SHIPBOARD NETWORK ANALYSIS

Alfred C. Weaver, Ph.D.
University of Virginia

Computer Science Report No. TR-86-11
April 29, 1986

SHIPBOARD NETWORK ANALYSIS

Alfred C. Weaver
Associate Professor of Computer Science
Department of Computer Science
University of Virginia

April 28, 1986

86-11

1. EXECUTIVE SUMMARY

In section 3 we first question whether the type of data transfers envisioned are best served by a local area network or a device interconnect; the conclusion is that a local area network is fully justified. We compare the primary LAN topologies and determine that the star configuration is risky due to its centralized switch; the baseband bus is better but subject to shorting; the broadband bus is better still but has a critical element in its head-end remodulator; the ring has inadequate reliability by itself, but can be made reliable by using (1) a modified star-ring architecture in which all devices are connected to wiring centers and only wiring centers are connected to the network backbone, and (2) enhancing the backbone reliability by using dual-cable media or even dual-redundant contra-rotating rings.

We examine media types and conclude that unshielded twisted pair is too error sensitive to consider; shielded twisted pair is adequate for short runs; baseband coax is limited to 500 meters; broadband has no effective limitation; and optical fiber supports runs of 2.5 km without repeaters. The noise immunity of fiber optics is far superior to the copper media and is strongly recommended for high noise areas.

In the remainder of section 3 we show that all modern LANs use a form of packet switching to exchange data; we examine the international standards and the standards-making bodies which are of concern; and describe how modern LANs can be interconnected with other LAN segments of the same type, or with LANs of different type, or with wide-area packet-switched networks by using bridges and gateways.

In section 3.7 we identify six factors which are crucial in deciding what type of LAN is best suited for shipboard use. The driving factors are: (1) system reliability; (2) baseband vs. broadband transmission; (3) copper vs. fiber optic cabling; (4) adherence to international standards; (5) cost; and (6) network performance. These factors are used as decision criteria in section 6.

In section 4 we review seven general types of networks: conventional star architectures; avionics busses such as MIL-STD-1553B; the proposed integrated systems digital network (ISDN) for the telephone system; the IEEE 802.3 contention bus; the IEEE 802.4 token bus; General Motors Manufacturing Automation Protocol (GM MAP) which uses the token bus; and the IEEE 802.5 token ring. Star architectures were discarded as being unreliable; avionics busses are not suitable for use as general purpose LANs; ISDN concerns can be handled by reserving capacity for 64 kbps digital channels; the contention bus is usable but does not appear to offer good value; the token bus is implemented in broadband technology and offers specific advantages; GM MAP does not offer any obvious advantages to offset its high price; and the token ring implemented with a star-ring architecture proves to be generally suited to our needs.

Section 5 reviews eighteen specific products of twelve network vendors. First are those which are contention-based: Sytek LocalNet, Bridge Communications Ethernet, Intel Ethernet, Interactive Systems LAN/II (Ethernet version), and Ungermann-Bass Net/One. Second are those which use token passing on a bus topology: Datapoint Arcnet, Interactive Systems LAN/I and LAN/PC, Concord Data Systems Token/Net, and Industrial Networking MAP/One. Third are those which use token passing on a ring topology: IBM token ring and Proteon ProNET. Fourth are those which use centralized control and master/slave polling: Intel Bitbus and Sperry ASCB.

In section 6 we divide the candidates into two groups: "not recommended" and "remaining candidates." For those systems not recommended, reasons are stated. Four candidates were chosen for further study and cost analysis: Concord Data Systems Token/Net; Ungermann-Bass Net/One (broadband version); IBM token ring; and Proteon ProNET.

Each candidate was reviewed for advantages and disadvantages. Each system was cost analyzed for a base configuration to support 5 PCs and 10 other I/O devices. Each network was ranked against the evaluation criteria developed previously in section 3.7. As an overall judgement, the ranking of these four products was:

Best: Proteon ProNET
Second: IBM token ring
Third: Concord Data Systems Token/Net
Fourth: Ungermann-Bass Net/One (broadband)

Based upon all the above information and analysis, the recommended network for the shipboard LAN is a token ring, using multiple wiring centers, and shielded twisted pair wiring (or optical fiber for long runs and/or high noise areas).

2. SCOPE OF WORK

2. SCOPE OF WORK

The following activities were undertaken:

(1) List the options available to a LAN designer in terms of network topology, media, service type, standards, and interconnects. Isolate the most important considerations (driving factors) for use as criteria in an eventual comparison of products.

(2) Evaluate seven general types of networks: conventional stars, avionics busses, integrated services digital network, contention bus, token bus, GM MAP, and token ring.

(3) Evaluate the LAN product lines of the following manufacturers: Sytek, Bridge Communications, Intel, Interactive Systems/3M, Concord Data Systems, Ungermann-Bass, Industrial Networking, IBM, Proteon, Datapoint, and Sperry.

(4) Recommend one or more systems as being best overall for a shipboard network. Evaluate each candidate with respect to the considerations developed in item (1) above.

3. GENERAL NETWORKING CONSIDERATIONS

3.1 When to use a local area network

3.2 Network topology

3.3 Network media

3.4 Network service types

3.5 Network standards

3.6 LAN interconnects

3.7 Selecting the driving factors

3.1 WHEN TO USE A LOCAL AREA NETWORK

A Local Area Network (LAN) is a communications vehicle which provides data connectivity among geographically adjacent computers, peripherals, and control devices. LANs are distinct from long-haul networks (also called wide-area networks) in several very specific ways:

(1) LANs are privately owned, generally by a single organization; they are user-administrated, which gives them flexibility; and they are not subject to regulation by the FCC or similar government agencies.

(2) LANs are structured; that is, the services they provide are integrated into a single hardware/software package which operates harmoniously in support of specific applications (e.g., word processing, electronic mail, file transfer, data distribution, real-time control systems) using a single cable plant.

(3) LANs are limited in geographic scope. A particular installation will generally be limited to servicing devices within a certain distance. This is often referred to as the worst-case end-to-end length limitation (i.e., the largest distance separating any pair of stations) and varies from 500 meters to about 25 miles. The limitation is a function of the media used and the signaling technique employed.

(4) LANs support full connectivity; that is, every device is both physically and logically connected to the same cable plant, thus providing a data communications path between any pair of devices. Data transmission is typically accomplished by *broadcast* where every device hears every transmission, but only the device whose address matches the destination address of the message bothers to copy the message into an internal buffer.

(5) LANs support high speed data communication, on the order of 1 to 10 million bits per second (Mbps). It is generally agreed that transmission rates of less that 1 Mbps disqualify an interconnection scheme from being a true local area network.

(6) LANs use serial transmission techniques; they are message (packet) oriented; they give the illusion of full-duplex transmission by implementing very fast simplex transmission between any pair of stations.

As a result, LANs are powerful and much more expensive than older, character-based, point-to-point communications schemes such as RS-232-C. Thus it behooves the designer to assure that his communication application is in need of LAN technology before specifying same. As an example, if the need is to allow a computer terminal access to any one of, say, three computers, that application is better served by a centralized PBX-style data switch (e.g. Rolm, Gandolf, Micom) than by a LAN. Since an ordinary terminal would use only one connectivity path at a time, would not support LAN-type data rates, and would probably use character-oriented (i.e. byte size) transmissions, the cost of a LAN is simply not justified.

A comparison of major characteristics for PBX systems and for LANs is shown in Table 1.

CHARACTERISTIC	PBX	LAN
Topology	star	star or bus or ring
Media	twisted pair	twisted pair, coax, fiber optic
Network control strategy	centralized	distributed
Data service	circuit switching	virtual circuits or datagrams
Transmission speed	typically 9600 bps	typically 1-10 Mbps
Wiring routes	fixed	flexible
Reliability	redundant switch	redundant cabling
Cost	high initial cost for switch; low cost for each station connected	moderate cost for each station connected

Table 1.
Comparison of PBXs and LANs

In summary, we find that the shipboard network requirements document fully justifies the generality of using a LAN to implement the desired connectivity. The hard part, of course, is choosing the right LAN.

3.2 NETWORK TOPOLOGY

LANs are generally constructed using three main topologies and some variations. The main topologies are star, bus, and ring. The main types of bus are linear bus, branching tree bus, baseband bus, and broadband bus; the main variations of the ring are star-shaped rings and several varieties of redundant rings.

3.2.1 Star topology

The star topology provides a single, centralized device, usually a switch, to which all network stations are attached. See Figure 1.

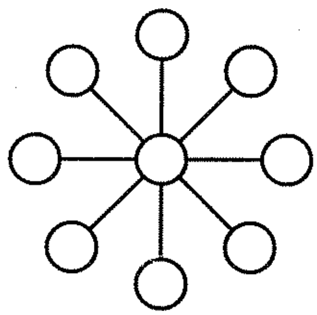

Figure 1.
Star Topology

The advantage of the star topology is its simplicity: every network device has exactly one connection, the one connecting it to the central switch. Thus every conversation is accomplished by simple two-hop routing: from the source to the central station, and from there to the destination. If the central station is indeed a switch, and if the speed of the conversations is such that they can be accommodated by round-robin service of the central node, then the star topology is a realistic choice. For example, all telephone systems use a star topology because the speed of the conversation (8 bits of digitized voice every 125 microseconds) is easily handled by a round-robin hardware switching system.

The disadvantage of the star topology, however, is a serious one. Since there is only one central switch, if it fails then the network fails. For any serious network (like a telephone system), the central switch must be dual-redundant or even triple-redundant to assure reliability. Since system control is centralized, redundancy is the only technique available to enhance overall system reliability.

For modern LANs, the star topology is outdated for at least three reasons:

(1) Data communication, unlike voice communication, tends to be bursty rather than constant; thus the (constant) round-robin service of the central switch is ill-suited to the actual communication rate. If the central switch is indeed fast enough to handle a burst of data from a station, then as a consequence there must be considerable wasted bandwidth when the station is silent.

(2) The topology demands that each station be wired to a single, centralized site.

This may not be possible, or it may not be convenient, or it may simply involve excess wiring costs when compared to one of the alternatives.

(3) Any time one introduces a single anything (in this case the central switch), you have a potential single point of failure. Redundancy of the critical element is one way to increase reliability, but it is not the only way, and it is not necessarily the least expensive way.

For these three reasons, star topologies are unpopular as the primary topology for a LAN. However, some newer *hybrid* topologies, such as the star-shaped ring (discussed later in section 3.2.3.1), do make legitimate use of the star concept.

3.2.2 Bus topology

Another simple topology is the *bus*. In its pure form, the linear bus looks like the diagram in Figure 2.

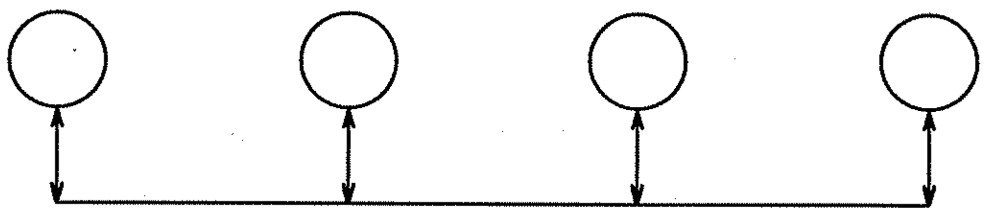

Figure 2.
Linear Bus

An alternative strategy is shown in Figure 3. This is a branching tree bus in which linear busses are themselves interconnected by a linear bus.

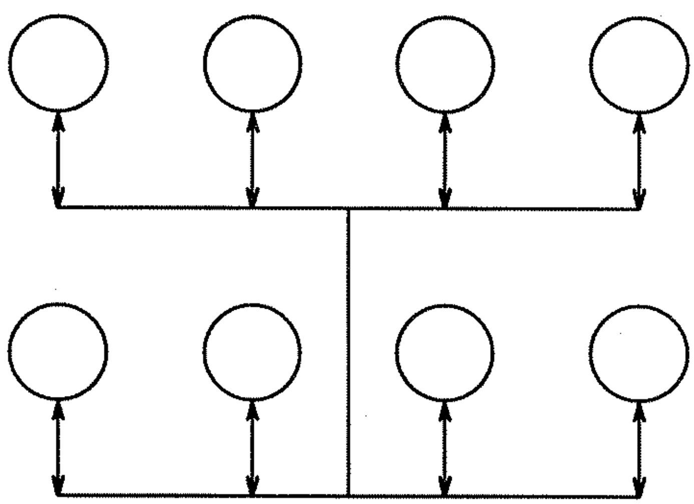

Figure 3.
Branching Tree Bus

Using the bus topology, the cable forms a single continuous path which interconnects all stations. A working definition of the "bus" is that it provides a *single, unique* data path between each and every pair of stations.

Depending upon the medium in use, tapping the cable may be simple or complex. If the medium is twisted pair, a tap is a simple splice which can be accomplished by virtually anyone; if the medium is fiber optic cable, tapping is complex and expensive, requiring special tools, techniques, and a qualified technician.

Figures 2 and 3 depict possible topologies for *baseband busses,* i.e., those using digital signaling, where the presence or absence of a specified voltage is interpreted as a "1" or a "0". These topologies work because digital signaling is inherently bidirectional, so energy (signaling) injected into the bus by a transmitter propagates out of the tap in both directions and eventually reaches all receivers.

The other class of busses is *broadband busses.* Signaling on a broadband bus is always analog, using the amplitude, frequency, or phase of the signal, or some combination of these three characteristics, to encode zeroes and ones. Due to the nature of the coupling and amplifying hardware necessary to carry the RF signal, broadband signals are effectively unidirectional. To enable each station to both transmit and receive unidirectional signals, one of two techniques is normally employed.

One technique is called *band-split* and utilizes different "forward" and "reverse" frequencies on a single cable. A transmitter transmits at a fixed frequency to a frequency translator at the "head-end" of the cable; this is the "reverse" direction. There the translator demodulates the signal, recovers the digital data, remodulates the data on a higher frequency, and retransmits the data in the "forward" direction (all of this on the same cable). By demodulating the signal and then remodulating it, some noise is eliminated in the process.

For single-cable systems, the place in the frequency spectrum which separates the "reverse" channel from its associated "forward" channel defines the type of band-split: *mid-split* or *high-split.* Depending upon the mix of data traffic, voice, and video channels being carried, one type of split may be preferable over another. A *single cable broadband bus,* also called a *head-end bus* (because of the frequency translator at one end) is shown in Figure 4.

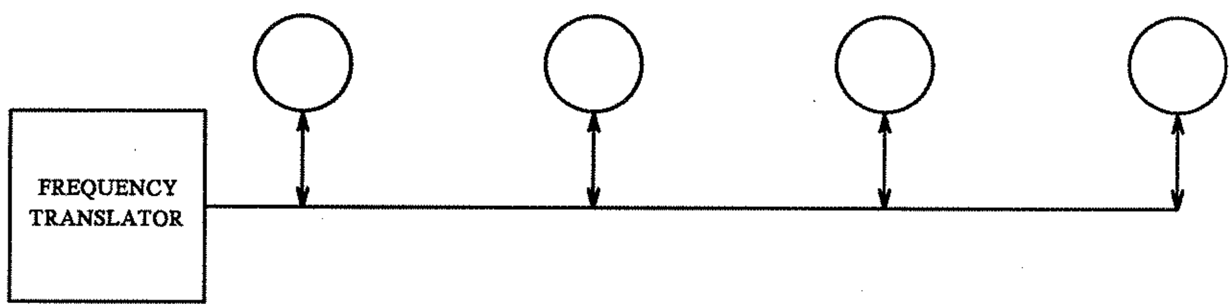

Figure 4.
Single Cable Broadband Bus

The other technique is to use two cables rather than one. One cable connects all of the transmitters to the head-end; the other connects the head-end to all of the receivers.

The advantage is that neither cable is band-split; the transmitters have the full frequency range of their cable while the receivers have the full frequency range of their cable. Thus the total network bandwidth is double that of a single-cable system. The disadvantage is, of course, that dual cable systems require twice as much cable as single-cable systems. See Figure 5.

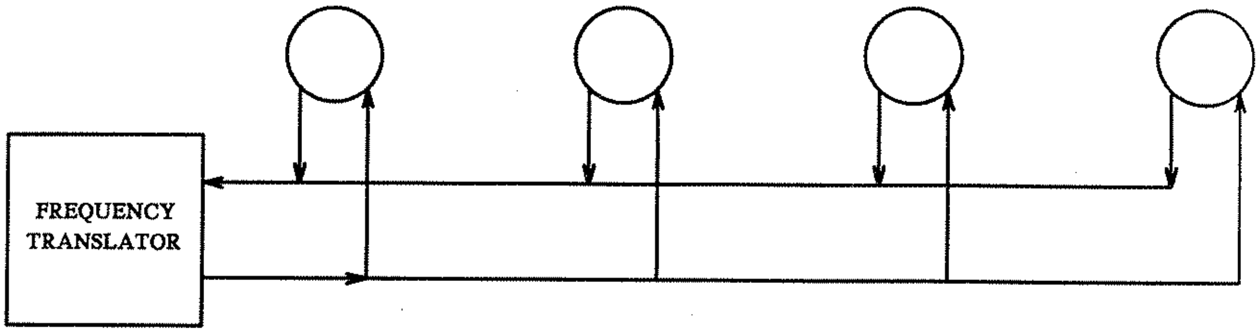

Figure 5.
Dual Cable Broadband Bus

In practice, broadband LANs carrying primarily data (as opposed to real-time voice and/or video) have enough capacity such that dual-cable systems are simply not justified.

Broadband networks use *modulator-demodulators*, or *modems*, to impose the digital signal on the RF carrier. Modems are available in two types: "fixed-frequency", which transmit on one specific frequency and receive on a different specific frequency, and "frequency-agile", which transmit and receive on any of several pairs of frequencies. Most frequency-agile modems select a pair of frequencies to use under software control.

Bus networks are sensitive to reliability issues. A single-cable bus can be rendered useless by a broken connection or a shorted link. There have been no major advances in bus reliability other than redundancy. Using a dual-redundant bus, software can switch from the primary to the secondary whenever test messages fail on the primary. For broadband busses, the "head-end" frequency translator can be made redundant.

3.2.3 Ring topology

The ring topology is the newest, although it had its beginnings in the 1960s as the Cambridge Ring in England. As its name implies, the ring is a closed linear bus as shown in Figure 6.

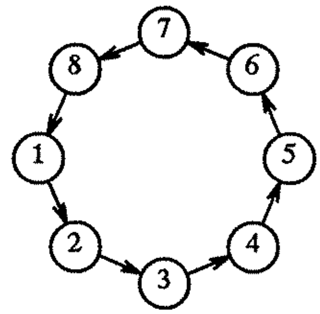

Figure 6.
Ring Topology

The wiring of the ring topology can be seen as either an advantage or a disadvantage. If the stations to be connected are in the same geographic area, then wiring the ring is as simple as wiring a bus. If one or more stations are quite distant from the others, then ring wiring could be twice as lengthy as the equivalent bus.

The primary advantage of the ring is that it lends itself to digital transmission, in which each "hop" from station to station is actually an independent point-to-point transmission cable. At least in principle, each point-to-point connection around the ring could utilize a different medium (twisted pair, coax, fiber optics, infrared, microwave, etc.). Also, because of the limited signal path, signal attenuation is rarely the problem on a ring that it is on a bus. This in turn reduces complexity and reduces the need for automatic gain control receivers, which in theory should lower interface costs.

A ring topology can utilize any medium, but is especially suited for fiber optics. The combination of point-to-point transmission (i.e. no taps), synchronous data signaling techniques, and high data transfer rates (10-100 Mbps), is ideally suited for fiber optics. As an aside, NASA is planning a fiber optic ring as the backbone LAN for the NASA space station; it will operate at 125 Mbps.

Each station on the ring has a receiver and a transmitter. In normal operation the transmitter simply repeats whatever is heard by the receiver. When the station wishes to transmit, it substitutes its own data in place of the received data. When the station wishes to receive, it merely copies the data that flows through its receiver. Details of network access are covered more fully in the section on the IEEE 802.5 token ring.

3.2.3.1 Wiring Centers

The primary concern of ring topologies is reliability. If each station is implemented with a simple transmitter/receiver pair organized as a repeater, then clearly a failure of any single station would cause the ring to fail. This would make the ring's reliability *n* times worse than the star, since there would be *n* potential points of failure rather than one. Obviously, such a simplistic design is not viable.

What we require is a robust way of assuring that any station on the ring which fails will be removed from the ring, rather than allowing that failure to disrupt the ring. The module which implements this assurance is called a "bypass". If a station fails, the bypass will electrically connect the station repeater's input to the repeater's output, thereby "bridging" around the failed repeater (station) and providing a direct electrical connection from the failed station's predecessor to the failed station's successor (i.e., its upstream and downstream neighbors).

