standard [EIA] the wiring center may be located in a dedicated closet. For smaller installations, the wiring center may be located with the server under a user's desk.

Table 2 shows a comparison of the supported media types and topologies for the four networks. The topological details are as follows:

The 100 BASE-T topology is significantly different from the 10BASE-T topology in that only one level of the building wiring hierarchy can be supported. Because of timing restrictions, the collision domain can contain no more than two repeaters, which must be separated by no more than 10 meters of cable in the default topology rules. The MAC limit of 1024 stations is not changed. The supported media types are planned to include category 3 UTP, category 5 UTP, and 150 ohm Shielded Twisted Pair (STP) [100BASE-T]. Support for multimode fiber is also included by incorporation of the FDDI PMD [PMD] spec for operation on multimode fiber, but the 2 km distance supported by that PMD is not allowed, due to the timing constraints previously mentioned.

VG-AnyLAN plans to support category 3 UTP, category 5 UTP and 150 ohm STP, as well as fiber optic links of 2km at 1300 nm, and 500 m at 850 nm [AnyLAN]. It is the only technology planning to support 25 pair category 3 cabling but, for this, the hub must include a store-and-forward function and the station is limited to half duplex operation. In The topology consists of cascaded hubs, as in the 10BASE-T architecture. The supportable end to end distance is not explicitly defined in the standard, but will cover an end to end extent of approximately 2.5 km, in a three level hierarchy [Grinham]. The maximum number of stations supported by the standard is also not explicitly specified.

FDDI provides support for category 5 UTP and 150 ohm STP [TP-PMD], multimode fiber [PMD] as well as singlemode fiber [SMF-PMD]. The topology is a dual ring of trees, in which the FDDI dual ring is extended to a tree structure by the use of concentrators. The number of levels in the tree is limited only by the allowable station count and the maximum extent of the network. FDDI protocols provide relatively thorough fault isolation capabilities as well as standardized ways of adding fault tolorance in the dual ring and, in trees, with dual homing [Hutchison][Ocheltree][Willebeek] . The total network extent cannot exceed 100 km, with a maximum of 500 stations.

ATM at 155 Mbps will include support for category 5 UTP, 150 ohm STP, multimode fiber [UNI3.1] and singlemode fiber [Bellcore]. The topology is a mesh of switches, which means a switch in the network may be reachable from another point using multiple routes involving independent sets of switches. The full exploitation of this flexibility requires completion of the switch to switch connection protocols currently in development.

| | 100BASE-T | VG-AnyLAN | FDDI | ATM |
|---|---|---|---|---|
| Topology | Star Wired | Hierarchical Star | Dual Ring of Trees | Mesh of Switches |
| Category 3 UTP | 100 m | 100 m | | |
| Category 5 UTP | 100 m | 100 m | 100 m | 100 m |
| 150 ohm STP | 100 m | 100 m | 100 m | 100 m |
| Multimode Fiber | | 2 km | 2 km | 2 km |
| Singlemode Fiber | | | 60 km | 40 km |

**Table 2: Summary of topology attributes**

ATM also provides a virtually unlimited end to end extent, since no MAC protocol sets the timing constraints for end to end extent.

Each of the technologies discussed in this paper will probably connect to other LAN types or a backbone of some kind. Some important properties for high speed backbones are the supportable end to end distance, the robustness and availability of the technology, and the incremental cost for connecting the workgroup to the backbone. In addition, the ability to minimize complexity through the use of one high speed solution to unite the high speed workgroups, as well as the existing low speed workgroups, should be considered. FDDI's reliability features, supported media types, and ability to support a large number of stations over a 100 km extent makes it a good choice for the backbone. ATM's future capabilities for a mesh topology, which will allow redundant connections to enhance reliability, and support for long distances will make it an ideal choice for a backbone technology.

## Performance:

We analyzed the four LAN technologies for protocol efficiency to compare the bandwidth available to transport data given the MAC protocols. The actual bandwidth available for user data is the claimed bit rate derated by the efficiency, above which the LAN is saturated and additional load cannot be carried. New models are presented herein for the VG-AnyLAN and FDDI protocols.

