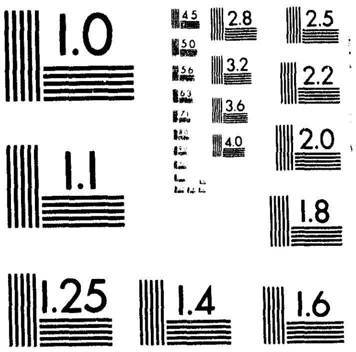
1.0
2.8
2.5
3.2
2.2
3.6
4.0
2.0
1.1
1.8
1.25
1.4
1.6

1 of 1

Conf-940571--1

LA-UR- 93-4438

Title: HIPPI, FIBRE CHANNEL, AND ATM
AS GIGABIT LANs

Author(s): Don E. Tolmie

Submitted to: NETWORLD + INTEROP'94 ENGINEER"S CONFERENCE
LAS VEGAS, NEVADA
MAY 2-6, 1994

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

MASTER

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Form No. 836 R5
ST 2629 10/91

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

HIPPI, Fibre Channel, and ATM as Gigabit/s LANs

Don Tolmie
Los Alamos National Laboratory
C-5, MS-B255
Los Alamos, NM 87545
Phone 505-667-5502, FAX 505-665-7793
Internet: det@lanl.gov

Abstract

Computer networks that operate in the gigabit per second speed range are becoming very important for interconnecting supercomputers and other high end equipment. Some trends and applications are examined and criteria for selecting an interconnection technology are developed. HIPPI is the current interface of choice, while Fibre Channel and ATM are emerging standards. These systems are examined as to their backgrounds, advantages, and shortcomings.

1: Needs

As computers get faster and more powerful, the networks that interconnect them must also increase in speed to keep a computer from going idle while data is being transferred to or from the computer. Not only must the networks be fast, but there may also be requirements for low latency from the sender to the receiver, for small as well as large messages, for long distances, or for continuous rate transmission such as voice or video.

File transfers are common occurrences, and file sizes always seem to grow with time. Files are growing from kilobyte size to megabyte size. File transfers are normally characterized as (1) large volume transfers, (2) not time critical, i.e., get it there as soon as possible but it doesn't have to start in the next microsecond, and (3) make it very reliable.

Remote procedure calls are often used to communicate between the processors of a multi-processor system. For example, a single problem is split up so that different parts of it execute in different processors, all at the same time. Remote procedure calls can be characterized as requiring (1) normally small amounts of data, (2) low latency, and (3) high reliability.

Voice and video transmissions are starting to be mixed with computer data, with the result called "multi-media". Transmission of voice or video requires (1) small bandwidth for voice and high bandwidth for video, (2) latency is not a problem, but any delay in the network should be consistent, and (3) dropping a small amount of voice or video is usually not fatal.

Computer visualization, e.g., computer movies of simulations in progress, have been shown to significantly improve a user's productivity by giving them additional insight into the calculation. Visualization requires (1) large bandwidths, e.g., 10's of megabits per second, (2) consistent delay, and (3) moderate reliability.

Some people are questioning the need for strict time-dependent isochronous data in local area networks (LANs). Present experiments with sending voice or video data using datagram protocols, e.g., using TCP/IP over Ethernet, have been quite successful. Hence, multi-media applications may not require isochronous transfer support.

2: Trends

The trend is for communications bandwidths to increase as computer power increases. It wasn't many years ago that 9600 baud (approximately 8 Kbit/s of data) was considered high speed. Today Ethernet runs at 10 Mbit/s, Fibre Distributed Data Interface (FDDI) at 100 Mbit/s, and the High-Performance Parallel Interface (HIPPI) at 800 Mbit/s.

Distances are also increasing. The era of stand-alone machines is ending, and interconnection is the norm, not only within a computer complex, but across the country and across the world. Long distance communications increase the end to end delay due to the speed of light limitation. For example, the round-trip time between New York and California is on the order of 35 milliseconds. This delay limits the applications and protocols that are practical over long distances. For example, a protocol that expected a hand-shake acknowledge message after each kilobyte transferred would be very inefficient on a high bandwidth communications link.