If the station wiring is copper cable, the bypass circuit is simplified; it consists of a normally-closed relay connecting the station's input to its output. Power to energize the relay and break the "short-circuit" is supplied by the connecting station. If station power fails, then power to the bypass relay fails, the relay returns to its normally-closed position, the station's input becomes directly connected to its output, and the failed station is effectively removed from the ring. An example of this strategy is the Proteon ProNET network. Each port on the wiring center is "protected" by a bypass relay. Whenever the port is unpowered (because nothing is plugged in or because power to the station which is plugged in has failed), the bypass circuit engages to assure ring connectivity. When the port is connected to a powered and properly operating station, the relay is energized which breaks the short-circuit and allows the station's repeater to be active.

For fiber optic or infrared links the scheme is somewhat different. A fiber optic connection supplies no electrical power to its port, therefore the bypass circuit must be powered from another source (sometimes from the ring itself). In this case the same normally-closed relay is used as before such that its unpowered state connects the station's input to its output, thereby removing the station from the ring. To energize the relay, the station must send a particular message (a *join ring* command) to the bypass relay's control circuits. The same message must also be reissued periodically while the station remains active. If *join ring* messages cease from a station for some predetermined period of time, then power to the bypass relay is lost, the relay contacts close, and the station is removed from the ring.

The use of a wiring center also solves two other problems. First, it elegantly solves the problem of dynamic topology. Adding a node to the network means connecting a new host to any existing wiring center, but not uncabling and recabling the ring connections themselves. Making the ring connections to the wiring centers semi-permanent is a reliability advantage, especially when using vibration-sensitive connectors such as fiber optics.

Second, the ring medium is independent of the host-to-wiring-center medium. While fiber optic is an excellent choice for the ring backbone, it is an expensive choice for interfacing a 4800 bps host connection. The wiring center concept permits simple, slow-speed, inexpensive, serial device connections using twisted pair while preserving the reliability and speed advantages of a fiber optic backbone ring.

As one would expect, these advantages are not without cost. Introducing the wiring center concept increases the per-device interface cost.

3.2.3.2 Redundant rings

More complicated is designing a ring which can recover from cable failures. Examine Figure 7 where the connection between station 7 and station 8 has been severed (dashed line).

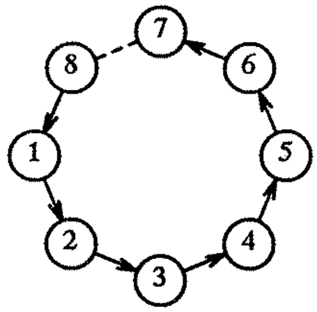

Figure 7.
Failure of Connection from Station 7 to 8

How can connectivity be restored? Clearly there is no way to accomplish this in a single-cable, unidirectional ring. One option is pure redundancy: add a second (or third, or fourth) ring (and all their interfaces) to each station, both rings transmitting in the same direction (either clockwise or counter-clockwise). When transmission fails on ring one, switch to ring two; when transmission fails on ring two, switch to ring three, and so on. Pure redundancy, as shown in the dual-redundant ring in Figure 8, *could* increase overall reliability *if* the whole system were very carefully designed.

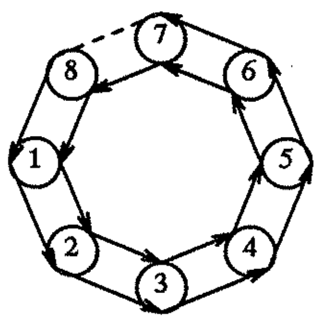

Figure 8.
Dual-redundant Unidirectional Rings

A dual-redundant ring, with both rings transmitting in the same direction, can survive any one cut, but not two cuts between the same pair of nodes. Common sense would dictate routing the dual path wiring separately, but they still share common endpoints, so there is some non-zero probability of a double cut between any pair of stations. Interestingly, a simple change of strategy will heal this fault: have the two rings transmit in opposite directions. Examine Figure 9.

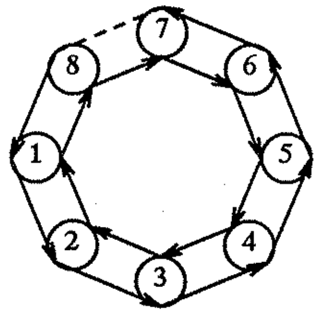

Figure 9.
Contra-rotating Dual-redundant Rings

In Figure 9, communication proceeds on the primary ring (say the outer one) until a break is detected. The station detecting the break sends a message on the secondary ring (here the inner one) directing all stations to switch to the secondary ring. So far the contra-rotating rings provide no advantage over the unidirectional dual-redundant rings. But suppose we have a double cut, as in Figure 10. Now the advantage is clear: with contra-rotating dual-redundant rings, we have restored full ring connectivity.

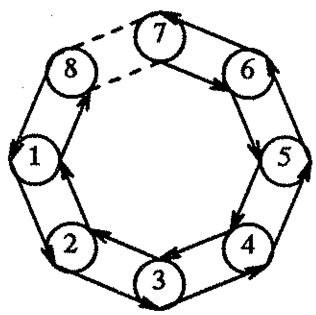

Figure 10.
Contra-rotating Dual-redundant Rings with a Double Cut

The final case occurs when we suffer three or more cuts. Figure 11 shows the situation with two double-cuts. Here, of course, it is not possible to restore full connectivity without resorting to triple-redundant rings or some other very complex mechanism. What we get with contra-rotating dual-redundant rings is that the network bifurcates into two separate sub-rings; each sub-ring has full connectivity, but the two sub-rings do not communicate.

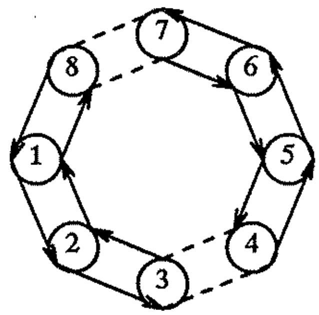

Figure 11.
Contra-rotating Dual-redundant Rings with Two Double-cuts

From the point of view of common sense, this situation is better than bringing down the whole network; however, the utility of two independent, non-cooperating sub-rings is questionable, and in any particular instance its utility would depend upon exactly where the breaks occurred. Most LANs can implement some form of dual-redundancy, but not more.

3.3 NETWORK MEDIA

LAN technology is supported by various media, including twisted pair, coax, infrared, and fiber optics. Infrared is a point-to-point technology like that used in remote controls for televisions, so its application is limited to being a potential point-to-point connection on a ring topology. Now we discuss the major three media in turn.

3.3.1 Twisted pair

Twisted copper wire, more commonly known as "twisted-pair," is common telephone cable. It is the least expensive medium available (from 5 to 25 cents per foot) and by far the easiest to install. It is readily available since it is in high-volume production for telephone use.

Its disadvantages, however, are serious. Since a pair of wires does not consume much space, they are typically bundled together, in tens or even hundreds, and given a common exterior covering for physical protection. The reason that each pair of wires was twisted in the first place was to minimize electromagnetic interactions among the pairs. However, any connection of any practical length also acts as an antenna, making this technology susceptible to outside electrical interference (EMI). While this has little effect on voice-grade analog transmissions, it creates two problems for data communications. One, it limits the speed of transmission, since a noise burst which scrambles a few bits at low speeds would garble many bits at high speeds. Two, it limits the distance a signal can travel. All signals attenuate with distance, but the problem is compounded for twisted-pair due to its antenna effect: the longer the connection, the more noise it picks up, until eventually the noise overpowers the original signal.

Two techniques are used to reduce this effect: shielding and repeaters. Shielded twisted pair is significantly less susceptible to EMI, but is of course more expensive. Repeaters can be used to increase overall communications distance by amplifying the signal and retransmitting it on another cable. Repeaters are expensive also, and are not foolproof: analog repeaters amplify noise as well as signal, so there is still a limit on overall circuit length.

Thus, twisted pair's primary advantages are its low cost and ease of installation. In a bus configuration, it is the easiest media to tap.

3.3.2 Coaxial cable

Coaxial cable, or "coax", comes in many forms, but all share a common structure. The copper wire carrying the signal is surrounded by a dielectric (insulator), which is surrounded by a solid or woven metal shielding layer, which is then covered with an external protective coating, typically plastic.

Baseband coax carries digital signals, i.e., the presence or absence of a voltage determines whether the data is "1" or "0". Baseband coax LANs generally support data rates of 1-10 Mbps, and of course normal time-division multiplexing can be used to create many logical channels from one physical channel.

Broadband coax carries many signals simultaneously, with each signal occupying a different part of the frequency spectrum. A single broadband (analog) LAN channel typically carries 1-5 Mbps, but by frequency multiplexing it is easy to get 20 or 30 such channels on a single cable. Broadband cables which carry data channels along with either voice channels (e.g. telephone) and/or video (e.g. television) channels divide the frequency spectrum in some agreed fashion so that the data channels occupy the otherwise unused frequency domains.

3.3.3 Fiber optic cable

Fiber optic cables are the newest and most promising LAN medium. The fiber optic core itself, with a diameter on the order of 0.006 inches, is covered by "cladding" which both protects the fiber and reflects light energy back into the fiber; the cladding is then covered with one or more protective layers to yield a cable of typical diameter 0.2 inches. Since the fiber optic cable is so small relative to its packaging, it makes sense to put many fibers in one physical cable. Two- and four-fiber cables are common, as is one 144-fiber cable used by the telephone company.

Light is injected into the fiber by either light-emitting diodes (LEDs) or lasers. Receivers use PIN diodes (with LED transmitters) or APD (avalanche PIN diode) devices (with laser transmitters) to achieve signaling speeds of 25 to 1500 MHz. The most common fiber uses a gallium arsenide light source to generate wavelengths of 800 to 900 nanometers; unfortunately, optical losses at this wavelength are rather high, about 3-4 dB per kilometer, so 10 km is a practical limit for repeaterless operation. A newer type of transmitter uses wavelengths of 1000 to 1600 nanometers which dramatically improves its optical properties; these systems can support repeaterless operation over 50 km.

The type of fiber used is important to performance. The oldest type is called "multimode step index" because the fiber core is surrounded by a glass cladding of lower refractive index. This discontinuity, or "step", in the refractive index causes light which hits the boundary to be reflected back toward the center of the core. Multimode step-index cables have a larger core than other types, which means that reflected light may take many paths through the core. This leads to *dispersion*, or *pulse spreading*, which limits the bandwidth of this type of fiber to 10-100 MHz.

A newer type of fiber is called "graded index". Rather than having a "step" in refractive index between core and cladding, the refractive index of the core gradually decreases from the center to the outside. Light is held in the core by refraction, rather than reflection, which minimizes pulse spreading and results in bandwidths of 200-1000 MHz. Although the physics of graded-index is more complicated than step-index, it is easier to produce. That, coupled with its large volume use for long-distance telephone lines, can sometimes make graded-index less expensive than step-index.

The newest technology is "single-mode" fiber. Its core is so small, a few micrometers in diameter, that light follows a single path and pulse spreading is eliminated; the resulting capacity is enormous, up to 50 GHz. This type of cable is very popular for long, unrepeated connections such as undersea transmission.

The benefits of fiber optic technology are numerous:

(1) very broad bandwidth
(2) inherent immunity from EMI and RFI
(3) no crosstalk or echoes or ringing
(4) inherently electrically isolated
(5) nonflammable
(6) able to withstand temperatures to 1,000 degrees C
(7) small physical size
(8) low weight (1 km of copper weighs 200 pounds, fiber 30 pounds)
(9) high security (no radiation leakage, taps easy to detect)
(10) high reliability (typical error rate is 1 bit in 10^9)

For short runs, fiber optic systems are more expensive than twisted pair or coax, largely due to the cost of the transmitters and receivers. For longer runs, fiber optic systems are competitively priced with coax.

There is one serious disadvantage: fiber optic cables are difficult to tap. Because the cores are so small, proper alignment of the cores requires special jigs, tools, connectors, and personnel. Although the situation is improving slowly, fiber will never be as easy to connect as copper. This is a concern if it is to be used in a maintenance environment which is low-technology.

Finally, we present a tabular comparison of the three technologies.

3.4 NETWORK SERVICE TYPES

Until the 1970s, data communication was performed using one of two service types: *message switching* or *circuit switching*. Using message switching, a host delivered an entire message to its interface message processor (IMP) and the network then used a store-and-forward technique to deliver the message to the destination. Each intermediate IMP along the path had to receive and acknowledge the full message from its predecessor before it would attempt a retransmission. See Figure 12.

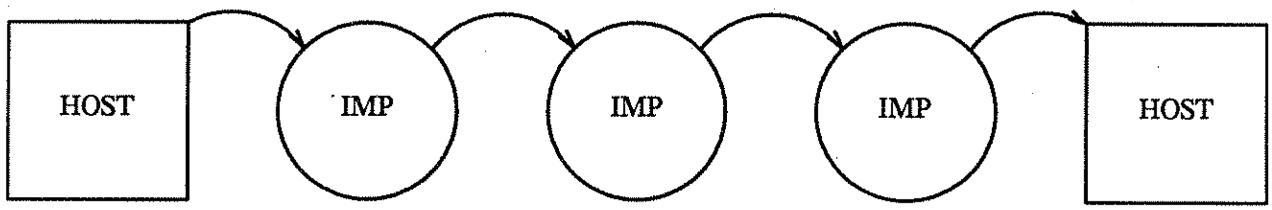

Figure 12.
Message Switching

The advantage of the technique was simplicity: either the whole message arrived at the destination all at once, or none of it arrived. The disadvantage was that it maximized delivery time while simultaneously maximizing the amount of network storage (main memory and/or disk) required to transmit the message.

Using circuit switching, all interior IMPs along the route from source to destination selected a path through the circuit switch (i.e., a mapping from an input line to an output line) before the message began transmission. Each IMP looked like Figure 13.

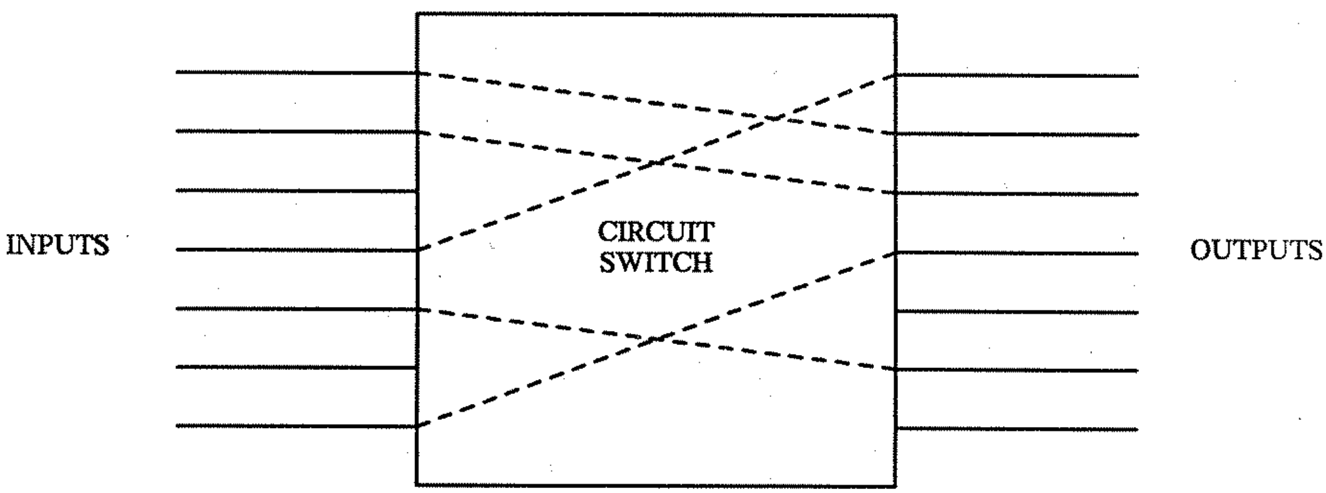

Figure 13.
Circuit Switching IMP

The telephone system still uses circuit switching to achieve an end-to-end "copper path" (in quotes because some routes might actually be microwave or fiber optic) between the source and destination. The advantages of circuit switching are that, once the path is set up, the route is stable (usually), the delays on the path are stable, and no other conversations can interfere with your portion of the channel bandwidth. The disadvantages are that there is considerable overhead in setting up the route before the circuit can be used, and any bandwidth not used by your circuit is unavailable to other circuits (i.e. wasted).

The compromise which emerged in the 1970s was *packet switching*. In the source host, messages are broken up into smaller units called *packets*. Packets can be either fixed size or variable size depending upon the implementation. The packets are then injected into the network and are routed to the destination. When all the packets have been received at the destination, they are reassembled into one message and delivered to the destination host. The software which breaks up the messages into packets at the source and which reassembles packets into a message at the destination is called the PAD -- the *packet assembler/disassembler*. See Figure 14. All local area networks use packet switching.

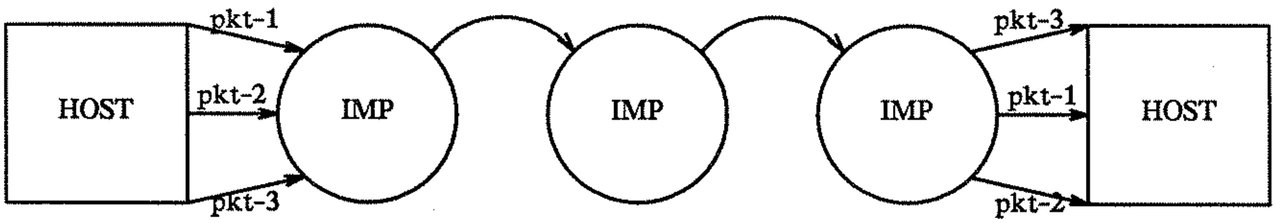

Figure 14.
Packet Switching

Depending upon the transport mechanism provided by the network, the packets may or may not arrive at the destination in the same order sent; in fact, they may not arrive at all. Thus it is the duty of the PAD to supply such features as *sequence numbers*, *acknowledgements*, and *automatic retransmission* of missing or flawed packets. However, these services are optional.

In a local area network data distribution is usually accomplished by *broadcast*, that is, the transmitter sends the packet to everyone, and only the receiver whose address matches the destination address of the packet bothers to copy it. This makes the *logical* channel seem circuit-like on a packet-by-packet basis, i.e., packets tend to be received in the order sent. Still, data errors can and do occur, so there must be a way to determine whether a transmitted packet reached its destination.

This need gives rise to two types of packet switching: *virtual circuits* and *datagrams*. A virtual circuit, sometimes called *connection-oriented service*, numbers and acknowledges each packet individually. It is called a *virtual circuit* because, even though it does not provide a physical end-to-end "copper path," it does provide a logical connection which maintains sequentiality. Datagrams are more like postal letters or telegrams and come in two styles: acknowledged and unacknowledged. If an unacknowledged datagram service, sometimes called *connection-less service*, is chosen, then packets are simply sent and arrive with some probability p (presumably p is high). If an acknowledged datagram service is chosen, the IMP transmits the packet and awaits an acknowledgement; if none is received, the packet is retransmitted.

The basic error rate for coax and fiber optic LAN media is very low, usually on the order of one bit in 10^9. (Twisted pair is not nearly so good.) By using a strong CRC checksum, the *undetected error rate*, i.e., those errors not caught by the CRC code, can approach one bit in 10^{12}. Thus, unacknowledged datagram service is not as risky as it sounds. While it is true that a message received in error will be thrown away by the receiver and the transmitter will never know, this is a very rare occurrence. Unacknowledged datagrams are also more efficient since there is no "stop-and-wait" for the message to be received and acknowledged. Unacknowledged datagrams are the basic service provided by the IEEE 802.3 contention bus and the IEEE 802.4 token bus.

One caveat is worth mentioning: the error rates of one bit in 10^9 or 10^{12} quoted above are the norm for *undamaged networks*. If the network is damaged or poorly tuned, the error rate can skyrocket.

In a wide-area network, virtual circuits, acknowledged datagrams, and unacknowledged datagrams are all distinct service types because of the interaction of packet sequencing and data errors. On a local area network where error rates are low and where (usually) only one packet is outstanding at a time, virtual circuits and acknowledged datagrams are very similar.

On most LANs, the choice of virtual circuit or datagram service is typically provided by the network interface software, and is logically located at the transport layer (ISO layer four).

3.5 NETWORK STANDARDS

In addition to the numerous marine standard-making organizations of which Sperry is already aware, there are a number of groups which set standards for data communications. We identify and briefly discuss the major players.

American National Standards Institute (ANSI)

Formed in 1918, ANSI is the principal standards-forming body for the USA. It is non-profit and non-governmental. ANSI is America's member body to the International Standards Organization (ISO). The lower level of ANSI consists of Technical Committees and Technical Groups; unlike IEEE, membership consists of qualified individuals rather than organizational representatives. Standards Committee X3 was formed in 1960 and chartered to investigate all standards for the computer industry. The Data Communications Technical Committee, X3S3, has seven task groups covering: planning, glossary, transmission formats, control procedures, system performance, signaling speeds, and public data networks.

Electronic Industries Association (EIA)

EIA is a trade organization with 4,000 members and 200 Technical Committees. EIA has produced over 400 standards since it was founded in 1924. Technical Committee TR-30, Data Transmission, has three subcommittees: signal quality, digital interfaces, and telecommunications network interfaces. Their work is primarily hardware oriented. EIA produced the RS-232 series and, more recently, RS-449 (HDLC).