The topology for comparison is a work group of limited extent. Due to differences in the offerings of each protocol, comparison is somewhat difficult. The configuration should provide stations at maximum separation for a worst case analysis but no two protocols have the same distance limit. We analyzed VG for two sizes based on physical extent - one size being moderately large for VG, the other being smaller, the maximum size of 100BASE-T. FDDI extent was chosen to be the same as the large VG for comparison. ATM, not being a shared access protocol, is configured as a single link between a switch and a station. The physical characteristics of the topologies are summarized in Table 3:

| | VG-AnyLAN | | 100BaseT | FDDI | ATM |
|---|---|---|---|---|---|
| station cnt | 20 | 20 | 20 | 20 | 1 |
| hub count | 1 | 3 | 2 | 3 | 1 |
| Diameter | 200m | 2.2km | 210 m | 2.2km | N/A |

**Table 3 : Summary of topologies analyzed**

The large VG-AnyLAN configuration, also used for FDDI, is shown in Figure 1: three hubs with one as the root, the hubs connected by 1km duplex cables. Stations are attached to the lower two hubs using 100 meter duplex cables. The overall span, or diameter, of both trees is the same and chosen to be 2.2km, near the limit of VG-AnyLAN [Grinham].

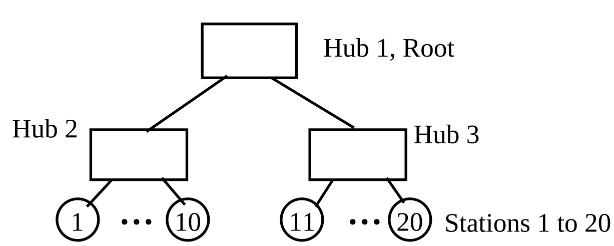


**Figure 1: Configuration for VG-Anylan and FDDI**

The 100BaseT topology, shown in Figure 2, contains two hubs, each connecting to stations with 100 meter duplex cable segments. The two hubs have a 10 meter cable connecting between them. This is roughly the maximum diameter for 100BASE-T.

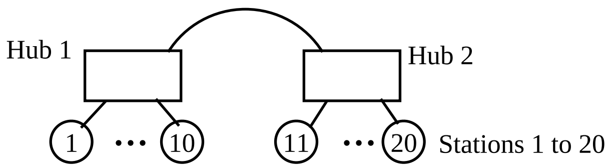


**Figure 2: Configuration analyzed for 100BaseT**

Details of the models:

Protocol efficiency is simply the time to transmit a frame, F, divided by the F plus overhead (OH)associated with the transmission:

$$E = \text{efficiency} = F / ( F + OH)$$

We tried to compare equal conditions. The analysis assumes that all stations on the LAN transmit packets as fast as allowed on the channel. Conceptually, any single station is capable of saturating the LAN. This gives a clear view of the protocol efficiency.

VG-AnyLAN:

We developed a model for VG-AnyLAN that accounts for cascaded hubs. A model for a single hub has been presented earlier [Kadambi]. Our model covers the more general tree topology. The model accounts for overhead incurred due to the processing by the hub of station requests to transmit, the address decode operation in the hubs, arbitration steps in the root hub of the tree, and the propagation delay of the media. Each is discussed below.

The model for the overhead for the multiple hub case was based on the overhead to transmit a packet after requesting and obtaining a grant from the hubs at all levels of the network up to the root hub. This models the case where a receiving station of a packet is on a destination hub, which also has the transmitting station for the next packet. This pattern of traffic flow may be typical of applications in a client/server environment.

A station must request access for each transmit packet and receive an acknowledgment. The time for this process with a single hub includes the time for the hub to process a request (t_req) plus the time to process the acknowledgment from the hub permitting the transmission.

$$t\_ack = t\_req + t\_grant$$

The t_ack time in the cascaded case is simply L times the value for a single hub where L is the number of levels of hubs:

$$t\_ack\_mult = L * t\_ack$$

The draft specification, [VG] does not yet include these key timer limits for the request/grant process. From the standards definition, we have carefully chosen reasonable values for t_req and t_grant to be 1.5 microseconds each.

The next contribution to overhead is the time for the hub to recognize the destination address in the packet. During this process the hub buffers the packet. This overhead, t_adrs, is chosen based on the need to receive approximately 14 bytes plus a lookup, the total of which we expect to be about 1.2 microseconds.

The total hub address lookup overhead is a little more complex: the address lookup has to progress up L levels of the cascade to the root hub, and then has to progress downwards to the 'destination hub' before it reaches the destination station. Thus, the individual packet incurs the store-and-forward delays at 2L-1 hubs (L hubs upstream and (L-1) hubs downstream). Thus,

$$t\_adrs\_mult = L * (t\_adrs) + (L - 1) * (t\_adrs).$$

The propagation delay is just the round trip propagation delay to the root hub and back from the transmitting end-station. Thus, for example, 2.2km of media at 5 uS/km contributes a delay, t_prop = 11 microseconds.