Local area networks are starting to use switches to interconnect end points rather than using a common media. A common media can be likened to a party-line telephone system – if someone else was already on the line then you waited until they were done, and once you got the line you continued until you were done. Hence, the total system bandwidth is limited to the bandwidth of the common media. Examples of LANs using a common media include Ethernet, FDDI, and the IBM token ring.

In contrast, switches can be likened to the user's view of a central office telephone system – multiple non-interfering conversations can occur simultaneously. An advantage of switches is that the system bandwidth increases with the number of switch ports, e.g., an 8 x 8 switch with 100 Mbit/s ports has a total system bandwidth of 800 Mbit/s. Point-to-point links are a natural for fiber optics, and copper point-to-point links are much easier to build than copper multi-drop links. Switches are used with HIPPI, Fiber Channel, and ATM.

Standards are becoming more important. Customers are no longer content to be held captive to a vendor's proprietary, but demand that standards be used to allow the interconnection of devices from many different vendors. Industry standards are a good idea, but we also have to be careful that the number of competing standards do not defeat the goal of having only a few proven interfaces that everyone supports.

The use of fiber optics rather than copper cable provides the opportunity for major improvements in the links. Fiber optics have much higher bandwidths (e.g., terabits/s where copper is megabits/s), and can span much longer distances without repeaters. The fiber EMI and RFI characteristics are also much improved over copper. The result is that the error rates for fiber optic links are orders of magnitude better than for copper links. For example, bit error rates going from 10^{-6} on copper to 10^{-9} or 10^{-12} on fiber. These improved error rates in turn affect the communications protocols, which were designed originally for noisy copper links. Fiber optics also provide smaller connector footprints, especially important when you consider how the electronics packages are shrinking. Fiber optic components are still somewhat expensive, but the prices are coming down as volumes increase.

3: Selection criteria

The main point to be made here is that the interconnection technology selected must work for your application. Just because it is the fastest, or has the most hype, does not necessarily mean that it is the best for you.

Questions that you need to ask about your application include the following. Note that there is no priority order implied in this list, and it is also probably not complete.

- What peak data rate is required?
- What average data rate is required?
- Does the data have any time dependence, e.g., voice?
- Is this a single stream of data, or do several streams need to be multiplexed on a single link?
- What level of error detection is required?
- What interconnection distances are involved?
- Is the environment benign or hostile?
- What interconnection topology is desired, e.g., bus, ring, point-to-point, switched?
- What interface is used with your present equipment, and can it be used for the new application?
- What interface(s) does the new equipment support?
- Is this a one-shot deal, or will the system be long-lived and expanded in the future, e.g., higher speeds, more connections, longer distances, etc.? If so, consider the growth as you develop your plans.
- Are the appropriate interfaces available commercially, or will custom interfaces need to be designed or procured?
- If custom interfaces are necessary, are the necessary silicon chips available to build them?
- Is security an issue, and if so how does this affect the interconnection system used?
- What is the time frame, i.e., does it need to be running in a month, or a year, or three years? Hence, can you wait for new products or improvements, or must you only count on what is available today?

The answers to these questions may not point to a specific interface, but will give you insight into the problem and allow you to talk to vendors and others about possible solutions.

4: Interconnect cultures

Often a simple name on an interface does not give a very good idea of what application area the interface was intended for. One way is to look at the background of the people who designed the interface -- what sorts of interfaces have they developed in the past. The interface world can be divided into five separate cultures. These are listed below followed by examples.

