Title
FDDI AS A BACKBONE FOR LARGE CAMPUS ETHERNET NETWORKS

Author
Fink, R.L.

Publication Date
1987-08-01

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA, BERKELEY

Information and Computing Sciences Division

To be presented at the 11th International Fiber Optic Communications and Local Area Networks Conference, Anaheim, CA, October 26–30, 1987

FDDI as a Backbone for Large Campus Ethernet Networks

R.L. Fink

August 1987

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

FDDI As A Backbone For Large Campus Ethernet Networks*

Robert L. Fink
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

Abstract: Many large campus environments rely heavily on multiple Ethernet LANs, interconnected with bridges and/or routers, operating over Backbone LANs also based on Ethernet. Although the use of Ethernets for direct user attachments will be viable for quite some time to come, the use of Ethernets in the Backbone is more limited. An FDDI Backbone, interconnected to existing Ethernets with FDDI to Ethernet MAC level bridges, will provide the principal evolutionary path for these large campus networks due to its combination of high speed, low delay time and reliability.

Large Campus Networks

Large campus network environments include those that are used for educational, research and/or commercial purposes. These sites are best characterized as those multiple building sites that are too large for a single LAN (Local Area Network) to span, have some practical way to be interconnected with private-use media (even if it is provided by a common carrier) and are administratively practical to manage as a single network. There are good examples of large campus networks in all the above environments, and it is reasonable to expect the number to grow.

Ethernet [1] (also known as IEEE 802.3 [2]), is the most widely used LAN in large campus networks today. This is primarily a result of its early acceptance as a standard by several large manufacturers and by a respected standards body (the IEEE). A case is not made here that Ethernet is perfect, or that it is the only LAN standard that should be used, only that it is widely used now and will continue to be a major LAN alternative for a long time to come (at least 10 years). General conclusions reached in this paper could equally well apply to other LAN technologies also coming into wide use (e.g., IEEE 802.4 and 802.5).

Constraints on the size of an Ethernet

The design of Ethernet has constraints that result in the necessity of creating a large campus network from a number of small, interconnected Ethernets, rather than attempting to build a single large network.

Ethernet was defined for use in quite large environments to begin with. Two Ethernet attached hosts (e.g., a work station and a file server) can be as far apart as 2800 meters with the use of Ethernet remote repeaters. Practical routing constraints in buildings limit this distance to about 1000 meters in diameter "as the crow flies".

Ethernet presents a physical hazard if its coaxial cable is extended too far, since it is a single point of ground system. If an Ethernet coax is used to interconnect two building with separate power/ground systems, which is common in many large campus environments, personnel and equipment hazard situations arise, dictating smaller Ethernet implementations.

Ethernet has limited bandwidth. Just how limited has long been a matter of contention, partly due to its non-deterministic collision avoidance scheme. Although an Ethernet may theoretically carry over 16,000 packets per second, the practical limit is far below this; 1000 packets per second is often used as a realistic limit.

In fact, Ethernet capacity depends mainly on the protocols being used, how the attached hosts implement them, and the use being made of them at any given time. Variations in these factors account for the widely divergent claims made about Ethernet performance. What remains constant is that Ethernet has practical limits that dictate dividing a large Ethernet into smaller ones to minimize traffic overload on any single Ethernet.

A significant issue concerning Ethernet size is the complexity of a single, physical Ethernet from a maintenance perspective. Since Ethernet is a "bus" design, any host on the bus can cause an electrical disturbance through interface hardware failure that can affect the entire Ethernet. The fewer hosts on an Ethernet, including all of its repeatered segments, the more maintainable it is.

Backbone Networks and Interconnection Systems

Interconnection of several Ethernets in a large campus network is done with a network whose primary purpose is that of interconnection; networks of this class are often called Backbone networks.

Backbone networks capable of interconnecting multiple Ethernets must, as a general rule, have at least as good a performance as the Ethernets they interconnect. This is due to the requirement that all hosts perform as if

* This work was supported by the U.S. Department of Energy under contract No. DE-AC03-76SF00098.

they were on a single Ethernet, and that hosts be accessible anywhere. This requirement is becoming even more stringent with the advent of workstations that rely on network file servers, a trend that will continue.

Equally important as the choice of a Backbone network is the choice of the systems used to interconnect the local Ethernets to the Backbone. These interconnection systems must have high performance and low delay. The most common approaches to this have been *routers* and *bridges*.