European Computer Manufacturers Association (ECMA)

ECMA has 14 European manufacturers as members, plus other institutional members representing other standard-making bodies. ECMA has worked closely with CCITT.

Federal Information Processing Standards (FIPS)

FIPS identifies standards adopted for the federal government. The National Bureau of Standards drafts FIPS specifications. Over 80 FIPS have been adopted, with the goal of achieving commonality of data processing equipment and services.

Federal Telecommunications Standards Committee (FTSC)

FTSC is a federal government interagency advisory body. FTSC normally does not develop its own standards, but selects others for use in the government telecommunications industry. They have been active in developing standards for modems, system performance measurement, and the Data Encryption Standard.

Institute of Electrical and Electronics Engineers (IEEE)

Established in 1884, IEEE is an international organization headquartered in New York with additional offices in Washington, D.C., and in California. Its most ambitious effort is

Project 802, an attempt to define local area network standards. IEEE 802 is divided into six sections:

802.1 — the mapping between IEEE and the ISO seven-layer model
802.2 — logical link control
802.3 — contention bus
802.4 — token passing bus
802.5 — token passing ring
802.6 — metropolitan local area networks

International Organization for Standardization (ISO)

Founded in 1947, ISO is comprised of 90 member nations. ANSI is the USA's representative to ISO. Their most important work was produced in 1978 when they released the "Open Systems Interconnect Reference Model" — the seven layer model for all computer networks.

International Telegraph and Telephone Consultative Committee (CCITT)

CCITT handles data communications and consists of 15 study groups such as data communications interfaces, services, and transmission; digital networks; telephony; and tariffs. Study Group VII, formed in 1972 and now called the Data Communications Networks Study Group, is the most active group in the history of CCITT; they are responsible for public data network standards. Although the CCITT issues "recommendations", they have the effect of law in many European countries.

Altogether, these organization have produced a number of very important data communications standards. A few of these are summarized in Table 3.

ORGANIZATION STANDARD NUMBER	TITLE / DESCRIPTION
CCITT	
V series	Electrical characteristics; data signaling rates; modem standards; synchronous circuit standards
X series	Standards for public data networks
X.3	Packet assembler/disassembler
X.20	DTE/DCE interface for asynchronous transmission
X.21	DTE/DCE interface for synchronous transmission
X.22	DTE/DCE multiplex
X.24	DTE/DCE definitions
X.25	ISO lower three layers for public data networks
X.75	International packet switched services
ANSI	
X.344	Performance of data communications systems
X.357	Message formats using ASCII
IEEE	
802.1	Mapping from IEEE to ISO OSI model
802.2	Logical link control
802.3	Contention bus
802.4	Token bus
802.5	Token ring
802.6	Metropolitan networks
ISO	
2593	Connector pin allocations
4903	Assignments for 15-pin connectors
EIA	
RS-232	Asynchronous communications frame formats
RS-363	Signal quality standards
RS-449	HDLC framing

Table 3.
Some Representative Data Communications Standards

3.6 LAN INTERCONNECTS

Throughput this discussion we present the concept of the local area network as if it were a single entity, i.e., as if every device were connected to the same cable segment. This is a convenient abstraction and we will continue to use it.

However, any LAN which hopes for commercial success (as opposed to specialized LANs for connecting, say, laboratory instruments) must be able to interconnect with and talk the language of other vendors' products. This is achieved by *bridges* and *gateways*.

A *bridge* is a device which interconnects two segments of the same "style" network (i.e., both segments conforming to the same standard, even if they are made by different vendors). Thus Xerox markets a bridge for interconnecting one or more Ethernet segments; Ungermann/Bass makes a bridge for connecting token bus segments; IBM will soon market a bridge to interconnect token ring segments. A bridge is primarily a hardware device, using a microprocessor, buffer memory, and network interface circuits to move data traffic from one LAN segment to another. Bridges look like Figure 15.

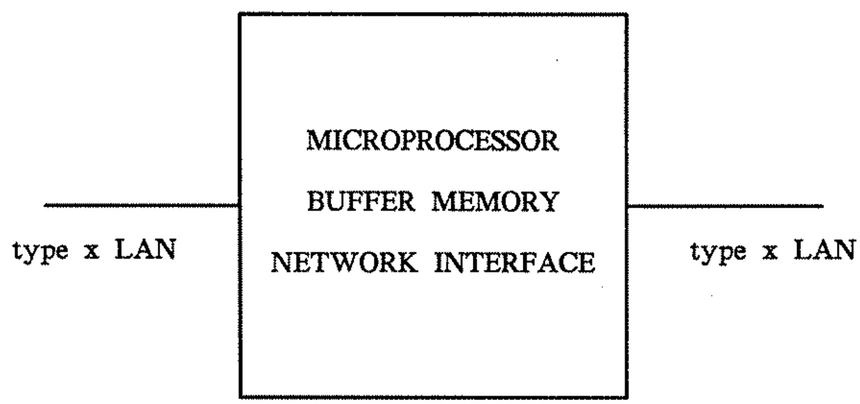

Figure 15.
A LAN Bridge

A *gateway* performs a similar function, but is a more complicated device because it interconnects networks of different types. In fact, the interconnected networks need not even be LANs; they could be wide-area networks. Gateways need the same internal parts as a bridge (microprocessor, buffer memory, network interface circuits) but have much more complicated software; the gateway serves as a *protocol converter* from network X to network Y. All gateways are bidirectional.

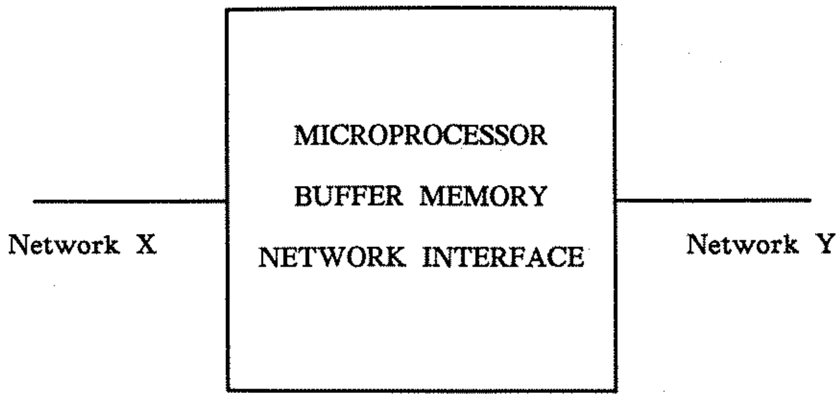

Figure 16.
A Gateway Between Different Computer Networks

Gateways are usually designed to interconnect two specific types of networks (the complexity of making the gateway truly general purpose would make it uneconomical). Because the Xerox Ethernet (nearly IEEE 802.3) has a five-year headstart on other LANs, the most popular gateway is an Ethernet-to-X gateway. The second most popular is a gateway to a public packet switched network which conforms to CCITT standard X.25; these are X.25-to-something gateways.

In the very near future we will see all three gateway combinations for interconnecting IEEE 802.3 (contention bus), IEEE 802.4 (token bus), and IEEE 802.5 (token ring).

Performance modelling of a bridge is fairly straight-forward; the bridge just looks like a common node on two otherwise similar LANs. However, performance modelling of gateways, where there is protocol conversion and packet reformatting in addition to mere buffering and repeating, is an open question.

3.7 SELECTING THE DRIVING FACTORS

We have identified six factors as being the driving forces behind the eventual decision of which LAN to implement. These six factors are:

(1) System reliability
(2) Baseband vs. broadband transmission
(3) Copper vs. fiber optic cabling
(4) Adherence to international standards
(5) Cost
(6) Network performance

As expected, no single LAN or even LAN type is best in all six categories. Factors which make a system reliable, for example, also make it expensive. For the moment, we simply relate the general networking considerations which we have discussed thus far to the six selected criteria; the actual selection of an appropriate engineering compromise appears in section six.

System reliability

Reliability is enhanced by distributed control, redundant cabling and transceivers, a topology free from single-point-of-failure devices, strong error-detection and error-recovery software, and low error rates on the cabling medium. Table 4 shows a gross comparison of some issues.

RELIABILITY ISSUE	LEAST RELIABLE	MORE RELIABLE	MOST RELIABLE
Basic topology	star	bus or ring	
Modified topology	star with redundant switch	dual-bus with redundant transceivers	contra-rotating dual-redundant rings
Media	twisted pair	coax	fiber optics
Service type	unacknowledged datagram	acknowledged datagram	virtual circuit
Protocol	centralized	distributed	
Error checking	parity	CRC-16	CRC-32

Table 4.
Reliability Issues

Baseband vs. Broadband

This decision hinges upon whether or not the network will carry real-time voice (telephone) or real-time video (television) in addition to its primary function of data distribution. If the decision is to use baseband, costs are lowered, maintenance and maintenance skill level are both decreased, and whatever voice/video signals are present on the ship have to be carried by a completely separate cable plant.

If the decision is broadband, then integrating the voice and video channels on the cable plant used for the network is very simple; the RF maintenance required for video is about the same as that required for the network; and total cable costs are reduced since there is only one cable plant.

A hybrid solution is to use baseband coax for data and twisted pair for voice, eliminating video. Using twisted pair for voice is much simpler than maintaining the second cable plant which would be required if video were included.

Finally, the baseband vs. broadband decision directly affects the choice of media. Twisted pair has limited broadband characteristics and thus is not recommended; baseband works fine on single-cable coax; broadband on coax needs a frequency translator for both single-cable and dual-cable topologies; and today's fiber optic systems support baseband much better than broadband. Fiber optics can support broadband by multiplexing different wavelengths of light on the same fiber, but the transmitters and receivers are very expensive. For the moment, broadband fiber optic systems are impractical; nevertheless, since fiber optics support very high data rates, voice and video which are normally transmitted using RF techniques can instead be digitized and sent baseband over fiber optics using simple TDMA strategies.

Table 5 summarizes these results.

CHARACTERISTIC	BASEBAND	BROADBAND
Supports voice/video	no (unless digitized)	yes
Topology	star, bus, ring	bus
Media	twisted pair, coax, fiber optics	coax
Skill level for installation and maintenance	lower	higher
Cost	lower	higher

Table 5.
Baseband vs. Broadband Characteristics

Copper vs. fiber optics

Shielded twisted pair offers economy and simplicity of installation for short runs; for long runs it has inadequate noise immunity. Baseband single-cable coax systems offer ease of installation and simplicity of transmitter/receiver; they are, however, subject to failure if the cable is shorted. Broadband single-cable coax systems offer the same ease of installation but introduce a critical element: the head-end frequency translator. The frequency translator can be made fairly reliable since we have 30 years of experience with this device in the CATV industry. Fiber optic cables are very easy to run, but the skill level for making connections is considerable higher than for installing coax. Maintenance of fiber optics in a low-technology location is a genuine concern.

Table 6 summarizes these characteristics.

CHARACTERISTIC	TWISTED PAIR	BASEBAND COAX	BROADBAND COAX	FIBER OPTICS
Supports voice/video	limited	no	yes	only if digitize
Running cable	easy	easy	easy	easy
Installing connectors	easiest	easy	easy	difficult
Maintenance	easiest	moderate	harder	hardest
Failure modes	shorted cable; broken connector; subject to EMI and RFI interference	shorted cable; broken connector	shorted cable; broken connector; failure of head-end retransmitter	crushed cable; misaligned con
Cost	lowest	moderate	moderate	highest

Table 6.
Copper vs. Fiber Optics

Standards

Either the network adheres to an applicable international standard or it does not. Both strategies have advantages.

CHARACTERISTIC	NON-STANDARD	STANDARD
Cost to produce	lower	higher
Level of effort to produce	higher	lower
Interoperability	difficult	assured
Multi-vendor networks	unlikely	possible
Network diagnostics	must be created	can be purchased
Installation and maintenance	more difficult	less difficult
Network performance level	higher	lower
Gateways	not available	available
Performance models	initially unknown	known
Proprietary market position	assured	weakened

Table 7.
Considerations for Standardization

Cost

Cost is an extremely important, but elusive, factor. A network with vast functionality but high interface cost may not sell, yet a low cost network with not enough features may not sell either. As always, it is an engineering trade-off to determine what functionality can be sold for what price.

COST FACTOR	DECREASES COST	INCREASES COST
Basic topology	star	bus, ring
Bus topology	single cable	dual cable
Ring topology	single cable	redundant ring
Ring switching	unidirectional	reversible
Ring interface	in wiring center	in each device
Media	twisted pair	coax, fiber optic
Installation and maintenance	twisted pair	coax, fiber optic
Reliability	non-redundant	redundant
Protocol	centralized	distributed
Voice/video	not supported	integrated with data
Baseband/broadband	baseband	broadband
International standards	ignore	adhere

Table 8.
Factors Affecting System Cost

Hidden costs resulting from topology

Depending upon the topology chosen, there are always costs unique to that topology which might be hidden in a high-level discussion of costs. This section attempts to add more detail on this issue.

Star topology

A conventional star topology has no hidden costs. A straight-forward implementation needs a single central switch, one cable from each device to the central switch, and appropriate connectors.

Bus topology

Baseband bus:

A baseband bus system requires circuitry to convert digital information into baseband signals and vice versa. This is typically done outside the device or computer interface itself in an external transceiver. In Ethernet systems, for example, each connected device needs a transceiver (about $300) and a transceiver cable (from $10 to $60). The transceiver costs are sometimes overlooked in a baseband transmission system.

Broadband bus:

The most common broadband bus topology uses a single cable and a head-end remodulator (frequency translator). The costs here include the device or computer interface which must include one RF modem per interface. The interface will be less expensive if the modem is fixed-frequency and more expensive is the modem is frequency-agile. Generally, for a given network vendor there is no choice; only one type or the other is supplied. Of course the head-end remodulator must be of the same type (fixed or agile). The frequency translator itself is a one-time cost per network. Typical units are about $3,500.

The less common type of broadband topology uses dual-cable, with separate cables for the head-end's inbound and outbound traffic. If this is chosen then cable costs are doubled.

Ring topology

Unmodified, the ring topology is generally less expensive than the equivalent broadband bus topology because no head-end remodulator is necessary. However, the simple ring is often judged inadequate from the point of view of reliability. To counter that objection, the star-ring topology is used which employs the wiring center concept. As explained earlier, the wiring center allows the backbone topology (and media) to be static, thus less subject to injury and hence more reliable. All device connections are made to the wiring center, not the backbone, and each such connection is "protected" by a bypass circuit.

For the Proteon ProNET network, the wiring center adds a cost of $295 for a 4-port connection, $630 for an 8-node connection, or $950 for a 12-node connection. This increases the per-device interface cost, but is essentially required if reliability is to be achieved.

Network performance

One important measure of network performance is its *throughput*, defines as the fraction of bus capacity which is actually usable for carrying data. Thus the throughput measure excludes the portion of bus capacity used for packet scheduling (e.g., polling, collisions, tokens). For any particular product, throughput is determined by both the network type and the efficiency of the vendor's implementation; since the latter is difficult to determine without extensive measurement, we rate four types of networks based upon experience and expectation:

Network Type	Throughput (fraction of bus capacity)

contention bus	0.3 — 0.5
avionics bus	0.4 — 0.6
token bus	0.6 — 0.8
token ring	0.8 — 0.9

4. GENERAL NETWORK TYPES

4.1 Conventional star

4.2 MIL-STD-1553B

4.3 Integrated Services Digital Network

4.4 IEEE 802.3 Contention Bus

4.5 IEEE 802.4 Token Bus

4.6 GM Manufacturing Automation Protocol

4.7 IEEE 802.5 Token Ring

4.1 Conventional Star

The star topology is not a local area network but a device interconnection scheme. As shown in Figure 17, the star requires that each attached device be connected by its own separate cable to the central switch.

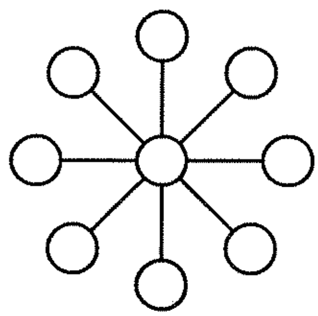

Figure 17.
Conventional Star Topology

There are no local area network standards based upon a star topology. There are three primary reasons:

(1) The star architecture introduces a reliability question because of its unique, crucial component (the central switch).

(2) The wiring length for a star topology is always equal to, and usually much more than, the comparable bus or ring topology.

(3) Stars usually use a master/slave, polling-based communications protocol. This restricts attached devices to transmitting only when they are polled, rather than whenever they have data to share.

4.2 MIL-STD-1553B

Mil-Std-1553B is not a local area network standard, but a device interconnection standard for military avionics systems. Mil-Std-1553 was first developed to provide some commonality to aircraft avionics systems and to reduce the wiring costs associated with then-common star topologies (this was 1968!). Modified to version A in 1972 and version B in 1978, the standard is still in use, primarily for the military. A special task force on high-speed avionics data busses has been formed by SAE to draft a replacement standard, but their work is easily two years from completion.

The current version of 1553B provides for communications at 1 Mbps over twisted-pair media using a centralized bus controller. Manchester encoding is used for additional noise immunity. There can be up to 31 "terminals" (meaning sub-controllers) on a bus, each of which can handle up to 30 "stations" (devices). The protocol is a basic master/slave polling sequence administered by the bus master. In version B there is an alternative mode in which bus mastership moves from "terminal" to "terminal". The topology is shown in Figure 18.

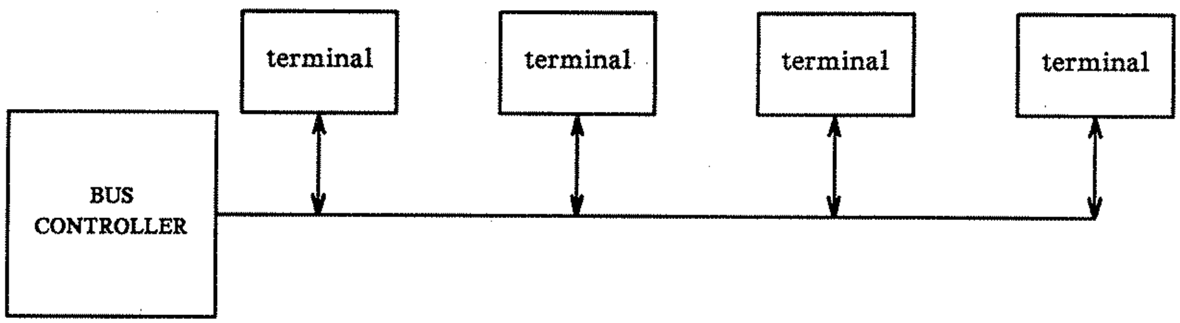

Figure 18.
MIL-STD-1553B Topology

The format of transmissions is as follows. Each "word" is composed of 20 bits: 1 is a parity bit, 3 are header bits, leaving 16 for commands/data. The three header bits are used as a signal to receiving terminals that a word is forthcoming, and a special encoding is used to signal whether the word is a control command or user data. The 3-bit header defines a transition at the 1.5 bit-time position. If the transition goes one way, it defines the next word as a command; if it goes the other way, it is user data. Note that the protocol is at most 80% efficient since 4 of every 20 bits are overhead; efficiency drops to more like 50% when one considers interframe gaps, polls, and control messages.

A command word sent by the bus master at the beginning of each message provides the 5-bit address of the receiving terminal, a transmit/receive bit, a 5-bit subaddress field (for devices attached to the terminal), and a 5-bit data count field. Thus the longest transmission is a command word followed by 32 16-bit data words. The specification also calls for a 4 microsecond intermessage gap between sequential messages.

As always, the introduction of a single crucial element, the bus controller, gives rise to a concern for reliability. Since this concern is accentuated by the military and by the application environment (aircraft avionics), this problem has been well-studied. As a result, all crucial implementations use dual-redundant bus controllers, and dual-redundant busses are often provided in addition.

As a shipboard system, Mil-Std-1553B seems to provide few advantages. The hardware is standardized, but because it is mil-spec it is generally expensive. There is no standardization of software. The 20-bit word format and the limit of 32-word messages would be tolerable in a control-oriented application like avionics, but general purpose data sharing such as file transfers would be slow and inconvenient. The lack of software standardization, and the lack of adherence to any networking model such as the ISO OSI model, is a serious drawback to its utility as a LAN standard.

4.3 Integrated Services Digital Network

ISDN is not a local area network standard but a projected worldwide public telecommunications network that will service a wide variety of user needs. The ISDN will be defined by the standardization of user interfaces, and implemented using digital switches and digital data paths, supporting a broad range of traffic types (i.e. mixed voice and data). Because the ISDN does not yet exist and because it is defined by a still-evolving set of standards, we merely summarize its intent.

An integrated digital network (IDN) was proposed as early as 1959, but needed the push of the first large scale digital telephone switch (the Western Electric 4ESS) in 1976 to get started. Since then the telephone industry has largely standardized on using pulse code modulation to digitize incoming analog signals, and them using TDMA to multiplex many such signals on one high-speed digital trunk.

