LAN and I/O Convergence: A Survey of the Issues

Martin W. Sachs, Avraham Leff, and Denise Sevigny
International Business Machines

Local area networks (LANs) and computer I/O are both interconnects that move information from one location to another. Despite this shared purpose, LANs have traditionally connected independent and widely separated computers. In contrast, computer I/O has traditionally connected a host to peripheral devices such as terminals, disks, and tape drives. Because these connection tasks were different, the architectures developed for one task were not suitable for the other. Consequently, the technologies used to implement one architecture could not address the issues faced by the other, and the technologies were seen as fundamentally different.

However, an examination of the architectural requirements of modern I/O and LANs shows that the differences between the two technologies are now disappearing. We believe that LAN and I/O architectures are in fact converging, and that this convergence reflects significant changes in how — and where — computing resources are used. To illustrate this convergence and its implications, this article examines several modern LANs and channels.

Once two distinctly separate technologies, LANs and I/O are becoming more alike through similar distances, media, and purposes. What few differences exist may disappear in the next decade.

Environment and architecture convergence

Today's I/O channels and LANs are characterized by a configuration size of less than 50 kilometers. Within this area, the environments under consideration include

- back-end networks (machine room environment),
- front-end networks (office environment),
- client-server networks, and
- campus backbone networks.

Modern I/O channels do not obviate the need for wide-area networks (WANs) or even large metropolitan-area networks (MANs). The general-purpose I/O channels that we discuss later in this article also do not lessen the need for optimized channels for real-time applications such as embedded systems.

Figure 1 (on page 26) depicts the evolution of the relationship between interconnect type and distance. Historically, I/O and communication network interconnects partitioned the space at the machine room boundary. In the 1980s, the communications space was further subdivided into LANs and WANs, followed by the introduc-

0018-9162/94/$4.00 © 1994 IEEE

Glossary

ATM (asynchronous transfer mode) — Packet-switching technology, using 53-byte cells, developed for telecommunications and supporting voice, video, and data transmission.

Back-out — Refers to various techniques for recovering from errors by restoring the state of a process to a state that existed before the error occurred.

Broadcast network — A network in which data are simultaneously transmitted to all destinations.

Class-1 service — A Fibre Channel service that establishes a circuit-switched connection between two communicating entities (N_Ports).

Class-2 service — A Fibre Channel service that multiplexes frames to and from N_Ports with acknowledgment provided.

Class-3 service — A Fibre Channel service that multiplexes frames to and from N_Ports without acknowledgment.

Connection-oriented — A service in which a connection between source and destination must first be established before communication can take place. Once the connection is established, messages arrive at the destination in the order that they are transmitted.

Connectionless — A communication service in which every message is transmitted independently of any other.

CSMA/CD (carrier sense multiple access with collision detection) — A bus network in which the MAC protocol requires a station to detect whether another station is already transmitting before transmitting its own frame and in which error conditions resulting from simultaneous transmission by more than one station are resolved through retransmission.

Cut-through — A technique used in frame buffering that permits the beginning of a frame to be moved out of the buffer before the whole frame has arrived in the buffer.

Data link layer — The OSI layer that controls data transfer over a link between two nodes and performs error control for the link.

ESCON (Enterprise Systems Connection) — A fiber-optic I/O channel developed by IBM that transmits data at 17 Mbytes/sec. It provides point-to-point connections of up to 40 km and uses a nonblocking circuit switch.

Fabric — The part of a network that transmits data from one node to another, usually including routing function.

FDDI (Fiber Distributed Data Interface) — A high-performance, ANSI-standard fiber-optic token ring LAN running at 10 Mbytes/sec over distances of up to 200 km with up to 1,000 connected stations.

Fibre Channel — A proposed ANSI serial I/O channel standard capable of transmitting at gigabit rates. It provides both circuit and frame switching using space division switches or loops.

HiPPI (High-Performance Parallel Interface) — A high-speed ANSI-standard parallel interface that transmits either 32 or 64 bits in parallel and transmits data at up to 800 Mbits/sec.