- Backplane
  VME, Futurebus +, SCI
- Peripheral I/O Channel
  SCSI, IBM Block Mux, IPI, Fibre Channel
- Local Area Network (LAN)
  Ethernet, FDDI, HIPPI
- Metropolitan Area Network (MAN)
  FDDI, IEEE 802.6 (DQDB)
- Wide Area Network (WAN)
  Telecommunications, X.25, Internet, SONET, ATM

The cultures can be differentiated by the problems they are trying to solve and the goals used to approach the problem. They can also be categorized by "Control" and "Trust". For example, in a backplane environment, a single user has complete control, and gives complete trust since usually the card being installed can read or write anywhere in the system. Compare this to the wide area telecommunications network -- here a single user has very little control (does your cross-country phone conversation go through St. Louis, Chicago, or Dallas?), and shows the lack of trust by using protocols with lots of checking and firewalls.

Another example is charging and data reliability. On a backplane, Peripheral I/O channel, or LAN, once you buy the equipment the vendor does not care if you use it at 1% or 90% of the total bandwidth. But in these environments the data must be delivered correctly, or if an error occurs, this should be a very rare exception and should be flagged. In contrast, the end user does not directly buy telecommunications network central office equipment, but instead rents bandwidth. Hence, the goal for the telecommunications providers is to keep the channels as full of data as possible since this is where they get their revenue. To achieve the full bandwidth, the providers may oversubscribe the links, depending upon statistics over a large user population to avoid dropping data. Since the telecommunications WANs were developed primarily for voice traffic, dropping small amounts of voice data is not catastrophic; but dropping packet data is catastrophic.

If an interface is developed for a particular "culture", it is not unrealistic to use the interface for a different culture, but it may not be optimized for that culture. It is highly unlikely that an interface developed for a backplane would be appropriate for a WAN, and vise versa. If someone touts a particular interface as the "best for all applications", you should be very cautious – it is probably wishful thinking.

5: HIPPI

The High-Performance Parallel Interface (HIPPI) was developed in American National Standards Institute (ANSI) Task Group X3T9.3. The work started in 1987, and today HIPPI is the interface of choice for high-end applications. [1] The HIPPI physical layer specifies an 800 or 1600 Mbit/s (100 or 200 Mbyte/s) simplex interface for distances of up to 25 meters using copper twisted-pair cables. Other HIPPI documents define the packet format, mappings for upper-layer protocols, control of physical layer switches, and a fiber-optic extender. [2]

The documents defining HIPPI include (1) the physical layer (HIPPI-PH, ANSI X3.183-1991), (2) a framing protocol that defines the packet format (HIPPI-FP, ANSI X3.210-1992), (3) a mapping to IEEE 802.2 to support communications protocols such as TCP/IP over HIPPI (HIPPI-LE, ANSI X3.218-1993), (4) a definition for physical layer crossbar switches (HIPPI-SC, ANSI X3.222-1993), and a fiber-optic extender for distances up to 10 kilometer (Serial-HIPPI, not an ANSI standard). The ANSI documents are also being processed as international standards. [3]

Additional documents being developed in ANSI X3T9.3 include (1) mappings to the IPI-3 generic command sets for magnetic disks and tapes, (2) a mapping allowing upper-layer HIPPI protocols to use a Fibre Channel physical interface, and (3) a mapping between HIPPI and ATM.

HIPPI came from the high-speed local network culture. Attempts to include features to better support direct disk and tape I/O were expressly omitted. HIPPI is intended as a memory-to-memory interface, and is used in this mode with high-end stripped disk and tape systems.

5.1: HIPPI advantages

HIPPI is the current interface of choice, largely because it was the first standard at close to the gigabit speed. It came to fruition quickly because of a "keep it simple" goal, and a well-focused direction in the standards committee that avoided adding lots of bells and whistles. Some of the advantages of HIPPI include:

- It is simple, elegant, and easy to understand.
- It has a good physical level flow control. The flow control even works with very long links by the addition of extra buffering at the receivers (approximately 1 kilobyte per kilometer of distance).
- A good tester was developed early on which allowed vendors to test implementations in-house so that interconnection with other vendor's equipment was usually a plug-and-play.
- A variety of products with HIPPI interfaces from a fair number of vendors currently exist. Many are second generation designs, incorporating improvements from earlier designs.
- HIPPI crossbar switches are available from multiple vendors.
- HIPPI specific integrated circuits are available. Even so, some vendors find that small scale integration parts are more suitable due to the simplicity of the physical interface and limitations of the HIPPI specific ICs.
- HIPPI to SONET adapters are available for very long distance links using telephone network facilities.