The basic service to be provided to end-users is (presently) a 64 Kbps digital channel. This channel speed was chosen as being the best for its primary target, digitized voice. Ordinary telephone channels support a bandwidth of approximately 4,000 Hz. Digital telephone channels take 8,000 7-bit samples per second, then add 1 bit for control information, for a total data rate of 64 Kbps per voice-grade channel. If interfacing slower devices, multiple devices could be multiplexed onto one of these channels for economy. If interfacing faster devices, the data could be "split" into multiple streams, with the streams carried in parallel by multiple channels. Figure 19 shows some of the basic uses envisioned for ISDN.

Bandwidth	Telephony	Data	Text	Image
digital voice	telephone	packet-switched data	telex	
(64 Kbps)	circuit-switched data	teletex		
	leased circuits	leased circuits	leased circuits	
	info retrieval	telemetry	videotex	
	funds transfer		facsimile	
		info retrieval	info retrieval	info retrieval
		mailbox	mailbox	surveillance
		electronic mail	electronic mail	
		alarms		
wide-band	music	high-speed computer		TV conferencing
(> 64 Kbps)		communications	teletext	
				teletext
				videophone
				cable TV

Figure 19.
Potential Services for ISDN

In terms of shipboard systems, we need only be aware that these services are contemplated; they will take 10-20 years to be in common use. By holding open the possibility of providing 64 Kbps digital channels on the shipboard LAN we will have a future option of using satellite-switched ISDN.

4.4 IEEE 802.3 Contention Bus

IEEE 802.3 is an international standard for baseband communication using a bus topology and Carrier Sense Multiple Access with Collision Detect (CSMA/CD) as the network access mechanism. Signaling on the medium uses Manchester encoding for the data; as a result any one transmission produces constant energy on the bus regardless of the data stream. Multiple simultaneous transmissions produce "collisions" and are detected by sensing a rise in bus energy level.

IEEE 802.3 specifies a 50-ohm baseband coax cable (see Figure 20) as the acceptable network medium. It provides low impedance and allows a large number of stations to be connected to the medium through simple coax taps. Any one network segment is limited in length to 500 meters, although segments can be interconnected by bidirectional repeaters. Even so, the maximum end-to-end separation of any pair of devices on the network is limited to 2,500 meters, and the whole network is limited to 1,000 attached devices. Transmission speed is fixed at 10 Mbps, although that is not a fair indication of the overall network throughput due to the operational characteristics of CSMA/CD access schemes.

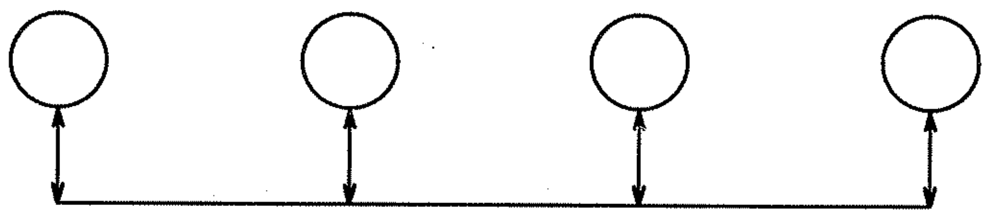

Figure 20.
Linear Bus Topology

Figure 21 shows a possible transmission scenario for an arbitrary pair of nodes.

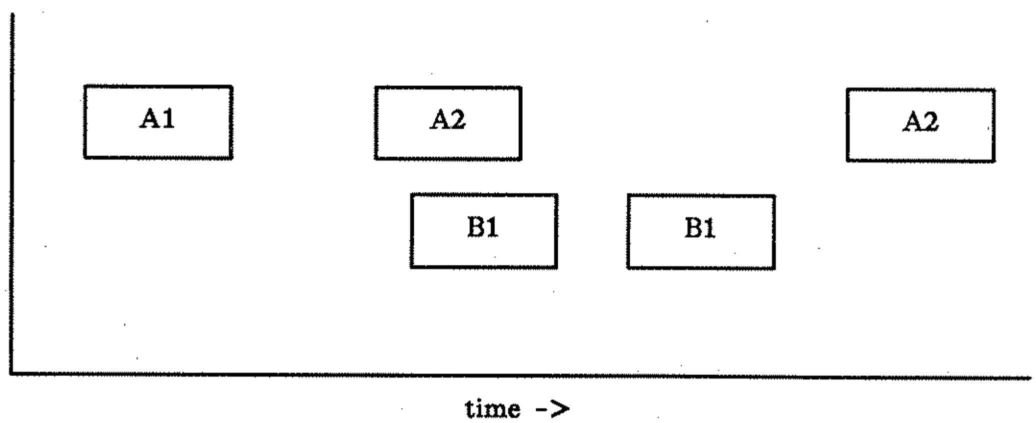

Figure 21.
CSMA/CD Transmission

The x axis shows the passage of time from left to right; the two horizontal displays on the y axis are the transmissions of two arbitrary nodes called A and B. At time t_0, the bus is idle. Shortly thereafter node A senses the channel, finds it idle, and transmits message A1. Transmission finishes and the bus goes idle. Somewhat later node A has another message to send; it senses the channel, finds it idle, and begins transmitting message A2. At approximately the same time node B receives a message to transmit. Because node B is located distantly from node A, it listens to the channel and hears it idle because message A2 has not yet physically propagated from node A to node B. Thus node B legitimately begins transmission, but messages A2 and B1 do overlap, or collide. Both nodes sense the collision, stop transmitting, and both enter the rescheduling discipline called "binary exponential backoff". It is this backoff mechanism (discussed shortly) which gives the contention bus its characteristic performance. After a delay mandated by the backoff algorithm, message B1 is retransmitted (this time without collision), and at a still later time, message A2 is retransmitted.

The binary exponential backoff algorithm is unique among retransmission schemes. It works as follows:

For the j^{th} collision, compute:

```
j := min (1, 10)
k := random ( 0 .. 2^j )
wait k slot times (= 512·k bit times)
sense channel
transmit when channel is next idle
```

For each of the first ten collisions, the delay time is a random number of "slot" times (multiples of 512 bit times which, at 10 Mbps, makes a slot time about 0.5 microsecond) where the upper bound of the random number field is increasing (binary) exponentially with each successive transmission. Thus on the first collision, the wait might be 0, 0.5, or 1 μsec; on the tenth, it is between 0 and 1 ms. Collisions 11 through 15 are treated the same as the tenth; if there is a sixteenth collision then it causes the packet to be abandoned and an error to be reported to upper level software.

This algorithm is clearly intended to optimize the channel performance at low loads. When loading is light, collisions are rare, and those which do occur are subjected to short backoffs (delays). As load increases, the backoff procedure artificially drives down the network offered load by forcing retransmissions to be spread over longer and longer time periods. This approach directly implements the classic tradeoff between throughput and delay — namely, that network throughput can be increased at the cost of increasing average message delay.

As network offered load increases, throughput increases until it reaches a maximum; at that point it remains steady (i.e. the protocol is stable). The maximum throughput is less than bus capacity by an amount dependent primarily upon bus length and message size. As bus length increases, propagation delay increases proportionately and the "collision window" (the amount of time each message is vulnerable to a collision) for each message increases, causing throughput to decrease. For any fixed bus length the size of the collision window is fixed; thus for increasing message size the percentage of the message subject to collision decreases. As a result throughput increases with increasing message size. Throughput vs. offered load is depicted in Figure 22(a).

End-to-end message delay comes from three sources: a queueing delay as the message awaits transmission, a propagation delay when it is transmitted, and a backoff delay whenever it collides and is later retransmitted. The resulting delay is exponential with offered load and is shown in Figure 22(b). What is unique about the performance of the contention bus is that, at low loads, end-to-end delay is nearly zero. But as offered load grows, delay grows exponentially.

As a result, throughput using CSMA/CD reaches a practical maximum at 20% to 40% of bus capacity. For a 10 Mbps bus, this means an effective capacity of 2-4 Mbps.

Figure 22(a).
Throughput vs. Offered Load for CSMA/CD

Figure 22(b).
Delay vs. Offered Load for Binary Exponential Backoff

4.5 IEEE 802.4 Token Bus

IEEE 802.4 is an international standard for bus topology networks which use token passing as their means of network access. All 802.4 networks to date have been implemented using broadband signaling. All stations attach with simple coax taps.

Network access is controlled by a special data element called the *token*. The token holder is momentarily the network master and may transmit messages, or even engage in sub-protocols, for a limited period of time. When a station has transmitted all its messages, or when certain timers expire, the station must pass the token to a known successor. The orderly progression of the token from station to station actually forms a logical ring on a physical bus. See Figure 23.

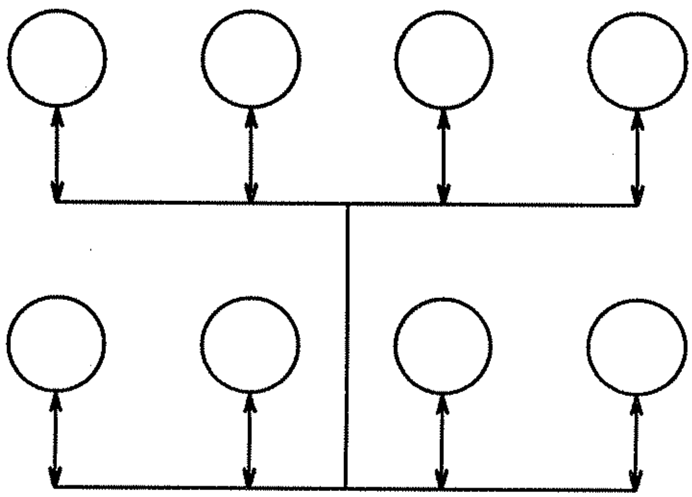

Figure 23.
Branching Tree Bus Topology

At network startup each station is assigned a unique logical address. During the startup period each station adds itself to the logical ring using a contention process. When ring membership has been established, token passing begins and the participating stations share the network capacity sequentially.

The token is a crucial data element since it alone confers bus mastership. It is a data message consisting of several fields:

preamble (1 or more octets depending upon bus speed)
start delimiter (1 octet)
control information (1 octet)
destination address (2 or 6 octets)
source address (2 or 6 octets)
optional data (0 or more octets)
frame check sequence (4 octets)
end delimiter (1 octet)

Within the above fields, it is the "control information" field which denotes that this frame is indeed a token.

Message frames have the same format, with the control information field denoting "data message". There are various types of protocol messages (those used for controlling the token passing sequence); they have the same format and the control information field defines their type and function.

The protocol automatically handles dynamic topology. Each station has a counter, initialized to a network-wide variable, which counts down by one each time the token passes through the station. If the network load is not too high and this counter counts down to zero, the station sends a special protocol message called *solicit_successor* whose purpose is to determine if there is a currently passive station which wishes to become active. After sending this protocol message the counter is reset to its initial value. Any station wishing to join the logical ring responds to *solicit_successor* with another control message. If the station sending the message hears no response, it passes the token normally; if it hears exactly one response, it passes the token to the new node and the new node will pass it to the original station's successor; if multiple responses are heard, then a contention algorithm is used to determine which one station will be allowed to join in response to this solicitation; all unsuccessful stations must wait until another *solicit_successor* message is issued. In this way the protocol gracefully adjusts to dynamic station membership.

IEEE 802.4 supports four priorities of messages. In descending order of priority they are called synchronous, urgent asynchronous, normal asynchronous, and time-available. Of these, only synchronous messages are guaranteed a level of service; the other classes are carried on a "best effort" basis. Any 802.4 station must either implement the synchronous class alone or else it must implement all four priority classes. Note that the concept of priority applies to a message, not a station.

When a station receives a token, it resets a timer to a network-wide variable which defines how long a station may service its synchronous class. Then service begins for the synchronous class, and continues until either this timer expires or all the synchronous traffic has been transmitted. If the station is not implementing the priority option, then the token is passed at this point.

If a station is implementing the priority option, it determines whether or not it may service any of its asynchronous classes. Each asynchronous class has associated with it a timer which tells the elapsed time from the last receipt of the token at this priority at this station until now; if that elapsed time is less than a network-wide parameter called the *Target Rotation Time* for that access class, then a timer is loaded with the difference between the target rotation time and the time already expended on the last token cycle (as measured by this priority class at this station). This algorithm is repeated in turn with different timers and different *Target Rotation Times* for each asynchronous class. A station can serve each asynchronous class in turn until either it has transmitted all its messages at that priority or the *Target Rotation Time* has been reached; in either case the station must then pass the token to the next lower priority or to the next physical station.

The net result is that the network token rotation time varies with offered load, and is responsive to the loadings at each priority class. If a network uses only the synchronous class, determining token rotation time is easy; if the priority option is used, it is difficult.

Like contention-based network access, token passing has a characteristic performance curve. Figure 24(a) shows throughput vs. offered load for 802.4. Throughput is normally equal to offered load until some limit is reached; the difference between that limit and bus capacity is a measure of protocol efficiency; that is, it is the network overhead (tokens, propagation delays, and network control messages) which prevents throughput from being equal to bus capacity.

Figure 24(b) shows the end-to-end delay as a function of offered load. Note that even at near-zero loads, network access is non-zero (unlike contention systems) because a station must first wait for token arrival before transmitting. But as offered load grows, end-to-end delay with token passing increases more gradually than with contention schemes.

Overall, IEEE 802.4 can be expected to deliver about 80% of its bus capacity as usable data throughput.

Figure 24(a).
Throughput vs. Offered Load for Token Passing

Figure 24(b).
Delay vs. Offered Load for Token Passing

4.6 GM Manufacturing Automation Protocol

The development of GM MAP is an event which will change the course of history with regard to factory automation. Before 1980, every plant or assembly line which GM built involved a specification for the control system which was bid on and won by one of the major suppliers of industrial automation equipment (e.g. Allen-Bradley, Modicon, etc.). But these products would never talk across vendor boundaries unless GM financed custom hardware and software to overcome the problem. In some instances the cost of developing the custom interfaces was one-half the cost of automating the operation. In addition, the custom interfaces took a long time to build, and once in place were inflexible. Even the wiring was a nightmare.

Faced with stiff international competition, GM found that it simply could not change its production steps or rates fast enough to respond to a rapidly changing marketplace. In 1980 GM set up the MAP task force whose charge was to define hardware and software standards which would assure interoperability among the various vendors' products.

The task force decided not to invent new standards, but to build on existing ones; thus MAP is primarily a collection of existing standards used to achieve a specific functionality (there are some parts of the standard which were invented by the task force). In a nutshell, what MAP accomplishes is that it defines the hardware mechanisms and the software framing such that specific types of information can be routed from one type of industrial control equipment to another.

A pictorial overview of MAP is shown in Figure 25.

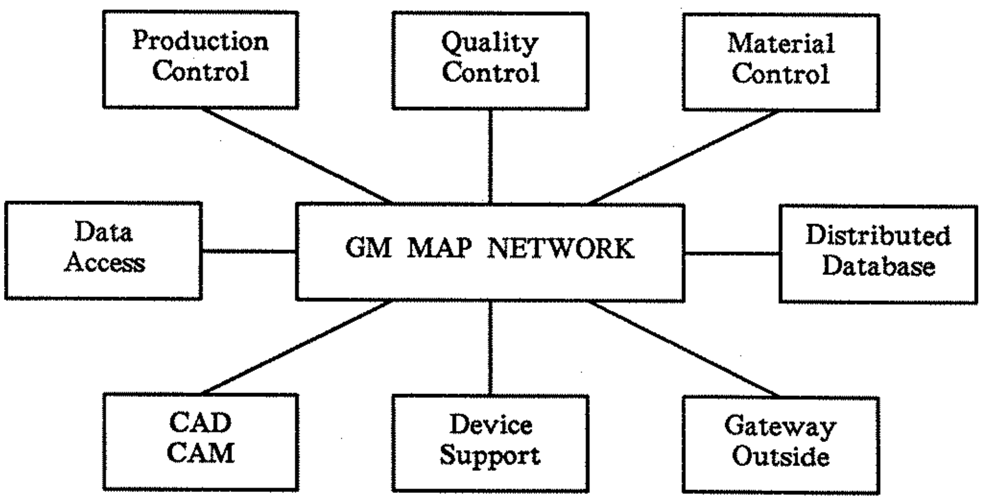

Figure 25.
Pictorial Overview of GM MAP

GM MAP has been through four versions so far; it will undoubtedly continue to change. The most recent version is numbered 2.1 and was adopted in March 1985. The very first MAP 2.1 networks will be installed in 1986.

What does MAP provide? For each ISO layer, it selects one or more services (standards) which are defined to be acceptable. Figure 26 shows the mapping taken from MAP version 2.1.

OSI LAYER	MAP VERSION 2.1
7 - Application	Network Management Directory Service Manufacturing Message Format Standards (MMFS) File Transfer and Access Management (FTAM) ISO Communication Application Service Elements (CASE) Kernel
6 - Presentation	Null
5 - Session	ISO Session Kernel
4 - Transport	ISO Transport Class 4
3 - Network	ISO Connectionless Network Services (CLNS)
2 - Data Link	IEEE 802.2 Link Level, Control Class 1 IEEE 802.4 Token Access
1 - Physical	IEEE 802.4 Token Bus on Broadband Media

Figure 26.
ISO Model vs. GM MAP 2.1

The two software standards which were defined specifically for MAP are MMFS (Manufacturing Message Format Standards) and FTAM (File Transfer and Access Management). MMFS defines the framing of specific control commands (e.g. program upload and download, read status of device, write control command to device, etc.) while FTAM defines the structure of files and how they are to be accessed.

Besides the MAP task force, the major players are the manufacturers of the MAP-compatible products. Concord Data Systems was originally chosen as the implementor of the 802.4 broadband network to support all the data transfers; their "lock" on the networking components has just recently been broken by Industrial Networking Inc. (INI), a new company founded by Ungermann-Bass and General Electric. Both companies are now marketing 802.4 broadband products which use software and/or firmware to implement the higher-level MAP specifications. While Concord has momentum from its choice as the first vendor, INI is gaining rapidly as a result of their introduction of gate-array products which implement some of MAP in firmware. Someday all of MAP will be implemented in a chip set, but when that will occur cannot yet be predicted with any certainty.

It is important to note that although MAP products can be purchased, MAP is still a moving target. To encourage vendors to build products, GM is committed to supporting version 2.1 and all successive versions for a minimum period of two years. However, given the rapid changes of technology, it is doubtful that the MAP specification will go unchanged for that length of time.

As an example, MAP currently specifies 802.4 broadband as the physical layer; Concord Data System's first implementation was a 5 Mbps version. Upon reexamination, the task force now thinks that this standard is more appropriate for a backbone than for the actual device interfaces on a production line, and that 10 Mbps would be a better choice for a backbone than 5 Mbps. In addition, the choice of broadband implies an RF modem in each device interface which is fairly expensive; an alternative is carrier-band technology which would be less expensive to implement and would also be media-independent, thus paving the way for fiber optics. Most knowledgeable observers expect a change in 1986 to incorporate some version of a carrier-band network.

In summary, MAP is a very important development for factory automation. Its application to shipboard systems is tenuous at best. It could be an advantage *if* shipboard control systems resemble factory automation control systems (i.e. if the message and file formats defined by MMFS and FTAM are meaningful) *and* if high-volume production of MAP-compatible hardware brings down the price. Since MAP is still changing, it is really too soon to count on the latter.

4.7 IEEE 802.5 Token Ring

IEEE 802.5 is an international standard for networks arranged in a closed circular loop; see Figure 27. In general, stations inserted into the ring act as repeaters; i.e., whatever is received as input is simply repeated as output. Periodically a station may transmit, in which case it substitutes its own messages in place of repeated bits. Because of the closed loop, the message will eventually return to the transmitter who is responsible for purging it from the ring. Permission to transmit is conferred by a token, similar to the token bus.

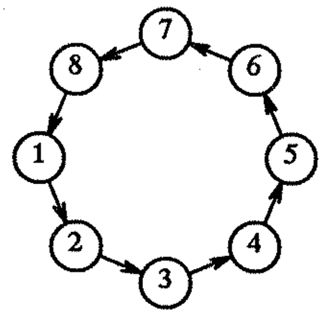

Figure 27.
Ring Topology

There are several main differences between the token bus and the token ring: a station plays an active role in the ring even while not transmitting his own messages (i.e. the station is a ring repeater); in the ring a station's physical neighbor and logical neighbor are the same, unlike the bus; even though the ring uses a bus architecture, it is comprised of point-to-point connections, thus permitting mixed media; the 802.5 standard permits eight priorities as opposed to four in 802.4; since path lengths (and hence propagation delays) are generally short, and station latencies are short (nominally 1 bit time), and the token itself is short (24 bits), the ring has higher inherent efficiency than the bus.

Each station has a receiver and transmitter separated by a (small) shift register. As bits are received from the ring they are examined and put into the shift register; shortly thereafter they emerge from the shift register and go to the transmitter. Depending upon the functionality of the station the shift register varies in length from 1 to 30 bits. Thus a station "sees" ring transmissions through a very small "window".