Routers and bridges send packets from one network to another based on routing algorithms and knowledge about network topology. Routers are ISO OSI [3] Network Layer interconnection systems that interact directly with a host trying to reach another host through the router. Since the host is actively communicating with the router the process is protocol sensitive, i.e., the host and router must use the same protocol at the network layer to provide the routing function from one network to another. Since routers are protocol sensitive they typically only support one protocol at the network layer (e.g., the Decnet routing layer). Routers may allow quite complex multiple path interconnection between LANs, but at the cost of limiting the large campus network to one, or a few, network protocols. Performance and cost of routers have been a problem for large campus networks because the complexity and instability of related internetwork standards has resulted in inefficient, and sometimes faulty implementation of the protocols used. This has not deterred the use of router technology in many large campus networks. Network routers are often only a small part of the overall cost of building and operating large campus networks, and in many cases it has been acceptable to operate only one network protocol over the network. As evolution continues toward fewer, more stable and standardized internetwork protocols, normal development will result in faster and cheaper routers.

Bridges are ISO OSI Data Link Layer interconnection systems that interact indirectly with a host trying to reach another. The bridge inspects all packets sent by all hosts on one of its attached networks for possible routing to its other attached network. Since the host is not actively communicating with the bridge the process is not protocol sensitive, i.e., the bridge can support any protocol above the data link layer. The performance of bridges can be very fast since the routing algorithm is simple, well understood and implementable in very fast hardware. One example, the Digital Lan Bridge 100™, performs at the full Ethernet bandwidth in certain situations. Bridges are also less expensive than routers, costing 25% to 50% of the cost of a router usable in a similar environment. Bridges typically have limited configuration possibilities due to the desire to keep the implementation fast and cheap. The result is usually a spanning tree network in which host to host communication involving different LANs can only be done over one path, i.e., multiple active paths between LANs do not exist. This single characteristic forces a higher bandwidth, or packet carrying capacity, on the Backbone than might be required of a router based network which may be considerably more complex in topology.

Most bridges in use today interconnect two Ethernets, leading to the natural choice of Ethernet as the Backbone network technology.

Experience from one implementation

At the Lawrence Berkeley Laboratory (LBL) of the University of California we have a physical site that extends approximately a mile by a half mile, with over 100 buildings housing over 3000 staff. The first use of Ethernet (in 1982) was a single coaxial cable in one building. In 1984, this was expanded to a single very large Ethernet using optical fiber based remote Ethernet repeaters reaching four major building complexes. By 1985, the rapidly approaching practical limits of topology, maintainability and performance led to the introduction of bridges. LBL became a test site for the Digital Lan Bridge 100, expanded the fiber plant to several more buildings and started the evolution toward a bridge based large campus network. There are currently ten large Ethernets with over 300 hosts attached, all interconnected with bridges using an Ethernet based Backbone.

By the end of 1988 LBL will have installed over 150 fiber miles of optical fiber cable interconnecting the entire Laboratory complex, including several LBL buildings on the adjacent UC Berkeley campus. Bridges will be used to interconnect more Ethernets over these optical fibers.

Traffic growth on this network, called LBLnet, continues to be fueled by the fast evolution of the personal computer, workstation and distributed medium scale computing systems in use at LBL. The largest recent contributor to this growth has been SUN workstations and the use of NFS (Network File System) over the Ethernet.

Overall traffic growth on LBLnet is larger than that on the Backbone, but the LBLnet Backbone traffic is steadily growing and has now exceeded 50 percent of the practical capacity of an Ethernet. We expect to exceed the capacity of an Ethernet based Backbone in the next two to three years.

A better Backbone technology needed

A better Backbone technology than Ethernet is needed for both higher performance and better reliability. Higher performance is a function of both packet per second capacity and packet delay. Raw speed (i.e., pps capacity) is necessary to support the higher rates supported by newer systems. Reduction of packet delay is important to minimize protocol problems associated with round trip delays. Fortunately, in most LAN technologies low delay is a positive side effect of the higher performance.

System reliability depends upon inherent reliability and redundancy. New LAN technology must be inherently reliable or the maintenance burden in the large environments to which we are evolving will become unwieldly. Redundancy is essential to provide for continued operation when the inevitable failures occur, which are most typically media tap, connector and splice oriented.

It is a widely held opinion that optical fiber token rings offer the best alternatives for future high performance LANs. One existing implementation of this technology is the ProNET 80 by Proteon Inc. This product is just coming into use as a router based Backbone network. It operates at 80 M bits per second over optical fiber and provides redundant paths between nodes. However, it is not an industry standard and currently only operates with a router.