Hop count — A unit of distance in a communications network. A hop count of 4 means that 3 nodes or gateways separate the source from the destination.

I/O channel — An I/O mechanism that manages the flow of data between a processor memory and the link to attached I/O devices.

IP (Internet Protocol) — The network layer of the Internet communications protocol. It defines the Internet datagram as the basic information unit passed across the Internet and provides a connectionless delivery service.

IPI (Intelligent Peripheral Interface) — An ANSI-standard I/O interface primarily used for attachment of data storage devices to processors.

LAN (local area network) — A communications system typically designed for use within a single organization, having a diameter greater than 10 m but less than several km.

LLC (logical link control) — One of two sublayers of the OSI model's data link layer. It includes functions unique to the particular link control procedures associated with the attached node and are independent of the underlying communication medium. The LLC sublayer uses services provided by the MAC sublayer and provides services to the network layer.

MAC (medium access control) — A sublayer of the OSI model's data link layer. It uses the services of the physical layer and supports topology-dependent functions, which it provides to the LLC sublayer.

MAN (metropolitan area network) — A communications system designed to cover city-wide areas (tens of km), using LAN technology.

Native I/O — The I/O system designed as an intrinsic component of a given computer architecture.

Node — A communications entity.

OSI (Open Systems Interconnection) *architecture* — A framework for coordinating the development of standards for the interconnection of computer systems. Network functions are divided into a hierarchy of seven layers; each layer represents a collection of related communication functions.

Quality of service — Parameters characterizing communication service that a service user either desires or requires as minimum satisfactory service. Examples include specifications for throughput, delay, and error rates.

SCSI (Small Computer Systems Interface) — An ANSI-standard I/O interface primarily used for attachment of data storage devices to processors.

Shared-medium topology — A communication network in which a single communication channel is shared among all the stations on the network. Examples are bus and ring topologies.

Station — A communication device attached to a network. The station is the component of a node that provides at least the MAC and physical-layer function.

Switch-based topology — A communications network that is based on one or more discrete switches, which may be either circuit or packet switches.

TCP (Transmission Control Protocol) — The transport layer of the Internet communications protocol. It provides reliable, full duplex service, and allows arbitrarily long streams of data to be transmitted. It provides a connection-oriented service and typically uses the IP protocol to transmit data.

WAN (wide area network) — A communications system typically designed to provide services to a geographical area that is larger than the area served by a single LAN.

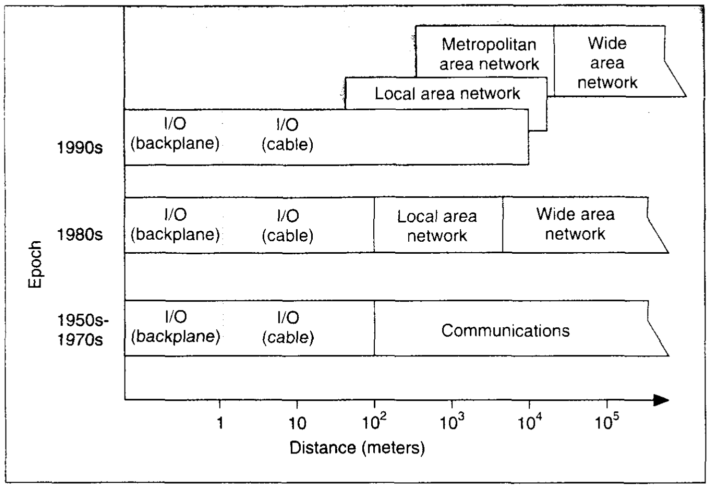

Figure 1. Interconnect type and interconnect distance for various epochs.

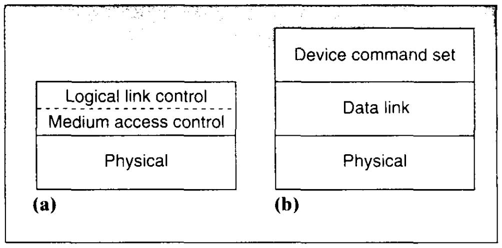

Figure 2. Layer structures of (a) LANs and (b) channels.