IEEE 802.5 requires a specific frame format as shown in Figure 28.

SFS		FCS					EFS	
SD	AC	FC	DA	SA	INFO	CRC	ED	FS

Figure 28.
IEEE 802.5 Frame Format

The fields are:

SFS — start of frame sequence field:
SD — start delimiter (1 octet)
AC — access control field (1 octet)

FCS — frame check sequence field:
FC — frame control [type] (1 octet)
DA — destination address (2 or 6 octets)
SA — source address (2 or 6 octets)
INFO — data field (0..N octets)
CRC — cyclic redundancy code (4 octets)

EFS — end of frame sequence field:
ED — end delimiter (1 octet)
FS — frame status (1 octet)

The token is a special case; it uses only three bytes of the frame: start delimiter, access control field, and end delimiter. When a three-octet token passes a station which wishes to utilize the token, it modifies it to change it into a full legal frame.

The access control octet contains a 3-bit priority sub-field and a 3-bit reservation sub-field. The priority field defines the current ring priority; the reservation sub-field allows stations to "bid" for a token of higher or lower priority. The general operation is that a station which enqueues a high-priority message will alter the next access control octet to "reserve" a token at the appropriate priority. If, after one token rotation, no other station has bid a higher priority, the requesting station can alter the token priority to equal that of its message. Having altered the token priority, however, this station is responsible for eventually lowering it to its original value.

The token holding mechanism is simpler in the token ring than the token bus. Each station operates a token holding timer (THT) and a return-to-repeat timer (TRR). The THT defines how long a station can transmit; the TRR defines how long a station can continue to hold the token after transmission has stopped while waiting for the station's last message to return to the sender. Thus the maximum waiting time in any one station is (THT + TRR). This make the ring's performance analysis simpler than that of the bus.

At all times ring performance is observed by an "active monitor". This station performs special watchdog functions and assures that the token is present, correct, and circulating at the appropriate priority. All stations are capable of becoming the active monitor, and one is "elected" at ring startup to serve this special function. If the active monitor should fail its duties are assumed by the "standby monitor".

The standby monitor has special duties similar to those of the active monitor, including watching the active monitor. If the active monitor fails to send special control messages periodically, then the standby monitor assumes the active monitor's role. Any station is capable of being the standby monitor, and one station is "elected" to serve at ring startup to serve this special function. If the standby monitor fails, some other station will assume its duties.

In terms of performance, the IEEE 802.5 token ring has throughput and delay curves which look similar to those of the token bus, although its absolute performance is better. The short path lengths, short station latencies, and small token size of 802.5 combine to give it inherently lower overhead than the token bus protocol.

As a result, the token ring can deliver 80 to 90% of bus capacity as data throughput.

5. SPECIFIC NETWORK PRODUCT EVALUATIONS

5.1 SYTEK LocalNet

5.2 BRIDGE COMMUNICATIONS Ethernet

5.3 INTEL Ethernet

5.4 INTERACTIVE SYSTEMS/3M LAN/I, LAN/II, LAN/PC

5.5 CONCORD DATA SYSTEMS Token/Net

5.6 UNGERMANN-BASS Net/One

5.7 INDUSTRIAL NETWORKING MAP/One

5.8 IBM Token Ring

5.9 PROTEON ProNET-10

5.10 INTEL Bitbus

5.11 DATAPOINT Arcnet

5.12 SPERRY Avionics Standard Communications Bus

5.1 SYTEK LocalNet

Summary

LocalNet is a general-purpose LAN using CSMA/CD techniques over broadband CATV cable. LocalNet supports RS-232-C asynchronous ports, but there is no circuit board version for direct insertion into microcomputers. The maximum data rate is a relatively slow 128 Kbps. An average RS-232-C interface costs about $500, plus $3,500 for the head-end remodulator, plus cabling.

General information

Sytek's LocalNet is a general-purpose, broadband LAN for use on CATV cabling. Founded in 1979 as a consulting company, Sytek acquired Network Resources Corporation in 1980 and, in 1984, agreed to build the IBM "PC Network" (not to be confused with the IBM token ring). Sytek is now affiliated with General Instrument Corporation which provides their CATV equipment. To date they have installed over 800 LocalNet systems.

Local/Net architecture

LocalNet is designed to provide process-to-process communication for end-user terminals and hosts, regardless of vendor. The bottom six layers of the ISO OSI model are provided by this system.

LocalNet uses broadband technology. The frequency domain can be divided into up to 60 separate channels, with each channel restricted to an aggregate data rate of 128 Kbps; a channel can serve up to 200 physical devices. At the physical layer, signals are modulated onto their respective carrier using frequency shift keying (FSK) techniques. The RF modem transmits to a head-end remodulator using a single mid-split cable. In the standard configuration, the "reverse" channel (to the head-end) occupies the 70-76 MHz spectrum while the "forward" channel (from the head-end) occupies the 226-232 MHz band. Each 6 MHz channel is divided into 20 300KHz data channels, and any one device may have up to three RF modems. Other frequency selections are available as extra-cost options.

At the data link layer, network access uses CSMA/CD techniques. The protocol is HDLC-like, provides automatic retransmission of collided packets, and uses CRC codes for error detection. The protocol is Ethernet-like, but apparently does not use binary exponential backoff and certainly does not run at Ethernet's 10 Mbps speed.

At the network layer, LocalNet provides a point-to-point datagram service from host to host. The transport layer provides virtual circuit connections among attached devices. Other transport layer services include sequence numbering, positive acknowledgements, and retransmission of flawed packets. The session layer controls network configuration and implements security, both through access control and data encryption.

The presentation layer provides three types of end-to-end services: (1) "Stream Transfer Service" uses virtual circuits for bulk file transfer between hosts; (2) "Virtual Terminal Service" handles bursty, asynchronous terminal-to-host transmissions; and (3) "Datagram Service" provides quick, low-delay data transfer with no flow control and no acknowledgements.

LocalNet hardware

LocalNet 20/100 Packet Communications Unit (PCU):

This is a microprocessor-based unit which connects two asynchronous RS-232-C devices at data rates from 300 bps to 19,200 bps. An optional circuit board supports synchronous transmission.

LocalNet 20/200 PCU:

A larger version of the 20/100, it supports eight asynchronous devices; it has no synchronous data option.

LocalNet 50/50 Head-end Translator:

The head-end translates the "reverse" channel frequencies to the "forward" channel frequencies. A redundant configuration is available, including a "Redundant Translator Auto Switch Unit" which automatically switches operation to the backup unit.

LocalNet 50/201 Packet Bridge:

This device permits interchannel bridging, i.e., devices tuned to one channel can communicate with other devices on separate channels.

Sytek 5101 Network Control Center:

This special computer contains a 32-bit microprocessor, 1 Mb main memory, 46 Mb hard disk, 1 Mb floppy disk, and 45 Mb cartridge tape. It includes 8 I/O ports with an option for eight more. It runs Unix and has battery-backup for its memory.

LocalNet 50/120 Statistical Monitor:

This device gathers statistical information. It attaches through an RS-232-C interface or through the 5101 Network Control Center. It monitors one channel at a time.

Sytek 6050 Network Translator Unit:

The 6050 connects a LocalNet with an IBM "Personal Connection" network.

Advantages and disadvantages

Advantages: creates a separate channel for each conversation; can implement security by access control and data encryption

Disadvantages: channel speed is 128 Kbps, much slower than other CSMA/CD networks (typically 10 Mbps); limited to broadband CATV media; supports only RS-232-C ports; no support for IBM PC or other microcomputers via circuit cards.

Pricing

HARDWARE	PRICE ($)
LocalNet 20/100 (2 RS-232 ports)	1,090
LocalNet 20/200 (8 RS-232 ports)	3,750
Synchronous data option	110
Modem channel groups:	
Single-cable mid-split group A	no charge
Other groups	50
Dual-cable mid-split group A	200
Other groups	250
LocalNet 50/50 Head-end Frequency Translator	3,500
Redundant Auto Switch Unit	3,500
Sytek 5101 Network Control Center, 8 ports	28,950
Sytek 5101 network Control Center, 16 ports	31,000
LocalNet 50/120 Statistical Monitor	1,895
LocalNet 50/10 Test Bed Cable Kit	100

5.2 BRIDGE COMMUNICATIONS Ethernet

Summary

Bridge Communications currently supplies an IEEE 802.3 compatible network using CSMA/CD access over baseband (Ethernet) coax, but with an Ethernet-to-broadband gateway which gives access to other networks. It is primarily designed for interconnecting large computers such as IBM, DEC VAX, and X.25 hosts, although there is a direct attachment for the IBM PC. The baseband network operates at 10 Mbps over Ethernet cable; the broadband gateway uses CATV RG-62 and RG-59 coax with a single cable and a mid-split or high-split configuration.

General information

Bridge Communications installed its first commercial network in January 1983. They are a second source for Ethernet products, hence they call this series the "Ethernet System Product Line (ESPL)." The basic product is indeed a baseband Ethernet, but it is integrated with a broadband product to give access to other major networks. Bridge supplies a large range of telecommunications products.

ESPL architecture

The Bridge version of Ethernet complies completely with all relevant specifications. It supports RS-232, RS-422, RS-449, IEEE 408, and V.35 device attachments. Direct attachments are available via a plug-in circuit card for the IBM PC and compatibles. Other direct attachments are available for IBM and DEC mainframes as well as X.25 hosts.

The topology is single-cable bus, using standard Ethernet cable for the baseband product. Signaling is 10 Mbps using CSMA/CD techniques. Integrated with the Ethernet is a broadband gateway which operates over CATV broadband coax such as RG-62 or RG-59. The broadband bus uses a standard head-end remodulator; RF modems in the device connections are frequency-agile over four channels, and the remodulator uses either a mid-split or high-split frequency allocation.

The goal of including the gateway is to encourage network communications among vendors other than Digital Equipment Corporation (who, along with Xerox and Intel, originally sponsored Ethernet), particularly IBM. For this reason direct network attachments are available for IBM mainframes, DEC mainframes, and IBM PCs.

The network supports up to ISO layer 5, session layer. Additional software in the hosts provides the remaining two layers.

Pricing

Not available.

5.3 INTEL Ethernet

Summary

The Intel Ethernet chip set (83586) *defines* Ethernet. The chip set is available by itself, or integrated into two board-level products with microprocessors and memories. Intel also supplies the software packages necessary to interface the boards into a Multibus-based host running either Xenix or RMX operating systems. Our experience with the 552 boards and the iNA 960 software is that they work, but their performance is slow and they are difficult to program and use.

General information

Intel makes the Ethernet chip set (82586) itself and several board-level products based on it. All of their board-level products are Multibus compatible (only) since that standard was originated by Intel. The two of interest are the iSBX 552 Ethernet Controller Board and the iSBX 186/51 "COMMputer". In cooperation with Intel, the 552 board-level product and its support software have been the object of intensive study and performance measurement at the University of Virginia. The 186/51 board contains an 80186 microprocessor and more memory; it can be used as a stand-alone communications computer. We have not yet evaluated the performance of the 186/51.

Ethernet product architecture

The 552 Ethernet controller packages the Ethernet chip set, an 80186 microprocessor, and 80K of memory on one Multibus card. This is designed for insertion into a host system such as the Intel 286/310 workstation which supports floppy disk, hard disk, RS-232 connections, etc. The board depends upon having a host processor to send it packets over the Multibus backplane. When packets arrive, they are transferred into the board's local "packet memory" and their location recorded in a linked list. Under firmware control (two ROMs for the 80186), packets are removed from the queue and transmitted on the Ethernet. One "hidden" cost in all baseband Ethernet systems is the transceiver and cable which connect the board-level product to the physical Ethernet. The transceiver is responsible for collision detection and similar duties and provides the coax tap into the Ethernet proper. Transceivers cost $250-$300, cables cost $10-60, and each attached device needs one of each.

The interaction between the host processor and the 552 board is simple in concept but difficult in its implementation. The concept is that the host generates packet frames, transfers them to the 552 board, and then delivery is automatic. While this is a correct abstraction, it requires substantial coding on the part of the user to attain the goal.

Two essential pieces of software are the iNA 960 (Intel Network Architecture) network interface and the Xenixnet or RMXnet package, depending upon the operating system in use. The iNA 960 is a generic product used for accessing multiple hardware interfaces (e.g., 552, 186/51) and so is available in a pre-configured version, iNA 961, just for the 552 boards. The iNA 961 package can be considered to be a subset of the iNA package with several options already chosen. Both iNA 960 and iNA 961 carry on an interpreted conversation (i.e., commands and responses which must be decoded at run time) with the 552 board. This is a good scheme to make iNA 960 generally usable by many products, but has a severe effect on ultimate system performance.

Within the host, iNA 961 provides the logical link control sublayer of the data link layer (i.e. the upper half of ISO layer two) and the transport layer. Under software control one can select three service types: transport virtual circuits, transport datagrams, and data link datagrams. Virtual circuits are the most secure; software sequence numbers each packet, transmits it, awaits an acknowledgement from the destination host, and retransmits it if no acknowledgement is received. Transport datagrams are generated at the transport layer (ISO layer 4) and are not acknowledged; data link datagrams are generated at the data link layer (ISO layer 2) and are not acknowledged. As expected, virtual circuits are slowest because of their inherent stop-and-wait protocol; data link datagrams are fastest because they are simply sent and forgotten after transmission. Transport datagrams are slower than data link datagrams (more software overhead) and appear to offer no compensating advantages.

One obscure point about the operation of the iNA 961 software is that delivering a framed message to the 552 board requires a "mailbox" style connection. Using an operating system call, a "communications request block" call is made which returns a buffer; the buffer is filled with the message; using arguments supplied in the call, the message is framed with source address, destination address, etc.; the buffer is released to the operating system for delivery to the 552 board (which appears to be accomplished using an interpretive command scheme, and by using word-serial communication over the Multibus rather than DMA); and the user then polls a "mailbox" until a signal is received indicating that the last packet has been accepted. "Accepted" does not mean transmitted; it means only that the packet was enqueued on the 552 board. One disquieting note is that, if the 552 board is given packets faster than it can accept them, the software hangs.

Intel system performance

Our performance studies of the 552 board indicate that while the electronics do operate at Ethernet speeds (i.e. the packets do emerge from the board at 10 Mbps), the hardware/software route from the host processor to the 552 board imposes a tremendous burden on the whole system. Thus, while a multi-station Ethernet might see traffic loadings in the megabits, no single node can come anywhere close to that loading.

The Figures in this section are taken from a research report which attempts to quantify the performance of Intel systems, connected via Ethernet, with all systems communicating through 552 controller boards and iNA 961 software.

Figure 29 shows the effect of protocol layer (transport virtual circuit, transport datagram, data link datagram) on mean delay. The upper graph shows that the time to complete the "communications request block" call was a uniform 5 ms regardless of packet size. The middle graph shows the time required to complete the request block call *and* get the signal from the mailbox that the packet has been accepted by the 552 board. Over a range of packet sizes from 32 to 1500 bytes, link datagrams took 7 to 10 ms; transport datagrams took 8 to 11 ms; transport virtual circuits took 20 to 28 ms. The bottom graph shows another important characteristic: end-to-end delay, the time elapsed from the moment the packet was enqueued in the memory of the transmitting host until the packet was successfully received in the memory of the receiving host. This measure covers all hardware, software, and transmission times. Over the range of packet sizes from 32 to 1500 bytes, end-to-end delays varied from 9 to 16 ms with data link datagrams; from 12 to 19 ms with transport datagrams; and from 13 to 20 ms with transport virtual circuits.

Figure 30 shows the distribution of times for these measures for a single packet size, in this case 32 bytes. Figure 31 shows the same information for a packet size of 1500 bytes. Note that in all cases, the time to complete the request block call is very near 5 ms (upper graphs). The time required to send a packet to the 552 board and receive the acceptance signal shows a substantial change as the packet length increases (middle graphs). Likewise, the end-to-end delays are strongly influenced by packet size (lower graphs).

Three observations emerge:

(1) Transport datagrams provide slower service but with no offsetting advantages, and so should never be used.

(2) Using Intel-supplied hardware and software, the maximum mean rate at which small (32-byte) packets could be generated was 1 packet every 7.41 ms, or 135 packets/sec, using data link datagrams, and was 1 packet every 19.49 ms, or 51 packets/sec, using transport virtual circuits.

(3) The maximum mean rate at which large (1500-byte) packets could be generated was 1 packet every 10.28 ms, or 97 packets/sec, for data link datagrams, and was 1 packet every 28.11 ms, or 35 packets/sec, using transport virtual circuits.

For this data we conclude that, regardless of how loaded the network is, no one node is going to handle data much faster than about 100 packets/sec.

Advantages and disadvantages

Advantages: complete hardware and software package available for purchase, including all seven ISO layers.

Disadvantages: available only for Multibus; 552 Ethernet controller board requires host processor to drive it.

Opinion: software is not user-friendly; considerable experience and insight needed to write the interface programs the first time.

Pricing

EQUIPMENT	PRICE ($)
Ethernet 552 controller board for Multibus	1,500
186/51 COMMputer for Multibus	3,000
Ethernet Transceiver (3COM)	300
Transceiver cable	50
iNA 960 Communications Software	3,325
iNA 961 Preconfigured Software	1,600

Figure 29.
Effect of Protocol Layer on Mean Delays

Figure 30.
Effect of Protocol Layer on Delay Distributions
(32 bytes/packet)

Figure 31.
Effect of Protocol Layer on Delay Distributions
(1500 bytes/packet)

5.4 INTERACTIVE SYSTEMS/3M LAN/I, LAN/II, LAN/PC

Summary

The LAN/I, LAN/II, and LAN/PC products actually provide four distinct local area network types. LAN/I is a 2.5 Mbps, token passing, broadband network which supports the Datapoint Arcnet protocol. LAN/II is available in two types. One version supports IEEE 802.3 CSMA/CD at 10 Mbps over baseband coax; the other supports the IEEE 802.4 token bus and GM MAP over broadband coax. Both operate at 10 Mbps. LAN/PC can serve up to 255 IBM PCs per channel pair on a 2.5 Mbps broadband bus, for a total of up to 1,275 PCs per network. Gateways are available to interconnect all four models.

General information

Interactive Systems installed the first commercial broadband data network in 1972. Since that time they have been trying to capitalize on their long experience with broadband equipment. The LAN/I product was introduced in May 1983 and utilizes the proprietary Datapoint token passing chip. LAN/PC was introduced in 1984. In June 1985 they introduced the LAN/II product in both its 802.3 and 802.4 models.

LAN/I

LAN/I is a broadband implementation of Datapoint's Arcnet token passing bus architecture. At the physical layer it uses 2.5 Mbps fixed-frequency RF modems with FSK modulation to transmit over one of five 6 MHz channels to a head-end remodulator; the remodulator shifts the signal to a higher frequency and retransmits it. At the data link layer, LAN/I uses IEEE 802.2 logical link control coupled with Arcnet channel access. LAN/I uses the same token passing chip as does the Datapoint RIM. Network membership is limited to 255 nodes because of an 8-bit address field. To the end user, LAN/I looks like a Datapoint Arcnet; however, due to LAN/I's broadband technology, it replaces the Datapoint star architecture ("hubs") with a single-cable, mid-split RF broadcast system.

LAN/I hardware

Network Interface Unit (NIU):

This NIU is available with 2, 4, 8, 16, or 32 asynchronous RS-232-C ports. It operates from 300 to 9,600 bps.

Network Monitor Unit (NMU):

This optional unit is an IBM XT running Concurrent PC-DOS. Supplied software allows simultaneous collection and display of network data. A network management function allows changes of configuration. The NMU attaches to the LAN/I through a NIU port.

LAN/II

LAN/II is third generation hardware and is available in two versions. One version complies with IEEE 802.3 using CSMA/CD on baseband (50-ohm) coax and operates at 10 Mbps. Network segments are limited to 500 meters.

The other version complies with both the IEEE 802.4 token bus standard and GM MAP protocol framing. The network medium is 0.5 inch solid shielded coax cable (not CATV cable); otherwise, components are CATV standard. LAN/II-802.4 operates at 10 Mbps.

LAN/II hardware

Network Interface Unit (NIU):

The LAN/II NIU is a VME bus product with local memory and network controller cards. It communicates externally over RS-232-C and RS-449 ports. An HDLC card is available for high speed communication using RS-232-C (2,400 to 19,200 bps) or RS-449 (2,400 to 460,800 bps). NIUs come in various configurations of from 4 to 16 RS-232-C ports and from 1 to 3 HDLC ports.

Head-end remodulator:

For LAN/II-802.4, this unit provides frequency conversion at the head of the cable.

Network Management Unit (NMU):

The NMU is a minicomputer which monitors network operation and attaches through a network bridge. It can list configuration status, collect and display statistics graphically, and conduct remote diagnostics.

Network Bridge Unit (NBU) and Wide Area Gateway:

These bridges and gateways are available: 802.4 broadband to 802.4 broadband bridge; 802.3 baseband to 802.4 broadband gateway; 802.3 baseband to 802.3 baseband bridge; and 802.4 broadband to LAN/I broadband gateway.