FDDI, the Fiber Distributed Data Interface, is an emerging American National Standard being developed by the Accredited Standards Committee X3T9.5 [6]. FDDI is being accepted by industry as the next generation high performance LAN due to its high performance, reliability and potential for low cost of manufacturing. It has a transmission rate of 100 M bits per second and is expected to sustain performance at 80% to 95% of this rate.

FDDI has been designed for extremely reliable performance in typical failure situations. For example, it will have optical switches that allow FDDI interfaces to self-reconfigure and to use redundant optical fiber paths, plus a management protocol that has been designed to take advantage of this capability. A typical FDDI network will be designed with two separate counter-rotating rings that are capable of completely independent operation, but which can be reconfigured into one large ring to compensate for a failure in any portion of the ring.

Component technology is developing apace with FDDI so that highly reliable LED transmitters and PIN diode receivers will be available when needed. In addition, highly reliable optical fiber splicing and connector technology is now available.

Many exisiting large campus Ethernet environments are using optical fiber based repeaters to interconnect multiple Ethernets. This has resulted in large plant investments in several different types of optical fibers. It is expected that FDDI will successfully operate over older optical fiber plants in most cases, even when differing sizes have been used. For example, LBL initially installed 100 micron multi-mode optical fiber, while newer installations are 62.5 micron. Both will work with FDDI. In addition, FDDI can be configured to accomodate existing physical layouts of cable due to its flexible ring topology. This reutilization aspect of FDDI can result in large plant savings.

A new bridge technology

A new bridge technology will be required to interconnect FDDI backbone LANs to Ethernet LANs. The interconnection of dissimilar LANs is a relatively new technology. Vitalink™ has done some work in this area with their TransLAN™, which uses lower speed point to point links to interconnect Ethernets. Also, Applitek™ has developed UniLAN™, which operates over different LAN media.

There are two ways an Ethernet to FDDI bridge might be developed: by translation or by encapsulation [7]. In translation bridges, two dissimilar LAN technologies are interconnected by function/address/format translation mapping. This is possible if the two LAN technologies are similar enough to easily map one onto the other. The advantage of this method is that hosts connected to either LAN can interoperate, e.g., a large Cray host can attach directly to FDDI while workstations can remain attached to Ethernet.

In encapsulation bridges, frames received are encapsulated and retransmitted across a LAN of disimilar type from the originating LAN, to be decapsulated by a distant bridge for transmission onto a LAN of the same type as the original. The disadvantage of this approach is that communication can take place only between hosts on the same type of LAN. Encapsulation bridges might be required if more than FDDI and Ethernet LANs are being interconnected. In this type of bridge FDDI would only be useable as a pure Backbone.

Conclusion

In many campus environments Ethernets are reaching their limits as Backbones. FDDI optical fiber token rings will offer an attractive match of performance and reliability for use as Backbones in large campus networks.

To make effective use of FDDI, however, a new generation of bridges needs to be developed that solves the problem of interconnection of Ethernets and FDDI. It appears now that the best possible solution would be a bridge that allowed for hosts located on FDDI LANs to communicate with hosts located on Ethernet LANs, i.e., a translation bridge.

References

[1] "The Ethernet - A Local Area Network, Data Link Layer and Physical Layer Specifications", Digital Equipment Corp., Intel Corp., Xerox Corp., Version 2.0, November, 1982.

[2] "An American National Standard - IEEE Standards for Local Area Networks -Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specification", ANSI/IEEE Std 802.3 - 1985.

[3] "OSI Reference Model - The ISO Model of Architecture for Open Systems Interconnection", H. Zimmerman, IEEE Transactions of Communications, vol. COM-28, April, 1980.

[4] "Bridges smooth troubled waters for wide-area networking", J. Hart, Vitalink Communications Corp., Data Communications, vol. 14, no. 3, March, 1985.

[5] "Remote Bridge Management Software Guide", Digital Equipment Corp., Order No. AA-FY93A-TE, May, 1986.

[6] "FDDI - A Tutorial", Floyd E. Ross, Unisys, IEEE Communications Magazine, vol. 24, no. 5, May, 1986.

[7] "Transparent Interconnection of Incompatible Local Area Networks Using Bridges", George Varghese and Radia Perlman, Digital Equipment Corp., ACM SigComm '87 Workshop, Stowe, Vermont, 11-13 August, 1987.