The efficiency also involves the time to transmit a packet,

$$F = (\text{ packet size in bytes}) * 80 \text{ nS/byte}.$$

and frame format overhead, (8 bytes preamble and 2 bytes of ending delimiter),

$$Tped = (10 \text{ bytes}) * 80 \text{ nS/byte}$$

Then,

$$E = \frac{F}{(F + Tped + t\_ack\_mult + t\_adrs\_mult + t\_prop)}$$

Results are presented for two cases, see Table 4. First, the results of the model for a single hub and 200 meter diameter are given. It can be seen that VG efficiency is similar to 100BaseT in this small configuration. 100BaseT is > 77% efficient for normal packet sizes, roughly 256 bytes or greater, which is acceptable.

Secondly, the results for the three hub case show that

| Packet Size (bytes) | VG-AnyLAN 1 hub, 200 m | VG-AnyLAN 3 hubs, 2.2 km | 100-BaseT (210 meter) | FDDI 3 hubs, 2.2 km | ATM (1 Link) |
|---|---|---|---|---|---|
| 64 | 46 % | 19 % | 65 % | 84 % | 58 % |
| 128 | 63 % | 32 % | 74 % | 91 % | 78 % |
| 256 | 77 % | 49 % | 80 % | 96 % | 78 % |
| 512 | 87 % | 66 % | 83 % | 98 % | 85 % |
| 1024 | 93 % | 79 % | 86 % | 99 % | 85 % |
| 1518 | 95 % | 85 % | 87 % | 99 % | 86 % |

**Table 4: Efficiency of protocol as a function of packet size.**

cascading, media delay, and small packet size have significant impact on protocol for this traffic pattern. We considered an alternative traffic pattern where the next station to transmit a packet is on a different hub than the destination of the previous packet. The efficiency improves significantly as this is almost a best case pattern. We found, for 3-hubs, the efficiency for 1518 byte packets reaches 94%, and for 64 byte packets to 40%.

100BASE-T:

We used a CSIM based simulation to examine the performance of the 100 Mbps Ethernet [CSIM]. Details of the station awaiting the channel to be idle before transmission and waiting for an IPG interval before beginning to transmit were modeled. The collision window was also modeled so that any two stations beginning transmit within that interval relative to each other experience a collision. The physical extent of the Ethernet was chosen to be the maximum allowed (5.12 microseconds slot time) which results in the highest penalty in throughput when a collision occurs. The number of stations was 13 because of simulation limitations. The results are based on 5 second simulation runs for each case. The results, calculated above, were then extrapolated to 20 stations using the known variation of efficiency with station count [Rama].

The dependence of efficiency on packet size is shown in Table 4. It is seen that very small packets yield low efficiency, but that the results are similar to the VG-AnyLAN results for the 200m case. 100BASE-T is > 80% efficient for normal packet sizes, roughly 256 bytes or greater, which is acceptable.

FDDI:

The often cited model for FDDI efficiency is [Jain]:

$$E = \frac{N * K * G}{N * (K * G + D) + D}$$

$$K = \frac{T - D}{G}$$

Where N is the number of stations, G is the frame size used in the reference, and D is the ring delay. The model is accurate for G larger than about 1K byte but is easily seen to be inaccurate at small packet sizes. This can be corrected by accounting for the packet preamble and ending delimiters (Tped) more correctly. Given,

$$F = G - Tped$$

The reference equation becomes

$$E = \frac{N * K * (G - Tped)}{N * (K * G + D) + D}$$

The equation can be further reduced for many cases where D is small. The term in denominator, (K*G +D) is simply TTRT and it can be seen that N*TTRT will be much bigger than the last term, D. Note that D scales with the number of stations, so, the above reduces to:

$$E \doteq \frac{K * F}{TTRT} = \frac{TTRT - D}{TTRT} * \frac{F}{(F + P)}$$

The latency of the ring for the configuration in Figure 1 is,

$$D = 2km * (M - 1) * 5uS/km + 1uS * (M + 1) + (.1m * 2 * 5uS/km + 2 uS) * N$$

Where M is the number of Concentrators. Table 4 shows that FDDI is very efficient for the topologies considered here, independent of packet size. For normal packet sizes, roughly 256 bytes or greater, it provides 16% to 30% greater throughput than other protocols of the same media bit rate.

ATM efficiency is evaluated only for the AAL, ATM and PHY layer protocols operating on a single link. Other protocols involved in ATM are not strongly related to this aspect of performance. The overhead affecting a single link is calculated based on AAL 5, the ATM Cell, and Sonet as described earlier. AAL requires 8 bytes of overhead per packet and Sonet requires about 3% of the link. The strongest impact is the 5 bytes of overhead per cell. It can be seen in Table 4 that ATM is somewhat lower in efficiency even for 1500 byte packets. This, due to the small cell size that was chosen for other reasons.