LAN/II software

Software for both versions of LAN/II is identical. It provides network functions at the local net, internet, transport, and application levels.

LAN/PC

The LAN/PC was developed in conjunction with Novell, Inc., and is a broadband LAN for linking IBM PCs, XTs, and ATs. It will connect almost any baseband system running Novel's NetWare software by using a broadband coax LAN.

Advantages and disadvantages:

Advantages: one version of LAN/II supports both IEEE 802.4 and GM MAP.

Disadvantages: LAN/I is slow at 2.5 Mbps; all network interfaces are through NIUs; there are no direct connections for microcomputers.

Pricing

EQUIPMENT	PRICE ($)
NIU, 2 ports	1,400
NIU, 4 ports	2,100
NIU, 8 ports	3,200
Network Monitor Unit (NMU)	5,700
LAN/I headend remodulator	2,500
modems are sold separately; various versions available; average price	900

5.5 CONCORD DATA SYSTEMS Token/Net

Summary

Concord Data Systems Token/Net was the first commercial product to comply with IEEE 802.4, implemented as a broadband token bus. It transmits at 5 Mbps over broadband CATV cable, using single cable mid-split or dual-cable topologies. It uses a frequency-agile head-end remodulator, covers six pairs of RF channels, and has a range of 25 miles. A special version of Token/Net is available which conforms to the GM MAP standard, version 2. The Token/Net MAP version has been successfully demonstrated at a major computer conference.

General information

The founder of Concord Data Systems, Mr. C. Kenneth Miller, has been an active member of the IEEE 802.4 standard committee. Not surprisingly, the Token/Net, introduced in March 1983, was the first LAN to fully conform to the (then emerging) 802.4 specification.

General Motors chose Token/Net as the backbone of GM MAP. Token/Net-MAP was demonstrated at the National Computer Conference in Las Vegas in July 1984 and successfully linked equipment from Allen-Bradley, DEC, Gould, HP, IBM, and Motorola (all of whom supply electronics to GM). Concord Data Systems has now signed OEM agreements with Allen-Bradley and Philips Telecommunications to supply components for the LANs marketed by those companies. Concord has recently agreed to collaborate with Gould AMI on the development of a Token/Net chip set.

The network connection is the Token/Net Interface Module (TIM). It connects devices to the network using RS-232 and RS-449 (HDLC) ports. As of this date Token/Net does NOT directly support an IBM PC. The TIM provides the packet assembler/disassembler needed by 802.4 as well as an RF modem, media access unit, control unit, and power supply. The RF modem operates on either mid-split or dual-cable CATV systems. The modem is frequency-agile over six channel pairs. The TIM also keeps network statistics on traffic, errors, and retries. A complete network also needs a frequency remodulator at the head-end; Concord supplies the HR-105 Head End Remodulator ($4,450) for this purpose. The HR-105 is frequency-agile over the same range as the TIM's RF modem. The HR-105 is also available in a dual-redundant configuration with a switch ($9,950); in case of failure, switchover is manual.

System software provides network control and selection of functions available to the user. There are two primary products, the Token/Net Network Control Computer (NCC) and the Token/Scope Network Analyzer. The NCC ($6,000) is based on an IBM PC AT and provides software with which the network manager can configure the network, select options, and display performance parameters. The Token/Scope Network Analyzer ($30,000) operates with a DEC VT100 display terminal and is the first protocol-level analyzer for 802.4 networks.

Architecture

The details of the IEEE 802.4 token bus standard are covered elsewhere and so are not repeated here. Each TIM provides RS-232 and/or RS-449 ports for connection to user devices. Network access control is fully handled by the TIM. Packets are transmitted over the RF modem at 5 Mbps, carried to the head-end by CATV cable, frequency translated to a higher "forward" frequency, and rebroadcast to all receivers. The user selects either a single-cable mid-split or a dual-cable topology. The RF modems and head-end remodulator are frequency-agile over six 6-MHz channel pairs.

Network hardware

There are four major pieces of hardware.

Token/Net Interface Module (TIM):

The TIM-200 provides two or four RS-232-C ports or one RS-232-C and one RS-449 synchronous port. The TIM-220 also provides two or four ports, but is field-expandable to 12 ports by adding one or two 4-port option boards. The unit is packaged in a tabletop enclosure and includes the ports, RF modem, access controller board, control unit, power supply, and medium interface. The TIM supports RS-232-C interfaces from 75 bps to 19,200 bps in asynchronous mode and from 1,200 bps to 19,200 bps in synchronous mode. The high-speed RS-449/442 synchronous interface operates from 9,600 bps to 230.4 Kbps. The RF modem provides software-controlled transmit power levels, channel selection, self-check, and an anti-jabber circuit.

HR-105 Head-End Remodulator:

The HR-105 is fully 802.4 broadband compliant. It demodulates the incoming signal and remodulates it at the higher outbound frequency to prevent noise accumulation on the channel. It is available in a redundant version with a manual switch.

Network Control Computer (NCC):

The NCC is an IBM PC AT with 512K memory, 1.2 Mb floppy, 20 Mb hard disk, and two asynchronous serial ports. The NCC uses the Unix operating system and a configuration data base library. Software is distributed on diskette which is then transferred to the hard disk.

The NCC provides the network manager with configuration capabilities. Associated with each TIM is a physical address, logical address, and configuration tables to define its parameters such as transmission speed and level of service. These tables are stored in non-volatile memory on the TIM and can be read and written by the NCC. The NCC provides a library of preconfigured profiles for common devices.

The NCC also provides real-time network monitoring. Information is displayed by layer (application, logical link control, medium access control, physical) and can be a high-level summary or very detailed. Diagnostics are supplied to test the whole network, any individual TIM, and any individual port. Each TIM has a resident Network Diagnostic Exerciser (NDE) to generate test frames.

Token/Scope Network Analyzer:

The Token/Scope is a protocol-level logic analyzer for the network. It contains a 5 Mbps frequency-agile RF modem, access unit, and a special control unit with Token/Scope firmware. It displays results on a user-supplied DEC VT-100 (or compatible) display.

Token/Scope monitors in real time and stores events in a 16K byte buffer memory. It operates in five modes:

(1) Command - set parameters;

(2) Monitor - real-time display of network traffic, error rate, token rotation time, and protocol events;

(3) Trace - record events based on user-defined trigger functions such as frame type, source and destination address, or condition of received frame. A command then switches to 'analyze' mode to display the data recorded.

(4) Analyze - chronological display of frames recorded in 'trace' mode

(5) Memory - displays raw frame data collected in 'trace' mode

GM MAP option

A special version of Token/Net complies with GM MAP version 2 in support of the Host Data Link (HDL). The HDL provides a high-speed RS-449 host-to-host interface operating at up to 128 Kbps and supports all four levels of priority in 802.4

Advantages and disadvantages

Advantages: fully 802.4 compliant; 5 Mbps over CATV cable with 25 mile range; GM MAP version available; company selected by GM as their vendor interface; redundant head-end available

Disadvantages: medium restricted to CATV cable; only one speed available; head-end switchover is manual; supports only RS-232 and RS-449 ports; no direct support (via plug-in cards) for IBM PC or any microcomputer.

Pricing

EQUIPMENT	PRICE ($)
Token/Net Interface Modules	
TIM-200 with four RS-232-C ports	3,485
TIM-200 with two RS-232-C and two RS-449 ports	4,135
TIM-220 with four RS-232-C ports, expandable to 12	3,950
TIM-220 with 12 RS-2323-C ports	5,940
MAP compatible TIM with HDL software	
TIM-200-MAP with one RS-232-C and one RS-449	4,735
TIM-220-MAP with one RS-232-C and one RS-449	5,200
CATV Equipment	
HR-105 Head-end Remodulator	4,450
Cable starter kit	350
Head-end redundancy switch	2,495
Redundant package: two HR-105s, one switch	9,950
Network Management	
Network Control Computer (NCC)	6,000
Token/Scope Network Analyzer	30,000
Starter Systems	
HR-105, cable, four TIM-220s	17,510
HR-105, cable, four TIM-200-MAPs	20,179

5.6 UNGERMANN-BASS Net/One

Summary

Net/One is a general purpose LAN which, by simple hardware replacement, supports baseband, broadband, and fiber optics. All Net/One systems are designed to be mutually compatible. "Ethernet Baseband Net/One" transmits at 10 Mbps, uses CSMA/CD, complies fully with the Ethernet specification, and operates on Ethernet 50-ohm coax, RG-58 A/U thin coax, or fiber optic media. "Broadband Net/One" transmits at 5 Mbps, uses CSMA/CD, employs fixed-frequency RF modems, and operates on industry-standard 75-ohm CATV coax. "Fiber Optic Net/One" transmits at 10 Mbps, uses CSMA/CD, is Ethernet compatible, and uses a fiber optic star coupler to simulate a baseband bus while actually using multiple point-to-point fiber optic connections. The basic NET/One product is the Network Interface Unit (NIU) which houses the intelligent network components and makes the physical connection to the network media; it is fed information through I/O ports. Other hardware such as repeaters and bridges is also available. Software runs on an IBM PC or PC-compatible.

General information

Net/One products are general purpose, designed to interface any vendor's equipment. This is a different strategy from, say, DEC or Xerox who make networks to interface their own products, and different still from, say, Corvus or Network Systems who target a specific market such a microcomputers. From the beginning, Ungermann-Bass sought to address the widest possible market, and they have achieved not only vendor independence but media independence as well.

Architecture

Net/One achieves independence through the modular architecture of its Network Interface Units (NIUs). Within an NIU, separate and replaceable circuit boards define the end-user devices (e.g. RS-232, IEEE-488) and the network medium interface (coax or fiber optics). Thus the difference between an Ethernet Baseband NIU and a Broadband NIU is the replacement of the Ethernet interface with an RF modem interface. Software support is provided by the Network Management Console (NMC) which is an IBM PC XT.

To compete with IBM's "PC Network," Ungermann-Bass offers the "Net/One Personal Connection." This single-board product plugs into an IBM PC expansion slot and provides a direct connection to the network medium without using an NIU.

All NIUs support the same five-layer architecture. The physical and data link layers are media-dependent so there are different circuit boards for the different network media: Ethernet 50-ohm coax, RG-58 A/U thin baseband coax, CATV-compatible 75-ohm broadband coax, and Ungermann-Bass-supplied fiber optics.

Ethernet Baseband Net/One

This system complies fully with the Ethernet specification which is very near to the IEEE 802.3 specification. The physical layer broadcasts digital, Manchester phase-encoded signals over the special Ethernet 50-ohm coax cable at 10 Mbps. The data link layer use CSMA/CD to arbitrate network access. Network repeaters are available to extend the range of the medium. This version interfaces both Ethernet cable and RG-58 A/U thin coax cable.

Optical Fiber Net/One

This system is an Ethernet-compatible baseband network using optical fiber rather than coax. Like the other baseband system above, it operates at 10 Mbps and uses CSMA/CD techniques. There is a wiring difference, however. Since baseband coax is inherently bidirectional and fiber optics are inherently unidirectional, a star coupler is used to make many fiber optic point-to-point connections look like an optical bus. Each transceiver has two fiber optic connections: one transmitting to the star coupler and one receiving from the star coupler. Whenever data is transmitted, it is sent to the star coupler on one inbound cable where it is then repeated on all outbound cables. Thus the star coupler looks like a "digital mirror" and provides the functionality of an electrical bus even though it is an optical device. The star coupler is functionally equivalent to the head-end remodulator on a broadband network.

There is a unique problem with fiber optic implementations of Ethernet — collision detection. Using baseband coax media, the signaling technique is Manchester encoding, which means that any data waveform produces constant energy in the channel. When two or more stations transmit simultaneously, the channel energy is increased proportionately and triggers and electrical collision detection circuit. Using fiber optics, collision detection is more difficult. The simple method of measuring the light energy in the channel is not always reliable. To assure reliability, an active device called a "collision reinforcer" is added to the star coupler. It actually compares the data being carried on the inbound fibers to determine if there is simultaneous transmission; if so a collision is signaled. This may be an elegant solution, but it is also expensive (price estimated at $10-15,000). The salesman's judgement was that the "collision reinforcer" was not needed for low load networks but was needed for highly loaded networks.

Broadband Net/One

This system broadcasts vestigial sideband amplitude-modulated (VSB-AM) RF signals over industry-standard 75-ohm CATV coax cable at a rate of 5 Mbps. It uses CSMA/CD for network access. It can be configured for mid-split, high-split, or dual-cable. The RF modem transmits on the "reverse" channel (TV channels 2 through 6) to a head-end frequency translator which shifts the signal to one of the "forward" channels (TV channels P through T). The RF modem is fixed-frequency and so utilizes only one pair of channels. By careful selection of modems, up to five separate 5 Mbps data channels can operate simultaneously, while still supporting voice and video in addition. Proven CATV technology provides good reliability and low error rates.

Net/One software

The Network and Transport layers of the ISO model are implemented in software. The Network layer provides the following services:

(1) Configure an NIU (select options)
(2) Download server (file transfer)
(3) Data link monitor (collects and displays statistics on NIU performance)
(4) Network debugger (permits dynamic alteration of user programs)
(5) Port reconfiguration (dynamic port changes)
(6) Broadcast/multicast (send message to all stations or to subgroup)

The Transport layer supports three service types: Virtual Circuit Service (VCS), Synchronous Communication Service (SCS), or Datagram Service (DS). VCS provides point-to-point communication between any two devices or programs; SCS provides a virtual circuit for synchronous devices; DS provides point-to-point, broadcast, or multicast communication between any set of intelligent devices.

Net/One hardware

NIU-130A (baseband), NIU-130B (broadband):

Compact, non-expandable, VLSI-based interface for two synchronous or asynchronous RS-232-C ports. Includes 128K (expandable to 512K) memory.

NIU-150A (baseband), NIU-150B (broadband):

Similar to NIU-130 but provides up to six ports chosen from among: RS-232-C; IEEE-488; parallel printer; serial V.35; 32-bit parallel. Includes 64K memory.

NIU-180A (baseband), NIU-180B (broadband):

Similar to NIU-130. Contains eight RS-232-C ports and 128K memory.

NIU-2A (baseband), NIU-2B (broadband):

Larger and expandable to accommodate up to 24 devices. Broadband model needs external modem. Supports up to three Application Processor Boards (user-supplied).

Network Repeater Units (NRU)

NRUs extend a baseband Ethernet segment to allow a maximum end-to-end station separation of 2,900 meters. NRUs connect segments separated by at most 100 meters; using inter-link repeaters this can be extended to 1 km.

Local and remote bridges

A local bridge provides store-and-forward service for geographically adjacent networks. A remote bridge connects segments via some high-speed telecommunications link (9,600 to 504K bps). An X.25 gateway is also available.

Network Management Console (NMC)

The NMC is an IBM PC XT with 512K memory and a 10 MB hard disk. It runs the network operating system and connects to the network via a Personal NIU (see next section).

Net/One Personal Connection

This is a family of hardware and software products. One hardware product is the *Personal Network Interface Unit (Personal NIU)*, a single board which plugs into the expansion slot of an IBM PC. It has an 80186 processor and 128K memory to off-load the PC's processor. The other hardware product is the *Personal Network Interface Controller (Personal NIC)*, which is like the Personal NIU except it has no processor. Instead the PC's processor supplies the computing power.

The software products provide network users with a disk server, printer server, electronic mail server, and a CROSSTALK terminal emulation package. Both products require operating system DOS version 3.1 (not the more common version 2.x).

Future plans

Ungermann-Bass has announced a token ring product which will be IEEE 802.5 compatible. Its expected availability is April 1986.

Advantages and disadvantages

Advantages: modular hardware and software; supports baseband coax, broadband coax, and fiber optics; supports mixed media; supports IBM PC; supports RS-232-C ports in the NIU.

Disadvantages: optical fiber Ethernet uses a centralized star coupler; heavily loaded optical networks need a $15,000 "collision reinforcement" unit.

Opinion: network is so flexible that it takes a sophisticated end-user to install and maintain it.

Pricing

EQUIPMENT	PRICE ($)
2-port NIU	1,400
8-port NIU	2,950
IBM PC Personal NIU (onboard processor)	1,095
IBM PC Personal NIC (no processor)	595
For baseband systems:	
Transceiver (one per connection)	275
Transceiver cable (one per connection)	65
For broadband systems:	
RF modem (one per connection)	650
Head-end remodulator (one per network)	3,500
For fiber optic systems:	
Star coupler (one per network) (estimate)	2-3,000
Optical transceiver (one per connection)	1,500
Collision reinforcer (optional, one per network)	10-15,000

5.7 INDUSTRIAL NETWORKING MAP/One

Summary

The MAP/One product is a joint venture of Ungermann-Bass and General Electric. It is a broadband 802.4 token bus which also supports GM MAP version 2.1. It is intended for factory automation and is an attempt to compete with Concord Data Systems' Token/Net. It communicates at 10 Mbps using a single-cable mid-split topology.

General information

General Motors declared in 1982 that, beginning in 1988, it would no longer buy industrial control equipment from vendors whose equipment could not communicate with other vendors' equipment. The problem at GM was that every new plant or assembly line they built soon became a captive market of whatever vendor was initially selected to provide the control system. From the point of view of the successful vendors (Allen-Bradley, Gould-Modicon, Texas Instruments, General Electric, etc.) this was good news; from the point of view of General Motors, it was unacceptable.

The manufacturers of industrial controllers were simply unable to agree on a standard for communications — each vendor had an investment in its proprietary scheme which it was unwilling to abandon. So General Motors took the initiative away from the manufacturers and developed its own set of communications standards: MAP (Manufacturing Automation Protocol). It was the first successful case of the end-user setting the standards instead of the vendors, and of course was successful only because of the enormous economic influence of GM.

MAP defines the lowest two levels of its seven-layer protocol to be the 10 Mbps broadband version of the IEEE 802.4 token bus. Concord Data Systems was chosen as the supplier of that product and given the responsibility of helping vendors of automation equipment interface to it. In the LAN world, this greatly increased the credibility and influence of Concord Data Systems.

To counter that influence, Ungermann-Bass (who lost their bid to be the main network contractor) and General Electric (who needed an industrial network of their own but didn't have one) agreed to form a new company, Industrial Networking, owned 60% by Ungermann-Bass and 40% by General Electric, whose responsibility would be to adapt an existing networking product to the needs of the GM MAP standard.

The result is MAP/One. It is a hardware and software adaptation of Ungermann-Bass' broadband Net/One product which has been very successful in the non-industrial marketplace (the University of Virginia uses the CSMA/CD version of Net/One as its backbone network). Thus, the performance of MAP/One could be expected to be similar to that of broadband Net/One, but with a totally different user interface as mandated by the MAP standard.

MAP/One architecture

MAP/One is a broadband 10 Mbps token passing bus conforming to IEEE 802.4 and GM MAP version 2.1. It uses common CATV cable and a single-cable mid-split topology. RF modems and the head-end remodulator are fixed-frequency, but three separate channel pairs are available. The frequencies themselves are 802.4 standard. Signal modulation uses a combination of amplitude modulation and phase shift keying (AM/PSK). Error control is provided by a cyclic redundancy code. The hardware provides services up through ISO layer four, transport layer.

One interesting note is that the network interface function is provided by a gate-array chip set; this is the first use of gate-arrays in support of MAP. Since the product must be prepared to function in environmentally poor conditions, it is designed to operate at ambient temperatures of up to 65 degrees C.

Another note is that this is a very new product with very limited field experience. The first commercial delivery was made in August 1985. Information on how many commercial networks have been installed was not forthcoming from the vendor.

MAP/One hardware

Microcomputer interface:

MAP/One hardware is similar to the broadband Net/One product line. Direct interfaces are supplied for the Multibus, IBM PC bus, and VMEbus. These interfaces consist of two boards, an RF modem and a digital controller, connected via ribbon cable. They occupy two slots in the backplane; the IBM version costs $3,450.

Head-end remodulator:

This is the standard frequency translator for broadband systems. The cost is $6,000.

Network Management Console:

This is an IBM XT with MAP interface board and provides network management software; it costs $25,000.

Network software

Message and file transfer:

Two programs will soon be available. MMFS provides message passing services among nodes; FTAM performs file transfer. Price is estimated at $500 each for binary copies.

Network monitor software option:

This is add-on software to the Network Management Console. It takes statistics kept by the Console and produces histograms and pie charts of network bandwidth utilization and individual node utilization; it is priced at $18,000.

Traffic analyzer:

This software package provides bit-level statistics for large (100-node) networks; it costs $40,000.

Future products

MAP is a young, changing standard. One move underway is to define a carrier-band option to replace the normal RF signaling technique in the hope of reducing cost at some sacrifice of performance. One possibility is to define a sub-net which would operate at 5 Mbps, over limited distance, with a maximum of 25 nodes. It would be media-independent and so could support fiber optics which are currently unsupported. It could connect to the backbone broadband via a bridge.