5.2: HIPPI shortcomings

HIPPI is not without limitations and shortcomings. Perceived shortcomings include:

- It is not a mass-market item, the number of applications that require the bandwidths are not that numerous. Hence the price is higher. It is questionable whether competing gigabit/s technologies, e.g., Fibre Channel or ATM would be any cheaper.
- It does not support speeds slower than 800 Mbit/s. Slower speeds would help make it more of a mass market item.
- It does not support multiplexing. If you transfer a megabyte over a HIPPI channel as a single entity then it will take at least 10 milliseconds. During this time the channel cannot be used for any other communications.
- HIPPI does not support time-dependent or isochronous data.
- The HIPPI specification limits the distance to 25 meters (82 feet) with copper twisted-pair cable. Serial-HIPPI defines a fiber-optic extender than is useful for distances up to 10 kilometers, but it is an added expense.
- The cable connector is large and somewhat fragile.

6: Fibre Channel

Fibre Channel is an emerging standard, also from the ANSI X3T9.3 Task Group. [4] The Fibre Channel work started in 1988, one year after HIPPI started. The first Fibre Channel documents are just now being completed, and the first products being delivered. Fibre Channel supports burst data rates of 100, 200, 400, and 800 Mbit/s. As the name implies, it is based on serial transmission over fiber optics, whereas HIPPI was based on parallel transmission over copper wires. The first products are being developed at the 200 Mbit/s speed, higher speed products will follow shortly.

You may see Fibre Channel referred to with different rates, for example, 133, 266, 531 or 1062.5 Mbit/s. These rates are the serial stream signaling rates that include the 8B/10B encoding and other overhead. The corresponding rates for the user data portion of the serial stream are 100, 200, 400 and 800 Mbit/s respectively.

The Fibre Channel Physical and Signaling Interface (FC-PH) document has completed development and is in the review process. The FC-PH draft document is also available from Global Engineering as ANSI X3.230-199x. Other Fibre Channel documents under development include:

- FC-EP, enhanced physical layer with support for isochronous, stripped physical layers (e.g., running three FC-PH physical layers in parallel for three times the bandwidth on a single transfer), and other things left out of FC-PH.
- FC-IG, implementation guide with state diagrams for FC-PH and a collection of folklore and helpful hints.
- FC-SB, mapping to single-byte command code sets, i.e., IBM Block Mux command sets
- FC-FP, mapping to HIPPI upper-layer protocols
- FC-LE, mapping to IEEE 802.2 for support of communications protocols such as TCP/IP
- SCSI-FCP, mapping for SCSI protocols to use Fibre Channel physical layer for higher speed SCSI devices
- SCSI-GPP, mapping for SCSI generic packetized protocols
- IPI-3 Disk and Tape, revisions to the existing IPI-3 standards to include running over the Fibre Channel physical layer
- FC-FG, fabric generic requirements
- FC-XS, crosspoint switch fabric

- FC-AL, arbitrated loop
- FC-DF, distributed fabric

6.1: Fibre Channel advantages

Fibre Channel came from the "mainframe I/O channel" culture, and it should provide an excellent solution for that application. Fibre Channel's success at penetrating the LAN environment remains to be seen. Fibre Channel is considerably more complex than HIPPI, but it also includes many more features. It will be interesting to see if this extensive set of Fibre Channel features turns out to be a boon or a bane. Advantages of Fibre Channel include:

- Very versatile; can do almost anything.
- Supports multiplexing of 2 kilobyte frames of different information transfers.
- Supports dedicated switched circuits.
- Supports datagram service, i.e., best-effort transfers without acknowledgments.
- Supports broad range of speeds with common integrated circuits.
- Defines a variety of interconnection fabrics, including (1) a crossbar for highest throughput, and highest cost, (2) distributed switching elements for maximum flexibility and ease of growth, or (3) arbitrated loop (a ring architecture with no extra fabric elements required) for lowest cost and lowest performance.
- The switch definitions allow for easy mixing of speeds in a single system. For example, a switch can simultaneously interconnect 200 and 800 Mbit/s end nodes. As you may expect, speed mixing is not supported on the arbitrated loop.
- Fibre Channel may support time-dependent isochronous data, e.g., voice or video, in future releases. Striping across multiple physical channels is also planned as a future enhancement.