Table 1. Interconnect requirements.

Requirement	Definition
Interconnect distance	Maximum distance between any two points of the configuration
Information model	Characteristics of the application information carried by the interconnect
Computation model	Relationship between the interconnected computers or devices (for example, master-slave, peer-to-peer)

tion of MANs in the 1990s. In today's systems, distances served by I/O channels and networks now overlap in local and metropolitan areas.

Although the traditional dichotomy is still valid at opposite ends of the scale, it blurs toward the middle. Some aspects of computer I/O, such as flow control, can be optimized for short distances through simple hardware protocols, but they do not perform well at very long distances. Similarly, communications systems are generally optimized for long distances and multiple hops, and they have too much overhead for high-performance I/O. However, in intermediate configurations for distances of 1 to 50 km, the issues faced by LANs and I/O are increasingly similar.

In the terminology of the OSI (Open Systems Interconnection) reference model,[1] LANs are characterized by the physical layer and by the medium-access-control (MAC) sublayer of the data link layer. The same logical link control (LLC) sublayer of the data link layer is used by more than one type of LAN. I/O channel architectures can be viewed as consisting of three layers. The lowest two are functionally equivalent to the OSI physical and data link layers. The highest layer specifies device command sets (for example, disk and tape commands) together with their associated protocols.

The layer structures of LANs and channels are illustrated in Figure 2. Typically, the software that supports a LAN implements a layered architecture, which may conform either to the higher layers of the OSI model or to some other communications architecture. In general, these higher layers are not specific to LANs and the same software may concurrently support LAN, MAN, and WAN communications. For a channel, the supporting software consists of an I/O supervisor and the device-specific and application-specific software.

Interconnect media and bandwidths are dealt with in the physical layer. The same physical layer specifications can be used for both LANs and I/O channels, independent of functional convergence issues.

Above the data link layer, protocols reflect specific application types. Examples are communications, such as the Internet Protocol, and device architectures, such as the American National Standards Institute Intelligent Peripheral Interface (IPI-3). These applications and their protocols are not converging. The data link layer, however, is the focus of our discussion on architecture convergence.

At this point, it is worth briefly mentioning architecture openness. Traditionally, a given I/O architecture was optimized for a given system architecture for the purpose of attaching I/O devices to a particular processor. In contrast to LAN architectures, I/O architectures were never open in the sense of being able to incorporate the identical I/O architecture with different system architectures. Early exceptions to this were the ANSI Small Computer System Interface (SCSI) and IPI architectures.

The SCSI and IPI standards both include a *device* model, which defines the characteristics of the device (for example, a disk drive) and its command set, and a *channel* model, which describes the architecture of the channel that connects the device to the host. However, most implementations of these standards were made on processors that already had native I/O systems to which SCSI and IPI adapters were attached. The emerging ANSI Fibre Channel[2] standard goes further: It is a true open I/O architecture that is both processor and device independent. Moreover, some implementations in which Fibre Channel is the native I/O system are likely.

Interconnect requirements

Among the requirements that must be addressed by any interconnect architecture are the expected interconnect distance, the information model, and the computation model (see Table 1). As shown in Table 2, both the interconnect distance and the computation model have traditionally differed for I/O and LANs. For short distances, I/O channels efficiently connected a smart host to a few dumb peripherals in a centralized, master-slave manner (Figure 3). Where longer distances were a factor, LAN architectures were needed to connect many autonomous, smart processors in a distributed, peer-to-peer manner (Figure 4). With respect to the information model, both I/O channels and LANs were primarily used for transferring data.

Interconnect requirements are changing. Both LANs and channels have benefited from advances in fiber-optic technology that greatly extend the combination of bandwidth and interconnect distance. At the same time, the information model — driven by multimedia applications using voice and video — is evolving to include information with very different characteristics.

In the evolving I/O computation model, data are increasingly off-loaded from the host to intelligent file servers that form the basis of client/server architectures. The result is that a traditional LAN peer-to-peer interconnect quite naturally supports the current I/O computation model. There are few differences between LAN and I/O architectures developed to serve the demands of multimedia and client/server applications, and even these differences may disappear in the next decade.