### Summary of Performance Analysis

The simple comparison of the efficiency of the four different LAN protocols is summarized in Table 4. Note that the environment considered across all of the technologies was that of a small workgroup of stations that have essentially an infinite load to transmit on the channel. The more efficient a particular technology, the better the overall user perceived performance. Packet sizes currently observed on existing networks are predominantly small, in the range of 512 bytes or less. Although this is anticipated to change in the future, high speed LANs deployed in the immediate future need to be reasonably efficient for small packet sizes. From this perspective, FDDI appears to provide better efficiency as compared to the other high speed LANs. When considering these 'worst case' workloads and configuration for the workgroup, the efficiency of VG-AnyLAN is almost comparable to what is observed with 100BASE-T. The typical claim of poor efficiency with CSMA/CD protocols for small packet sizes is usually with a large number of stations. We observe that for a small workgroup LAN configuration (200 meters) that VG-AnyLAN and 100BASE-T efficiency is acceptable. The penalty of ATM (5 bytes of overhead for every 48 bytes of payload) is also substantial for workloads of small packets.

## Summary and Comparisons

We next make some observations separated into four topics of comparison. These topics are: the maturity of the technology, the cost of implementation, latency control based on a priority scheme, and the applicability of the technology to workgroup and backbone networks.

Of the four technologies, only FDDI is currently a mature technology with completed standards and proven interoperability. The other three protocols are at various stages of standards and technology development. For example, neither Fast Ethernet nor VG-AnyLAN have passed a letter ballot within their committees. ATM is still working out issues around congestion management, switch to switch protocols, and LAN Emulation, are still being developed. Typically, it takes some time for a standard to reach its maturity stage where relatively large scale implementations achieve the interoperability, stability, cost, and performance goals. 100BASE-T, VG-AnyLAN, and ATM are yet to achieve the same stage as in Ethernet, Token Ring, and FDDI.

Protocol complexity can be used as a basis for cost comparison, and is divided into two aspects - station costs and hub costs. Of the four protocols, 100BASE-T and VG-AnyLAN have relatively simpler station protocols, which can result in less cost. ATM is more likely to have higher relative cost due to cells assembly and additional services. On the hub side, the design complexity and level of semiconductor integration for 100BASE-T, FDDI and VG-AnyLAN may be relatively close, resulting in similar cost structures. ATM switches, currently expensive, will drop in price but are expected to be more expensive than the shared LAN's hub costs. However, the additional bandwidth and services provided by ATM may justify the added complexity and cost. Finally, note that additional cost, not directly related to an individual protocol, must be considered when configuring an infrastructure network where dissimilar LAN types interconnect the desktop, server, and backbone.

A form of priority service for latency control is provided by three of the protocols, the exception is 100BASE-T. The VG-AnyLAN protocol has Demand Priority and FDDI has synchronous priority. These two protocols have or plan to have bandwidth allocation and mangament schemes. ATM's UNI protocol provides multiple service classes and bandwidth reservation established during call setup of a switched VC. One issue in deploying these services is the need to coordinate between switches or LAN segments the bandwidth allocations made locally. Another issue in deployment of these services is operating system support for the multiple classes of service. Today's higher level protocols generally request that a packet be transmitted with no notion of varying classes of service on the LAN. Therefore, while the underlying datalink technologies provide a service, the usefulness depends upon development of new or modified higher layer protocols.

The design of backbone and server interconnects requires many special considerations. Key design considerations include high and scalable bandwidth, scalable network extent, robustness and stability under high load, availability, fault tolerant topology, and manageability. Two of the protocols, well suited for use in the workgroup, are not as well suited for the backbone. For 100BASE-T, the distances supported are too limited. For VG-AnyLAN, the centralized arbitration, rooted topology, and the lack of fault tolerant features (such as redundancy in hub or link) makes it less desirable as a backbone technology. FDDI supports large network distance, with good network efficiency, and several fault tolerant capabilities. Although, it's bandwidth is not scalable, it is a good choice as a backbone technology. ATM meets most of the backbone requirements, however, there is a *fair* amount of standard and technology work before these capabilities can be fully realized.

## Conclusions

Four protocols intended to provide 100Mb/s LAN services to individual users or work-groups have been described and compared. We introduced new or extended models for the VGAnyLAN and FDDI protocols.

Comparisons are made using several criteria: maturity, simplicity, cost, scalability, quality of service, and efficiency.

While vendors are pushing various technology driven solutions, the price, maturity, ease of use, and interoperability of a given solution will determine it's fate in the market.