Another subject under discussion is EPA — the Enhanced Performance Architecture, also called Proway, also called MicroMAP. It might be a 3-layer model supporting physical and data link layers and then an interface layer leading directly to the application layer. It might support priorities and provide an "instant" acknowledgement option.

These additions to the MAP standard are currently under discussion and their evolution into products could be reasonably expected to take another year.

Advantages and disadvantages

Advantages: implements both IEEE 802.4 and GM MAP; lower level functionality of MAP implemented in gate-array technology for increased speed; supports IBM PC; based on a mature broadband product, the Ungermann-Bass Net/One; provides 10 Mbps transmission rate and token passing determinism; specifically designed for harsh factory environments; tolerates temperatures to 65 degrees C.

Disadvantages: it is a new product and has a limited track record.

Opinion: The support products such as the network management console and the network traffic monitor are overpriced.

Pricing

EQUIPMENT	PRICE ($)
IBM PC interface (one per node)	3,450
Software license (one per node)	300
Head-end remodulator	6,000
Network Management Console	25,000
Message transfer program (estimate)	500
File transfer program (estimate)	500
Network monitor software additions	25,000
Traffic monitor	40,000

5.8 IBM Token Ring

Summary

The IBM token ring network is a 4 Mbps LAN supporting up to 260 IBM PCs when operating over shielded twisted pair and up to 70 PCs using unshielded twisted pair. The token ring complies with IEEE 802.2 logical link control and IEEE 802.5 token ring network access specifications. It is designed as an "open" architecture to encourage third-party vendors to build compatible equipment. Network interface is provided by a eight-port Multistation Access Unit and by single circuit boards which fit an expansion slot in an IBM PC. Software support is provided by the PC and includes SNA emulation, an asynchronous gateway, a bridge to the IBM PC Network, a Series/1 gateway, a print server, and NETBIOS functions. An Advanced Program-to-Program Communications (APPC) package supports peer-to-peer communications between PCs and SNA-based hosts. As of this date, the token ring hardware supports only PCs and, via the Multistation Access Unit, a few other intelligent devices.

General information

The world has waited a long time for the announcement of the IBM token ring network. A teaser was supplied in May 1984 with the announcement of the IBM Cabling System, a structured wiring scheme which supported five types of media. Then in August 1984 IBM announced its PC Network — but it did not use the IBM Cabling System! Instead, it was a broadband product (actually built by Sytek) which used CATV coaxial cable. Finally, on October 15, 1985, IBM unveiled the token ring network. It is generally believed that the IBM token ring, like the IBM PC, will become a *de facto* standard.

The IBM token ring operates over the IBM Cabling System (see next section) at speeds up to 4 Mbps. It supports up to 260 devices on Cabling System type 1, 2, and 6 media (basically shielded twisted pair) and up to 70 devices on type 3 media (telephone-style unshielded twisted pair). Note: as of this date, the token ring does NOT operate on type 5 media — fiber optics). The token ring conforms to IEEE 802.2 and 802.5, and also to the European Computer Manufacturers Association (ECMA, see section 3.5 on Standards) standard 89 for baseband, token ring LANs.

The basic components of a token ring network are: (1) the PC Adapter which plugs into the expansion slot of an IBM PC; (2) the Multistation Access Unit, an eight-station interface; and (3) the IBM Cabling System. As with the IBM PC, a key part of IBM's strategy is to have an "open" architecture in which specifications are public and other vendors are encouraged to build compatible products. The technical interface specifications are crucially important, and IBM has already made them available to selected independent vendors. They are to become generally available in first quarter 1986. The semiconductor components were developed by Texas Instruments and form a five-chip set: three processor chips and two interface chips.

IBM Cabling System

IBM announced the Cabling System in 1984 in an attempt to limit the intrusion of other vendors on IBM's "turf". It was hoped that by announcing an acceptable cabling plant, potential users would begin wiring applications in anticipation of the long-expected announcement of the token ring products.

The Cabling System defines the following types of media:

Type 1, Type 1 Plenum, Type 1 Outdoor:

shielded, two-pair twisted wire for data; 22 AWG solid conductors; type 1 and type 1 plenum are enclosed in braided cable shield; type 1 outdoor is enclosed in a corrugated metallic cable shield

Type 2, Type 2 Plenum:

shielded, six-pair twisted wire for voice and data; two pairs are 22 AWG solid conductors in a braided metal cable shield and are meant for data; four pairs are unshielded 22 AWG solid conductors and are meant for telephone

Type 3:

unshielded, two-pair twisted wire for voice and data (i.e. telephone cable); 22/24 AWG solid copper wires; minimum two twists per foot; NOTE that this type cable is not manufactured or sold by IBM

(There is no Type 4.)

Type 5:

fiber optic cable; two 100/140 micron optical fibers

Type 6:

shielded, two-pair twisted wire; 26 AWG; stranded; flexible, intended for patch cable

Prices:

CABLE TYPE	PRICE PER FOOT
Type 1	$ 0.40
Type 1 Plenum	$ 2.60
Type 1 Outdoor	$ 0.70
Type 2	$ 0.65
Type 2 Plenum	$ 3.00
Type 5	$ 3.00
Type 6	$ 0.45

Prices for IBM Cabling System Media

Token ring architecture

The token ring adheres to the specifications of IEEE 802.5; thus we will not repeat the details of the protocol.

An important note is that connection to the network follows the concept of the "wiring center" presented previously in section 3.2.3.1. This is used to enhance reliability. If a device port is not enabled through hardware signals or software handshaking, that port does not become an active member of the ring (in the words of section 3.2.3.1, its "bypass" is not energized so the port is "shorted out" to provide connectivity across it). The IBM phrase for "wiring center" is "wiring closet," but it means the same thing.

As one would expect with twisted pair media, there are limitations on the cable lengths between wiring centers and between the wiring center and its connected hosts. Type 1, Type 2, and Type 6 media support a maximum of 260 devices subject to these limitations: maximum distance between wiring centers—200 meters; maximum distance from host to wiring center—100 meters. Type 3 media (unshielded twisted pair) is subject to these limitations: maximum number of devices—70; maximum distance between wiring centers—120 meters; maximum distance from host to wiring center—45 meters. Type 5 media (fiber optics) is not yet supported.

Token ring hardware

At present only two pieces of interface hardware are available (besides the cabling):

Token Ring Network PC Adapter:

This card plugs into a full-size expansion slot of an IBM PC, XT, AT, or Portable. It contains a microprocessor which operates from microcode. The adapter implements IEEE 802.2 and 802.5 at 4 Mbps. On-board diagnostics verify proper operation and check out the cabling to the access unit. Permanent errors (such as loss of signal) generate a notification and initiate automatic network recovery. Transient errors (such as bit errors) are detected and reported through the ring diagnostic program. Each adapter includes a software diskette with two programs. The adapter diagnostic program is used to verify correct operation of the adapter; the ring diagnostic program is used as an aid in problem determination. Permanent and transient errors are reported, including information on the probable source of error.

Multistation Access Unit:

This unit provides connections for up to eight devices. Cables from the devices are plugged into the access unit; access units can be connected to other access units to form larger wiring centers. The access unit will automatically bypass a device which fails to provide the proper signal. The unit is designed for installation in a 19-inch rack or in a component housing.

Token ring software

Besides the two diagnostic programs mentioned earlier, there are two main programming interfaces, both of which run on an IBM PC under DOS.

Advanced Program-to-Program Communication for the IBM PC (APPC/PC):

APPC/PC is an application 8rogramming interface to the network which supports program-to-program communication over the token ring and also over synchronous data link control (SDLC) links. Program size is about 200K.

Token ring NETBIOS:

The Network Basic Input/Output Subsystem (NETBIOS) provides user-callable functions (operating system procedures). Application program requests for NETBIOS functions are translated into token ring protocol requests which are then passed to the network PC adapter for execution. The NETBIOS functions and the APPC/PC program can be used concurrently. Program size is about 50K.

Bridges and gateways

The "IBM Token Ring Network/IBM PC Network Interconnect Program" provides a connection (gateway) between the baseband token ring and the broadband IBM PC Network. The NETBIOS is required; the result is the ability to remotely access files, data, or programs on another system.

Although independent vendors have announced their intention to build gateways to other LANs, none are yet available.

Advantages and disadvantages

Advantages: supports an "open" architecture to attract more products; supports the IBM PC; can operate on inexpensive twisted pair.

Disadvantages: it is a very new product, so there is little experience with either its hardware or software; cable length limitations using twisted pair are severe; ring speed of 4 Mbps is not very fast; efficiency of software is unknown; inadequate support for devices other than IBM PC; fiber optics not yet supported

Opinion: it is an advantage to be supported by IBM; of the star, bus, and ring topologies, the ring is inherently the most efficient; connection via twisted pair will probably provide inadequate noise immunity for shipboard applications.

Pricing

PRODUCT	PRICE($)	NOTES
HARDWARE		
IBM Token Ring Network PC Adapter	695	one per PC
IBM Token Ring Network PC Adapter Cable	35	one per PC
IBM 8-port Multistation Access Unit	660	for RS-232 devices
Component Housing for above	99	optional
SOFTWARE		
APPC/PC Program	150	one-time charge
NETBIOS Program	35	one-time charge
Token Ring/PC Network Interconnect	495	one-time charge
Asynchronous Communications Server	495	one-time charge

CABLING		
Type 1	$ 0.40	all prices per foot
Type 1 Plenum	$ 2.60	
Type 1 Outdoor	0.70	
Type 2	$ 0.65	
Type 2 Plenum	$ 3.00	
Type 5	$ 3.00	optical fiber cable not yet available
Type 6	$ 0.45	

5.9 PROTEON ProNET-10

Summary

The ProNET-10 is a family of products which provide host-to-host communication at 10 Mbps over shielded twisted pair, fiber optics, or infrared, using a star-ring configuration. The ring provides a backbone between wiring centers; hosts are then connected to the wiring centers. Terminals and RS-232-C devices which run TCP/IP are connected to a wiring center through a Terminal Interface Unit. The network access protocol is currently proprietary for the 10 Mbps product; a slower 4 Mbps version (ProNET-4) conforms to IEEE 802.5 and is also IBM compatible. Media connections are external to wiring centers so that mixed media are permitted on the backbone. A redundant fiber optic cable supporting automatic reversible rings is available for interconnecting wiring centers. Host interface cards are available for UNIBUS, Multibus, Q-BUS, Universal bus, and IBM PC bus. A higher-speed 80 Mbps version is available in the ProNET-80 product line.

General information

ProNET-10 was developed in conjunction with the Laboratory for Computer Science at MIT. Proteon chairman Howard Selman said he had always intended to use the same architecture as IBM (the star-ring topology) and the same network access method (token passing), but that he was confident that he could come to market sooner than IBM. He was right; Proteon's first delivery was June 1981; IBM's announcement of their token ring was October 1985. Of course, IBM waited until the direction of IEEE 802.5 was clear before making their announcement; Proteon was able to beat IBM to the marketplace by (at least initially) not adhering to any standard. In response to market pressures, Proteon developed the 4 Mbps ProNET-4 which is IEEE 802.5 compliant and IBM compatible.

ProNET-10 takes advantage of the inherent simplicity and high performance of a ring topology and protocol while simultaneously addressing the ring's inherent weak spot with regard to reliability: ring connectivity. By using wiring centers to interface all hosts, terminals, and devices, the chance of any one device failure bringing down the whole ring is greatly reduced. By offering single-cable fiber optic wiring as a replacement for their normal shielded twisted pair cable, the overall integrity of communications is enhanced. By offering a dual-cable, reversible ring configuration the reliability of the entire network is greatly enhanced.

ProNET-10 architecture

ProNET-10 uses a ring topology for a backbone which connects wiring centers; wiring centers use a star topology for connecting hosts and terminal interface units. Each connection in the wiring center has a bypass circuit (relay) associated with it so that failures of remote hosts or terminals will not disrupt ring connectivity. The wiring center is passive; it draws no power from the ring. The bypass circuits use normally-closed relays to ensure that unused or defective ports are not a part of the physical ring. If a host is connected via twisted pair, the physical connection (through a 15-pin connector) and supply of power from the host disengages the bypass circuit and allows the host to join. If the host connection uses fiber optics (an option), then an exchange of data messages is necessary before the bypass circuit is disengaged. The overall reliability of the wiring center concept has been shown in the literature to be high.

PROTEON ProNET-10

The host interface is a plug-in circuit card; interfaces are currently available for UNIBUS (DEC), Q-BUS (DEC), Multibus (Intel), Universal bus (TI), and IBM PC bus. The connection from host to wiring center can be shielded twisted pair (standard) or fiber optic (optional).

Terminals which run TCP/IP are connected via RS-232-C ports to a Terminal Interface Unit (TIU) which in turn is connected to a wiring center. There is no direct RS-232-C-style connection to the network because that would violate the integrity of the wiring scheme. Terminals connect to the TIU using standard 15-pin connectors; the TIU connects to the wiring center using shielded twisted pair, fiber optics, or infrared. Terminals communicate over the network using the industry standard TCP/IP protocol.

Network size is limited to 255 nodes; each connection in the wiring center has a DIP switch or equivalent for setting the nodal address. Using shielded twisted pair, the maximum wire-center-to-wire-center distance is 480 meters; the maximum host-to-wire-center distance is 480 meters. By switching to fiber optics, both distance limitations are raised to 2.5 kilometers.

Network access is controlled by the unique token circulating on the ring. Although the current version is not IEEE 802.5 standard, it is similar. Upgrade to IEEE 802.5 in the future should be straightforward. The current version shows these major differences from the IEEE 802.5 version: (1) the physical address space is small, 255 nodes; (2) priority is not supported; and (3) the message is error-checked by a parity bit, rather than the much more powerful 32-bit cyclic redundancy code (CRC-32) specified in IEEE 802.2.

Host interface

The PC interface occupies one slot in a PC, XT, AT, or compatibles. The interface provides both the communications interface controller and the network control interface. The communications interface performs the communications processing for the host, including packet sequencing, buffering, and controlling; this places the minimal computational burden on the host processor. The network control interface handles all the network transmission functions with no host intervention; these functions include serial-parallel conversion, bit stuffing, error timeouts, and lost token regeneration. The network control interface also handles activation of the bypass relays in the wiring center; self testing is also provided under software control. If the PC supports DMA operations through an 8237A DMA controller, then the interface board can take advantage of that, passing packets to and from the host via DMA rather than word-serially over the backplane bus.

To send data, the host prepares a "Proteon packet" consisting of a beginning-of-message flag, destination address, source address, data bytes, end-of-message flag, and 'refused' bit, followed by a beginning-of-message flag and new token. When a packet is transmitted, the 'refused' bit in the tail of the message is set; if the message is copied successfully by the receiver, it clears the 'refused' bit. When the packet returns to the sender, it can tell whether the packet was correctly received. If not, the sender must pass the token, but it knows that the previous message was not delivered. Clearly, this 'refused' bit algorithm is a forerunner of (and the opposite sense of) the 'copied' bit in the frame status byte of the IEEE 802.5 standard.

Overall packet length is limited to 2046 bytes (same size as the hardware buffer on the interface cards). The host may send only one data frame for each token received, which increases the token circulation frequency (and reduces the total network capacity usable by data). When not sending packets or receiving its own packet, the host interface acts as a repeater; however, it decodes and recodes the data as it passes through the interface to maintain signal strength and refresh the clocking information.

SDLC data link protocol

At the physical layer, the modem uses synchronous, phase-locked-loop, Manchester-encoded signaling; at the data link layer, it uses SDLC (synchronous data link control). SDLC was chosen for the following reasons. Since data is encoded using Manchester encoding, rather than differential Manchester encoding, each bit time encodes one of only two values, "0" and "1" (differential Manchester provides four values). Thus a token (called 'flag' in SDLC) consists of a particular 8-bit string of 0s and 1s, which happens to be 01111110. But this byte could appear in ordinary data, which would confuse the protocol. So SDLC does the following: the bit stream being transmitted is continuously monitored; whenever five 1s in a row are detected, a 0 is inserted after them. Upon receipt of the proper command, the SDLC chip can generate the code for a token (01111110), but this code can never be generated from the user's data stream. At the receiver the input stream is also continuously monitored. Whenever the receiver sees five 1s in a row, it looks at the next bit. If it is a 0, it is removed; if it is a 1, a token has been received. This scheme, which was originated by IBM, has now been adopted by IEEE and ISO under the name HDLC (high-level data link control). Note that IEEE 802.5 eliminates this token recognition problem entirely.

ProNET-10 hardware

Host interfaces:

UNIBUS, Multibus, Q-bus, and Universal "Host Interface Systems" consist of two boards: the host-specific board and the network interface board. Only the IBM PC interface is a single plug-in card.

Wire center:

Wire centers contain the robust ring connection and bypass relay circuits. The are available with 4, 8, 12, and 16 external connections. The wiring from the host to the wire center can be a shielded twisted pair, fiber optic, or infrared link. The cabling between wire centers can likewise be a shielded twisted pair, fiber optic, or infrared link.

Network monitor:

The network monitor is software for an IBM PC and a hardware interface card which replaces the normal host interface card. The network monitor provides the following functions:

(1) network activity report — shows number of packets on the network for the current second, minute, hour, and day, as well as for the previous and the busiest second, minute, hour, and day.

(2) network loading — displays network loading, as a percentage of maximum capacity (10 Mbps), for current, previous, and busiest second, minute, hour, and day.

(3) network errors — counts all transmissions which do not conform to the data link standard (e.g. short packets, control characters, ring reinitializations); used as a

maintenance tool to detect deteriorating conditions in advance of actual failure.

(4) active host display -- keeps track of source addresses for all packets passing through the network along with the time of the most recent packet from each host; helps identify host-level problems.

Infrared communications link:

This module can replace twisted pair or fiber optics in a wire-center-to-wire-center link or in a host-to-wire-center link. It spans 350 meters in clear weather, 100 meters under all conditions.

Gateways:

ProNET-Linkways are a family of gateways. They are currently available in configurations to support ProNET-10, ProNET-80, Ethernet, Arpanet, and point-to-point telephone lines. They support two protocols now, IP (Internet Protocol) and XNS (Xerox Network System), with plans to support SNA, DECnet, IEEE 802.4, GM MAP, and others. A separate product, the SNA Gateway, provides a single plug-in card running SNA and 3270 communications links.

ProNET-10 software

TCP/IP network software:

TCP/IP means Transmission Control Protocol/Internet Protocol. IP is a datagram protocol, and TCP is a stream protocol layered on top of IP. It was developed by ARPA for the Defense Communications Agency and became a Department of Defense standard in 1983. Operating system support is available for UNIX, MS-DOS, and PC-DOS.

NetWare:

The NetWare operating system allows IBM PCs to share one or multiple file servers. This would be useful in a network of diskless PCs which wish to share a centralized hard disk on the network.

ProNET-80

ProNET-80 is an 80 Mbps version of ProNET-10. Almost every aspect discussed above is the same, except: (1) there is a wiring length limitation of 100 meters between wiring centers and between hosts and wiring centers; fiber optics will extend this length to 1 km; (2) host interfaces are available only for Multibus and Q-bus (none for IBM PC); (3) the transmission protocol uses a 4-out-of-6 error correcting code so that simple 1-bit errors can not only be detected but corrected.

Advantages and disadvantages:

Advantages: cabling from hosts to wire centers and between wire centers can be shielded twisted pair, fiber optics, infrared, or any combination; fiber optic wiring is available for a redundant link (both links operating in the same direction) and for a reversable dual link; the host interface is available as a single plug-in card for the IBM PC; network monitoring and diagnostic hardware/software is available; if still greater backbone speed is needed, an 80 Mbps version is available in the ProNET-80 family.

Disadvantages: using the wiring center increases the per-device interface cost; RS-232-C devices can not attach directly to the network, but must use a device concentrator; the current version of ProNET-10 is not IEEE 802.5 compatible.

Opinion: the ring topology and protocol are the most efficient; the wiring center concept solves the reliability problem of ring connectivity; using the wiring center the backbone connections can be semi-permanent, which should be an advantage for vibration-sensitive connections such as fiber optics; the 10 Mbps speed should be 90% usable by the hosts.

Pricing

EQUIPMENT	PRICE($)
IBM PC interface card	799
IBM PC network monitor interface	799
4-node wire center, no expansion	295
4-node wire center, expandable	330
8-node wire center	630
12-node wire center	950
shielded twisted pair cable, 3 meters	50
shielded twisted pair cable, 10 meters	65
shielded twisted pair cable, 100 meters	450
shielded twisted pair cable, 300 meters	1,180
fiber optic link, host-to-wire center	2,300
fiber optic link, wire center-to-wire center	2,300
redundant fiber optic link, host-to-wire center	2,900
redundant fiber optic link, wire center-to-wire center	2,900
Terminal Interface Unit, 8-port	9,700
Terminal Interface Unit, 16-port	11,000
Basic license for software (one per customer)	250
ProNET-10 diagnostic software for IBM PC	100
TCP/IP software for TIU	1,500
NetWare for MS-DOS	1,495
Manuals (each)	50

5.10 INTEL Bitbus

Summary

Bitbus is a master/slave communications protocol implemented on a bus architecture using serial SDLC communications at three different rates (62.5 Kbps, 375 Kbps, 2.4 Mbps) depending upon bus segment length. Bus media is twisted pair and supports a multi-drop arrangement (branching tree topology). Each segment supports at most 32 nodes. The highest speed is achieved with 4-wire twisted pair over 30 meters; the intermediate speed with 2-wire twisted pair over 300 meters; the slowest with 2-wire cable over 1,200 meters. Repeaters can be used to interconnect segments. At most 10 repeaters can be in use to achieve the network limit of 254 nodes distributed over 10 kilometers. The master node polls attached devices ("slaves") to solicit transmissions. Communications are restricted to 18 bytes maximum, of which 4 are overhead; only one message may be transmitted per poll. There is no provision for redundancy of the master controller or media.