6.2: Fibre Channel shortcomings

In trying to be "all-things-for-all-people", Fibre Channel included an extensive set of options, which some people label as "bells and whistles". Only the future will tell if this large option set made Fibre Channel stronger by being useful for a large set of applications. The large set of options could also be a detriment, resulting in vendors having difficulty making interoperable products, i.e., the set of options used by one vendor are not compatible with another vendor's equipment. Fibre Channel may have tried to do too much in one interface. Other perceived shortcomings of Fibre Channel include:

- The development process has taken a long time; Fibre Channel may miss its window of opportunity.
- Integrated circuits supporting Fibre Channel are just becoming available, and they may be made obsolete by later changes in the specification before it is an approved standard.
- Fibre Channel does not currently support time-dependent, i.e., isochronous, data. There are plans within the ANSI committee to add this support, but when it will be developed, and when it will be available in integrated circuits, are open questions.
- All of the options and capabilities resulted in a specification that is quite difficult to read and understand. The complexity will also make interfaces difficult to implement, check out, and verify against other vendor's interfaces.
- Fibre Channel is in a race with ATM, and if Fibre Channel products are not forthcoming soon, then ATM may win by default.

7: ATM

In this context, ATM stands for Asynchronous Transfer Mode, not automatic teller machine. ATM came from the telecommunications community, and defines a protocol for sending information in 53-byte cells. [5,6]

Note that ATM is not a physical level interface. SONET (Synchronous Optical NETwork) is the physical layer interface most often mentioned with ATM. SONET is a point-to-point interface supporting data rates from 51 Mbit/s to gigabits per second. [7] SONET does not support switching by the end users, ATM will provide this function. Hence, SONET is comparable to a leased line, and ATM to a dial-up connection. The SONET speeds on optical interfaces are designated as OC-n, where the serial speed on the link is about 51.8 Mbit/s times n, i.e., OC-3 is about 155 Mbit/s, OC-48 is about 2.4 Gbit/s. As with Fibre Channel, these rates include overhead; the actual user data rates are about 75% to 85% of the signaling rates.

The 53-byte ATM cell was designed for carrying many separate voice traffic connections over a single physical media. The ATM cell is composed of a 5-byte header with routing, control, and checking information, and a 48-byte payload. Adaptation layers, called AAL1 through AAL5, define the nature of the information in the payload. For example, AAL1 is intended for constant bit rate data, e.g., voice or video, while AAL5 is intended for packet data that has no specific timing requirements.

Standards documents for ATM are being developed in ANSI, in ISO, and in the ATM Forum.

7.1: ATM advantages

ATM was designed for wide area telecommunications networks, but there is also a lot of interest in using ATM technology in LANs. Some of the advantages would include:

- ATM already has good support for mixing time-dependent data, e.g., voice and video, with packet data.
- ATM is independent of the underlying physical media, but is most often mentioned in conjunction with SONET. This combination supports a wide range of speeds from megabits per second to gigabits per second.
- It is easy to mix equipment with differing speed interfaces, e.g., OC-3 and OC-12, in the same system.
- By using the same technology as the telecommunications industry, larger volumes of common components should result in lower prices.
- Bridging between LANs using ATM and wide area networks (which also use ATM) should be simpler than converting between two dissimilar standards.
- WANs presently have extensive network management tools and these tools may be available in an ATM based LAN.
- There is a lot of interest and momentum behind the ATM work – if it fails it will not be due to the lack of talented people working on it, or the lack of effort.