## References:

[100BASE-T] IEEE P802.3/d1 "MAC Parameters, Physical Layer, Medium Attachent Units and Repeater for 100 Mb/s Operation", Draft Standard, 17 June 1994.

[AnyLAN] IEEE P802.12/D3, " Draft Standard, Information Technology, Local and Metropolitan Networks - Part 12; Demand Priority Access Method and Physical Layer Specifications" , 28 June 1994.

[ANSIT1S1] T1S1.5/91-449, "A New High Speed Data Transfer AAL", T1S1 working group proposal, 8 November 1991.

[Bellcore] Bell Communications Research, Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria, Document Number TA-NWT-000253, Issue 8, Oct., 1993.

[CCITT363] ITU/CCITT, B-ISDN ATM Adaptation Layer Specification, Recommendation I.363.

[CSIM] H. Schwetman, "CSIM Reference Manual", MCC Technical Report, ACA-ST-257-87 Rev 14, March 1990.

[EIA] EIA/TIA 568-1991, Commercial Building Telecommunications Building Wiring Standard.

[FMAC] ANSI X3.139-1987, Fiber Distributed Data Interface (FDDI) - Token Ring Media Access Control (MAC).

[Grinham] J. Grinham & M. Spratt, "IEEE 802.12 Demand Priority and Multimedia", Proceedings of the 4th International Workshop on Network Operating Systems Support for Digital Audio and Video Lancaster, U.K.,(November 1990), pp75-86.

[GROW] R. Grow, "A Timed Token Protocol for Local Area Networks", Proceedings Electro 82, paper 17/3.

[Hutchison] J. Hutchison, "The Role of Concentrators in FDDI Rings", Proceedings of IEEE 15th Conference on Local Computer Networks, Minneapolis, MN (October 1990), pp. 24 - 40.

[ISO] ISO 8802-3:1989, Information processing systems - Local area networks - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.

[JAIN] R. Jain, "Performance Analysis of FDDI Token Ring Networks: Effect of Parameters and Guidelines for Setting TTRT", Proceedings, ACM SIGCOMM '90, Philadelphia, PA pp. 264-275.
[JAIN2] R. Jain, FDDI Handbook: High Speed Networking Using Fiber and Other Media, Adison Wesley, Reading MA,

[Kadambi] J. Kadambi & S. Barilovits, "Analysis of 100BASE-VG Demand Priority Protocol", Presentation to IEEE 802.3, 8 March 1993.

[LANEM] ATM_Forum/94-0035R4, "LAN Emulation Over ATM" Draft Rev. 4, 1 July 1994.

[Ocheltree] Keneth Ocheltree, "Using Redundancy in FDDI Networks," Proceedings of IEEE 15th Conference on Local Computer Networks, Minneapolis, MN (October 1990), pp. 261 - 267.

[PMD] ANSI X3.166-1990, Fiber Distributed Data Interface (FDDI) - Physical Layer Medium Dependent (PMD).

[RAMA] K. Ramakrishnan & H. Yang "The Ethernet Capture Effect: Analysis and Solution" Proceedings, IEEE 19th Conference on Local Computer Networks, Minneapolis, MN (Oct. 1994).

[RFC1483] Internet Engineering Task Force (IETF), "RFC 1483, Multiprotocol Encapsulation Over ATM Adaptation Layer 5" , July 1993.

[RFC1577] Internet Engineering Task Force (IETF), "RFC 1577, Classical IP and ARP Over ATM", January 1994.

[SForum] FDDI Synchronous Forum, "FDDI Synchronous Forum Implementers' Agreement, Revision 2", 1 December 1993.

[SMF-PMD] ANSI X3.184-1993, Fiber Distributed Data Interface (FDDI) - Single Mode Fiber Physical Layer Medium Dependent (SMF-PMD).

[SMT] ANSI X3.229-1994, Fiber Distributed Data Interface (FDDI) - Station Management (SMT).

[TP-PMD] X3T9.5/93-022 " Fiber Distributed Data Interface (FDDI) - Twisted Pair Physical Layer Medium Dependent (TP-PMD) Rev 2.1", 1 March 1994.

[UNI] The ATM Forum, "ATM User Network Interface, Version 3.1 (UNI3.1) Specification ", draft of 28 May 1994.

[UNI3] The ATM Forum, "ATM User Network Interface, Version 3.0 Specification ", September 1993.

[Willebeek] M. Willebeek-LeMaire & A. Ruiz, "FDDI Network Availability Analysis", Proceedings of IEEE 16th Conference on Local Computer Networks, Minneapolis, MN (October 1991), pp. 24 - 33.