Technology

Some of the classic differences between channels and LANs are independent of the interconnect's intrinsic architecture. Instead, they're technological —

Table 2. Traditional requirements for communication and I/O.

Requirement	LANs	I/O
Interconnect distance	> 10^2 meters	< 10^2 meters
Information model	Data	Data
Computation model	Peer-to-peer	Master-slave

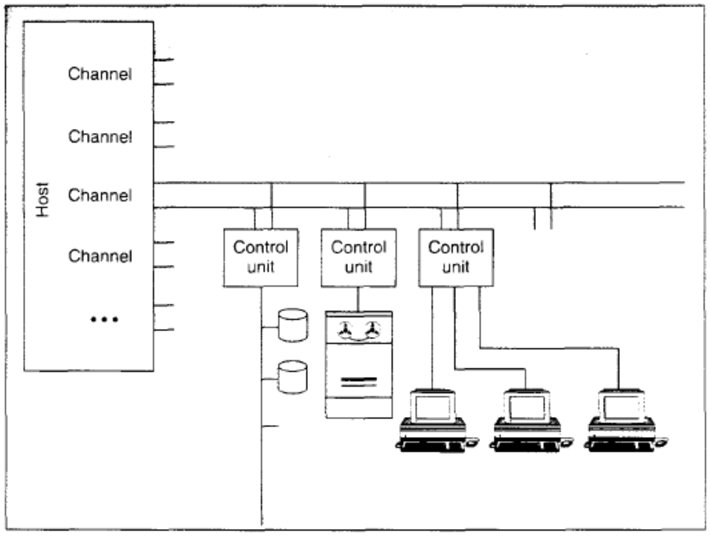

Figure 3. Typical I/O interconnect configuration.

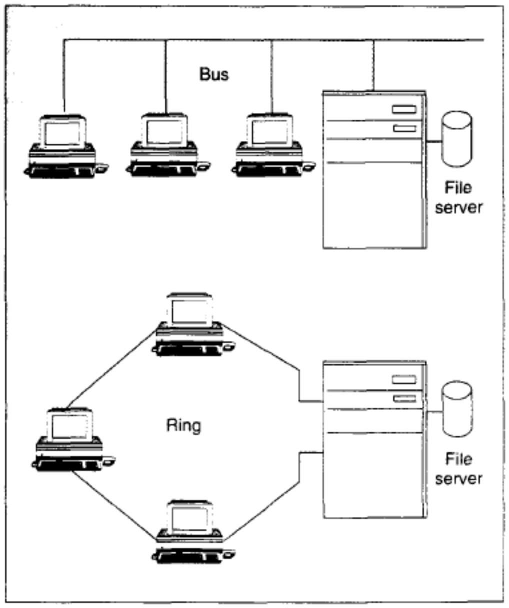

Figure 4. Typical LAN interconnects—bus and ring. Several workstations are shown interconnected to each other and to a file server.

Table 3. Two major technological differences between traditional LANs and channels (late 1980s).

Attribute	LANs	Channels
Interconnect distance	> 10^2 meters	< 10^2 meters
Bandwidth	1 Mbyte/sec	10 Mbytes/sec

for example, the use of parallel versus serial links or copper cable versus optical fiber. These design decisions depend on product-specific cost, performance, and distance trade-offs. In theory, both interconnects could use the same technology.

Table 3 compares two technologically determined attributes — interconnect distance and bandwidth — that distinguished channels from LANs in the late 1980s. Compared to channels, LANs have had lower bandwidth and a longer interconnection distance.

Table 4 compares channel and LAN bandwidth over three generations of technology. Through the mid-1980s, both used the same basic copper-wire transmission medium. Channels were designed to move large volumes of data rapidly between a host and its attached high-performance data-storage devices, while LANs were designed as low-cost, higher bandwidth alternatives to the existing communication networks over shorter distances. Therefore, channels were optimized for high speed and very low error rates, using parallel transmission for relatively short distances. LANs were optimized for economy and relatively long interconnect distances, and they featured serial transmission and tolerated higher error rates. LAN specifications also permitted repeaters to extend distance; indeed, repeating is an essential part of a ring, one of the common LAN topologies.