7.2: ATM shortcomings

ATM was not specifically designed for LAN usage, and hence has some shortcomings when used in that environment. Perceived shortcomings include:

- Vendors are building and delivering products before the standards and problems have been solved. This is largely a result of "over hyping". If too many troubles or delays occur, then there may be a backlash against ATM.
- The lack of flow control, which may not be a serious problem in wide area network, can cause serious lost data problems in LANs. Compared to wide area networks, LANs have many fewer connections, and the data is much burstier. Hence, in wide area networks, the statistics may play in your favor, but in LANs they do not.
- The loss or corruption of a single 53-byte ATM cell may result in the re transmission of a much larger entity, e.g., 16K byte packet.
- The early ATM equipment for LAN usage supports only Permanent Virtual Circuits (PVCs). This requires bandwidth to be dedicated on the basis of "might be needed sometime". Support for Switched Virtual Circuits (SVCs) is being developed, but is not available in early equipment.
- Setting up an SVC may take a fairly long time, e.g., milliseconds.
- The common speeds supported today with ATM are 100 Mbit/s and 155 Mbit/s. There is some equipment supporting 622 Mbit/s starting to appear, but higher speeds have yet to be tapped. The next higher speed used in the wide area networks will be 2.4 Gbit/s; it remains to be seen if this speed or an intermediate speed is used in LANs. Hence, true gigabit speeds are not available now, and probably not in the near future.
- Splitting large packets of data, e.g., 1M Byte, into many 53-byte cells for transmission seems intuitively wrong. Experience has shown that the fewer times you "touch" the data, i.e., the less the overhead, the faster things run.

8: Summary

Local area network speeds are increasing to keep up with the new generation computing equipment, and gigabit per second speeds are becoming reality. Switches are replacing shared media at the higher speeds, and fiber optics are changing the error characteristics. Standards are becoming more important with few customers willing to invest in proprietary solutions.

HIPPI is the current interface of choice for high-speed LANs, but it is being challenged by ATM and Fiber Channel. There are advantages and shortcomings for each of these interfaces and a potential customer will need to examine his requirements carefully in order to select the most appropriate technology. It is unlikely that a single technology will be the best for all applications.

Acknowledgments

The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

The author's work was performed under auspices of the U.S. Department of Energy.

References

1. D. Tolmie, J. Renwick, "HIPPI: Simplicity Yields Success," *IEEE Network - The Magizine of Computer Communications*, January 1993, pp 28-32.

2. D. Tolmie, "High Performance Parallel Interface (HIPPI)," in *High Performance Networks – Technology and Protocols*, A.. Tantawy, ed., Kluwer Academic Publishers, Norwell, Massachusetts, 1994, pp. 131-156.

3. Approved standards are available from the American National Standards Institute, Attn: Customer Service, 11 West 42nd Street, New York, NY 10036, (212) 642-4900, or from Global Engineering, 15 Inverness Way East, Englewood, CO 80112-5704, (800) 854-7179. The HIPPI documents are also available via anonymous FTP from the hippi directory of ftp.network.com.

4. M. Sachs, "Fibre Channel," in *High Performance Networks – Technology and Protocols*, A.. Tantawy, ed., Kluwer Academic Publishers, Norwell, Massachusetts, 1994, pp. 109-129.

5. R.Handel and M.Huber, *Integrated Broadband Networks – An Introduction to ATM-Based Networks,* Addison-Wesley, Wokingham, England, 1991.

6. M. Zeug, "Broadband Integrated Services Digital Network (B-ISDN) Standards," in *High Performance Networks – Technology and Protocols*, A.. Tantawy, ed., Kluwer Academic Publishers, Norwell, Massachusetts, 1994, pp. 183-204.

7. G. Shenoda, "Synchronous Optical Network - SONET," in *High Performance Networks – Technology and Protocols*, A.. Tantawy, ed., Kluwer Academic Publishers, Norwell, Massachusetts, 1994, pp. 205-228.

DATE
FILMED
4/19/94
END