Recent developments in serial fiber optics permit multikilometer attachment of peripheral devices to channels. For example, the IBM Enterprise Systems Connection (ESCON)[3] channel has 3-km point-to-point connections and a laser option that enables up to 20-km point-to-point connections. Fibre Channel specifies a gigabit/sec option with a 10-km point-to-point connection. Moreover, the actual maximum distance for both the ESCON channel and Fibre Channel can be greatly extended through repeaters and complex switching fabrics. At the same time, emerging LANs such as the ANSI Fiber Distributed Data Interface (FDDI) have applied fiber-optic technology to greatly extend the combination of bandwidth and interconnect distance available with earlier LANs, such as Ethernet and token-ring topologies. Thus, interconnect distance and bandwidth no longer distinguish channels from LANs.

Architectural aspects

To compare I/O and LAN architectures, we identify and discuss nine key features: topology, transmission latency, single-hop versus multihop configurations, connection-oriented versus connectionless service, real-time constraints, fair access, priorities, multiplexing, and bandwidth management.

Topology. Channel topology originally reflected the host/peripheral relationship very strongly. A host communicated in master-slave fashion to its attached devices. Devices could share the medium because the master-slave relationship allowed them to use a very simple access-control protocol to mediate contention. Device sharing among multiple hosts was provided by the device's controller. Each controller had multiple ports for channel connections and a switch that resolved contention among hosts. Newer channels, such as the ESCON channel, Fibre Channel, and ANSI High-Performance Parallel Interface (HiPPI),[4] have moved to a switch-based topology that is intrinsically peer to peer. Hosts set up connections through the switch to share devices.

In contrast to the master-slave topology of early channels, LANs evolved in a distributed, peer-to-peer environment. Shared-medium topologies (for example, rings and buses) were attractive because they provided sufficient bandwidth at low cost. However, the combination of peer-to-peer relationship and shared medium required more complex access protocols, such as carrier sense multiple access with collision detection (CSMA/CD) and token-passing. Moreover, in a shared-medium topology, the bandwidth of a single link is divided among all the stations.

Future applications will require LANs with more bandwidth per station than is now feasible through shared-medium topologies. Switch-based LANs — such as asynchronous-transfer-mode (ATM)[5] LANs and switched Ethernet — are attractive because they provide aggregate bandwidth, which is a multiple of the bandwidth of a single link. ATM technology is even being applied as an I/O interconnect.[6] In the extreme, a nonblocking circuit switch, similar to the IBM ESCON Director[3] or a Fibre Channel circuit switch, provides the full bandwidth of a single link concurrently to every station engaged in a connection. Channels and LANs are clearly evolving from their original shared-medium topologies to switch-based topologies.

Transmission latency. The intrinsic latency of all channels and LANs is similar for a given transmission medium, being determined by the signal speed in the transmission medium. For optical fiber, this is approximately 2×10^8 meters/sec, corresponding to a propagation delay of 5 nanoseconds/meter, that is, 50 microseconds at 10 kilometers. However, whereas channel latency is measured in microseconds, LAN latency is measured in milliseconds. This difference relates to aspects of architecture and implementation above the OSI physical layer.

Interconnect architecture affects latency because of the access protocols needed to mediate simultaneous transmission requests. For example, the medium-access-control protocol used in CSMA/CD and token-ring LANs causes considerable transmission latency, compared to the simpler switch-based access protocols of channels. In addition, the end-to-end latency of both channels and LANs may, in fact, be dominated by software that implements the applications and higher layers of the communications protocol stack. In the future, as both channels and LANs adopt similar switch-based topologies, the latencies in the MAC-layer function will become similar. Further, improvements in architecture will reduce the effect of software on latency.

Single-hop versus multihop. An interconnect must ensure that data are transmitted from the source to the destination.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.