General information

Intel Corporation has long been a well-respected leader in integrated circuit design and fabrication. Although first known for their dynamic memories, they are now known as the producer of the world's first microprocessor (Intel 4004), the best-known microprocessor (Intel 8080), the microprocessor for the IBM PC (8088), the Ethernet chip set (82586), and a series of higher-performance microprocessors (8086, 286, 386). More recently they have turned their corporate attentions to data communications, and have expanded their product line of RS-232 interfaces such as USARTs to include serial communications bus interfaces. The Bitbus utilizes an 8044 microcontroller, which itself is a combination of a 12-MHz 8051 microcontroller with a serial interface unit.

Recognizing that local area networks (such as Ethernet) are an expensive solution to the problem of distributed communications, Intel has devised a simple solution for moving short (i.e. control type) messages over short distances at slow-to-moderate speeds. Intended for factory automation and machine control, the Bitbus replaces a traditional star architecture, with separate wires connecting each controlled device to the device controller, with a serial bus architecture. The primary advantage of the architecture is the cost savings which accrue from reduced wiring. A major disadvantage introduced is the protocol dependence upon a master node.

Bitbus architecture

Bitbus substitutes a master node on a bus architecture for a central star switch with independent cabling to each connected device. The difference is real in terms of the efficiency of the wiring, but illusory in terms of the control logic.

At the time of network generation, each node is assigned a unique 8-bit address. The master node, under the control of the iRMX 51 real-time operating system (or similar), is responsible for determining what network addresses exist and establishing a polling order for them. Subsequently, the master node sends polling messages to all slaves in turn, after which each slave device is permitted to respond with a single message. Remote (slave) devices are never permitted to initiate a transmission; they must respond to polls.

Message formats are severely limited in anticipation of a factory control environment. Maximum message length is 18 bytes, of which 4 bytes are overhead (message length, flag byte, source address, and destination address). The remaining field of up to 14 bytes may contain arbitrary data.

Source and destination addresses are 8 bits each. One suggestion for managing the 8-bit name space of the network is to allocate the 256 potential addresses into 16 groups of 16 addresses. In the first fifteen groups, allocate 15 device addresses and one group address. In the sixteenth group, allocate fourteen device addresses, one group address, and a common "broadcast" address for the whole network. This scheme results in a maximum of 239 nodes. If the address space is considered unique, less one address for common broadcast and one address for subgroup broadcast, then the maximum membership is 254.

Recognizing that RS-232-C was a connection standard, not a protocol, Bitbus needed something more elaborate. The decision was to use RS-485 at the physical layer and pseudo-SDLC at the data link layer. RS-485 is a superset of the more popular RS-422, and is expanded to provide multi-drop support over twisted pair.

Data link control is provided by using a subset of the SDLC standard. Unfortunately, some of the statements made in support of Intel's implementation are misleading. One such statement credits the 16-bit CRC code in SDLC with providing error detection capabilities so robust that the probability of passing a message in error is reduced to 10^{-10}. While this may be true about the CRC code, this cannot be true about the end-to-end error rate. Using (unshielded) twisted pair in an electrically noisy environment will significantly increase the error rate. The errors will be detected and the messages rejected, but overall network throughput will be substantially reduced.

Because of signal attenuation in the twisted pair media, bus lengths are limited. No segment may exceed 32 attached nodes or 1,200 meters in length. Segments may be interconnected using repeaters; no more than 10 repeaters may be used, and the longest path may not exceed 10 kilometers.

Utility of Bitbus

Bitbus was designed for a special environment: factory control. Its primary advantage is the economy of replacing multiple individual wires with a single bus. Its master/slave design marks it as being appropriate for data collection, where intelligence is centralized (the master node) and where nodes are dumb (sensors, actuators). Kept in this context, Bitbus is usable.

But Bitbus is not a LAN. It does not support distributed intelligence; it does not even support self-initiated transmission outside the master node. Its limit of 14 bytes of user data per message is acceptable for transfer of control information (for which it was designed), but would make file transfers extremely tedious. There is no software support above the data link layer other than the normal tools (compilers, linkers, loaders, operating systems, real-time executives) distributed by Intel.

Advantages and disadvantages

Advantages: compared to an equivalent star architecture, it reduces wiring costs.

Disadvantages: limit of 254 nodes; no support of dynamic topology; implemented on noise-sensitive twisted pair; data rate limited by bus length; maximum data rate 2.4 Mbps over 30 meters; maximum bus length of 10 kilometers only achievable by using 10 active repeaters; communications must be initiated by master; data transfers limited to one message per poll; messages limited to 14 bytes of user-data; data link protocol is only a subset of SDLC; no software support above the data link layer; no network analysis tools; not a standard.

Opinion: the centralized bus controller introduces an unreliable element; the master/slave communications protocol is inappropriate for a local area network.

Pricing

Bitbus is deemed unsuitable for the intended shipboard application, so pricing was not pursued.

5.11 DATAPOINT Arcnet

Summary

Arcnet is the communications mechanism which connects Datapoint computers and IBM PCs. The whole network is called ARC, Attached Resource Computer, to indicate that the collection of computers is treated as one large, distributed computer by the operating system. Arcnet uses a proprietary token-passing scheme to allow communications at a rate of 2.5 Mbps over coaxial cable at distances to 4 miles. Up to 255 processors can participate in an ARC system.

General information

Datapoint introduced the ARC concept in 1977, and thus was an early leader in computer networking. Since the Datapoint network predates the ISO and IEEE standards, there is no compatibility between Arcnet and the three IEEE protocols. Responding to market pressure, Datapoint abandoned its original scheme of interconnecting Datapoint-only devices by introducing hardware and software for the IBM PC and PC XT. Still, Arcnet remains a computer-to-computer connection scheme.

Arcnet architecture

Arcnet is a baseband, token-passing scheme which transmits digitally over coaxial cable at 2.5 Mbps among up to 255 stations over distances up to 4 miles. Arncet uses a branching tree topology.

Each attached computer needs a Resource Interface Module (RIM) which contains a microprocessor, 1k or 2k of local memory, clock, interface to the network cable, and an interface to the local ARC processor. RIMs can be viewed as specialized, serial interfaces.

Arcnet uses a proprietary form of token passing to arbitrate network access. Each station has a unique 8-bit network address, and the token passes from one station to another in order. When a RIM receives a token, it checks its transmit buffer; if empty, it passes the token to its successor. If the transmit buffer is not empty, the RIM send a "free buffer request" to the destination which asks it to reserve buffer space for the message soon to be transmitted. If the destination station has a free buffer it responds with an ACK, otherwise with a NACK. If the receiver is not operational, the sender waits a specified amount of time (about 75 microseconds) and then passes the token.

When a new station adds into the network, it sends a special signal which is longer than the longest legal transmission. Upon hearing that signal, each station waits an amount of time based on the result of subtracting its network address from 255; thus the highest numbered station waits the shortest time. That station creates a new token and transmission resumes.

The Arcnet protocol is character-based, rather than bit-oriented. Messages are encoded as 11-bit words, and messages are framed with ASCII characters such as start-of-header (SOH), end-of-transmission (EOT), and enquire (ENQ), and uses the ASCII ACK and NACK characters to respond to free buffer requests.

Arcnet uses a star topology; all computers have a coax cable running to a central hub. The hub may be passive, in which it is just a physical connection among cables, or active, it which it amplifies incoming signals and repeats them on all outbound cables except the inbound one. There is a 2,000 foot limit on cabling between the computer and the hub; there can be no more than 10 hubs on the longest Arcnet path; the longest path cannot exceed 4 miles; and no connection can result in a loop.

Datapoint has recently introduced a fiber optic link which doubles the maximum hub-to-hub spacing to 4,000 feet; the fiber optic hub has one fiber optic port and two coax ports.

IBM PC

In September 1984 Datapoint announced the Intelligent Network Executive-PC (INX-PC) option. Consisting of network hardware, software, file server, and an adapter card which fits the expansion slot of an IBM PC or PC XT, it allows the PC under DOS to access the network and communicate with other ARC processers operating under Datapoint DOS.

External communications

Datapoint supports general purpose asynchronous interfaces which operate from 37.5 to 9,600 bps and which fit inside the ARC processors. There is no RS-232 style connection to the Arcnet itself.

Advantages and disadvantages

Advantages: integrates up to 255 Datapoint or IBM processors.

Disadvantages: limited to baseband coax media; limited to 2.5 Mbps; designed to work only with Datapoint and IBM PC microcomputers; no RS-232 style direct network connections.

Opinion: the use of the star architecture raises the issue of reliability, especially since redundant components are not available.

Pricing

COMPONENT	PRICE ($)
RIM (external)	1,650
Passive hub	125
8-port expander card	650
50-foot coax cable and connectors	90
Asynchronous communications adapter	1,950
INX-PC adapter card	695

5.12 SPERRY Avionics Standard Communications Bus

Summary

The Sperry ASCB is a proposed standard for aircraft avionics busses. ASCB uses a bus topology and a master/slave control philosophy. Two bi-directional busses are used for enhanced reliability. The bus master always transmits on both busses; all slave devices listen on both busses; a slave device can only transmit on one bus; and the expectation is that device errors will be more common than bus controller errors, so it is unlikely that both busses will be incapacitated by a user device error. A master bus controller (possibly redundant) is used to poll connected devices and request data transmission. "Slave" devices may not initiate transmission unless polled by the bus master. The data link protocol uses HDLC framing. Network bandwidth is partitioned among four groups of users: those who will be polled 40, 20, 10, or 5 times per second. The basic data rate is 667 Kbps.

General information

Sperry has developed the ASCB as an alternative to the ARINC 429 bus, the MIL-STD-1553B bus, and local area networks using either CSMA/CD or token passing network access mechanisms. The claim is made that ASCB is more efficient and more cost-effective than these other three approaches. The ASCB has been developed, installed, and certified in 6 aircraft.

ASCB architecture

The argument is made that bus topologies reduce wiring costs over traditional (at least in the military) star topologies; this is undoubtedly correct. Thus the ASCB uses a dual bi-directional bus design (busses A and B) with half the attached devices on bus A and the other half on bus B. A centralized bus controller is connected to both busses and is used to allocate network access to the attached devices. Using the ISO-standard HDLC protocol, the master issues polling messages to the attached devices which gives them the opportunity to transmit.

A crucial argument concerns the inherent reliability of the bus controller. The ASCB documents state that distributed control schemes are more difficult to design, more costly to implement, and less reliable that centralized control schemes. There is a narrow range of applications for which this is true, and aircraft avionics busses might be one. But the generally accepted notion is that *carefully designed* distributed control is more reliable that centralized control because no one device, or no one failure type, can bring down the whole system. That is in contrast to the bus controller concept in which failure of the bus controller causes the whole network to fail.

Bus controller failure is obviously a concern, since the documentation talks about dual-, triple-, and quad-redundant implementations of the bus controller being an acceptable technique for raising the reliability of the bus controller to any prescribed level. This is more likely to be true in theory than in fact. For instance, there is no indication that the ASCB actually has redundant bus controllers, only that it is capable of supporting them. Further, there is no mention of how a failed controller might be detected, or how the network would switch to a redundant controller, or what station would be responsible for the switchover, or how attached devices would distinguish between messages from a flawed primary controller vs. messages from a redundant controller. Also unmentioned is how any number of redundant controllers would gain access to the data busses if the primary controller had an error like an infinite transmit loop (jabber).

One of the most interesting comments made was that, for "real" reliability, redundant controllers cannot be simple duplicates; that would leave them vulnerable to a common software fault. The suggestion is made that at least one redundant controller should have a different processer and be coded somehow "differently" (in fact this is the approach taken with the five computers on the NASA space shuttle — one is different hardware and programmed independently). However, these issues are not resolved, and there is no indication that these observations have been reduced to practice.

Turning to the dual bi-directional busses, the controller is attached to both, half the devices are attached to A and the other half to B. Here the argument is that "any one user device failure could at most incapacitate one bus, A or B, but not both, and the probability of two independent failures on separate busses is extremely unlikely." Perhaps this is indeed an observation of the avionics group, but it lacks statistical verification. Also, devices tend to share a common design for the bus interface, so a common design error, common hardware fault, or common software error would invalidate the assumption of independence.

One assertion which can be verified is that the ARINC 429 and MIL-STD-1553B busses are high in overhead and low in throughput due to an elaborate polling structure. While the ASCB also uses a polling strategy, it is more efficient than the others.

The performance gains attributed to ASCB over the others comes primarily from the *a priori* division of network bandwidth into four classes: those devices which will be polled 40, 20, 10, or 5 times per second. The other networks allow more flexibility but at the cost of more overhead.

Finally, with regard to architecture, one must ask whether a shipboard environment is in any way equivalent to an avionics environment. Aside from the common desire to share data among multiple devices, the two environments seem distinctly different. One goal of the ASCB is to gather data from remote devices and centralize that data somewhere (like an air data computer). In that sense ASCB resembles the functionality of Intel's Bitbus. If that is indeed the application then one can make a stronger argument for the topology and the protocol. But if the goal is to share data, rather than accumulate data, then the utility of a "real" local area network outweighs that of a "device interconnect network".

ASCB hardware

Specific details on ASCB hardware and interfaces are not given. It is clear that it is intended to interconnect devices more than to interconnect computers. The shipboard system must be responsive to both needs.

Advantages and disadvantages

Advantages: in the context of avionics busses with repeatable, predictable information transfer, the pre-scheduled approach of ASCB is efficient.

Disadvantages: the allocation of network bandwidth into four groups based on polling frequency is inappropriate for stochastic information transfer.

Opinion: the centralized bus controller creates a reliability issue; the master/slave polling concept is inappropriate for a local area network; the dual-bus architecture creates a reliability issue; the data rate of 667 Kbps is probably too slow for good overall performance.

Pricing

The ASCB is deemed inappropriate as a shipboard LAN, so pricing was not pursued.

6. RECOMMENDATIONS

6.1 Not Recommended

6.2 Cost Analysis of Remaining Systems

- 6.2.1 CONCORD DATA SYSTEMS Token/Net
- 6.2.2 UNGERMANN-BASS Net/One
- 6.2.3 IBM Token Ring
- 6.2.4 PROTEON ProNET

6.3 Ranking of Candidate Systems

6. RECOMMENDATIONS

6.1 Not Recommended

The following networks are not recommended for the reasons stated.

NETWORK	PRIMARY REASONS FOR NON-SELECTION
SYTEK LocalNet	128 Kbps channel speed; CSMA/CD access scheme; no microprocesor bus PC support
INTEL Ethernet	available for Multibus only; inefficient software
BRIDGE COMMUNICATIONS Ethernet	CSMA/CD access scheme
INTERACTIVE SYSTEMS/3M LAN/I (based on Arcnet)	Arcnet standard not widely accepted; no direct connections for microcomputers
INTERACTIVE SYSTEMS/3M LAN/II (Ethernet version)	CSMA/CD access scheme; no direct connections for microcomputers
INTERACTIVE SYSTEMS/3M LAN/II (token bus version)	head-end frequency translator is a single point of failure
UNGERMANN-BASS Net/One (Ethernet version)	CSMA/CD access
UNGERMANN-BASS Net/One (Optical fiber version)	requires crucial star coupler; possibly requires "collision reinforcer"
INDUSTRIAL NETWORKING MAP/One	compatibility with MAP does not appear to justify increased price
INTEL Bitbus	master/slave protocol; 14-byte user message limit
DATAPOINT Arcnet	standard not widely accepted no RS-232 ports
SPERRY ASCB	master/slave protocol; 667 Kbps channel speed

6.2 Cost analysis of remaining candidates

Each of the following networks was judged to be worthy of further consideration. Each network had some particular advantage, but each one also had some particular weak point; these are summarized.

A cost analysis was made for each network. The costs included for purposes of comparison were the minimum hardware and software costs incurred to interconnect ten serial I/O devices. Not included in the cost comparison was the cost of cabling, or of optional network software, or of network diagnostic equipment. If particular hardware or software was *required* to make the network operate, it was included.

6.2.1 CONCORD DATA SYSTEMS Token/Net

Primary strengths: fully IEEE 802.4 compliant; reliable company; proven broadband technology; excellent network diagnostics.

Primary weaknesses: achieves only 5 Mbps; no PC direct connection, only RS-232 ports.

Quantity	Equipment	Price Each ($)	Total ($)
1	TIM-220 with 12 RS-232 ports	5,940	5,940
1	TIM-200 with 4 RS-232 ports	3,485	3,485
1	HR-105 head-end remodulator	4,450	4,450
1	Network Control Center (NCC)	6,000	6,000
	Total		19,875

6.2.2 UNGERMANN-BASS Net/One (broadband version)

Primary strengths: reliable company; proven broadband technology; directly interfaces microprocessors.

Primary weaknesses: uses CSMA/CD access; speed only 5 Mbps; *not* IEEE 802.3 compatible.

Quantity	Equipment	Price Each ($)	Total Price ($)
5	Network Interface Unit	$ 1,095	$ 5,475
1	2-port NIU	1,400	1,400
1	8-port NIU	2,950	2,950
15	RF modems	650	9,750
1	Head-end remodulator	3,500	3,500
	Total		23,075

6.2.3 IBM Token Ring

Primary strengths: uses ring topology; uses shielded twisted pair (economy); uses "wiring center" concept via Multistation Access Unit; software available for one-time purchase; fully IEEE 802.5 compliant.

Primary weakness: operates at only 4 Mbps; wiring-center-to-wiring-center length limitation of 200 meters; host-to-wiring-center length limitation of 100 meters; fiber optic interfaces not yet available; network diagnostic hardware/software not yet available.

Quantity	Equipment	Price Each ($)	Total ($)
5	PC Adapter	695	3,475
5	PC Adapter Cables	35	175
2	8-port Multistation Access Units	660	1,320
1	NETBIOS program	35	35
1	Asynchronous Communications Server program	495	495
	Total		5,500

6.2.4 PROTEON ProNET

Proteon has just announced a new 4 Mbps product, ProNET-4, which is IEEE 802.5 compatible. Their 10 Mbps product, ProNET-10, runs a proprietary protocol which is similar to 802.5. They have also restructured their equipment prices to compete more directly with IBM.

Primary strengths: uses ring topology; IEEE 802.5 version available; supports wiring center concept; reliable company; full 10 Mbps communications rate; directly supports microprocessors; operates on shielded twisted pair (economy); fiber optic link available; redundant fiber optic links available; reversible dual-ring fiber optic link available; supports mixed media between wiring centers.

Primary weaknesses: does not interface RS-232 devices directly.

Quantity	Equipment	Price Each ($)	Total ($)
5	microprocessor interface card	799	3,995
3	4-node wiring center	330	990
	Total		4,985
	Optional fiber optic configurations: (additional equipment)		
3	wire-center-to-wire-center fiber optic connections	2,300	6,900
	Total using fiber optic backbone		11,885
3	wire-center-to-wire-center redundant fiber optic connections	2,900	8,700
	Total using redundant fiber optic backbone		13,685
6	fiber optic repeaters for dual-redundant, contra-rotating rings	1,700	10,200
	Total using dual fiber rings		15,185

6.3 Ranking of Candidate Networks

In section 3.7 we presented six criteria for selecting a shipboard LAN:

(1) System reliability
(2) Baseband vs. broadband transmission
(3) Coax vs. fiber optic cabling
(4) Adherence to international standards
(5) Cost
(6) Network performance

We now reexamine the four candidate networks in light of these criteria.

System reliability

Best: ProNET — due to its mature wiring center concept
Second: IBM Token Ring — second because its bypass circuit is slower than ProNET's
Third: Net/One — cannot recover from a shorted bus
Third (tie): Token/Net — head-end remodulator is a critical component

Baseband vs. broadband transmission

If broadband is desired such that voice and video can be carried simultaneously with data, then the only choice is the Token/Net; it is the only broadband system of the four.

Copper vs. fiber optic

We recommend fiber optic cabling in high-noise areas such as the engine room, although specific experiments are proposed in section 6.6 to determine the noise threshold at which fiber optics are required. In terms of their ability to support fiber optics as an optional medium:

Best: ProNET — single-cable, dual-cable, and reversible ring fiber optics easily accommodated, although expensive
Second: IBM Token Ring — although not currently available, fiber optics are part of the IBM cabling system (type 5 media) and will become available
Third: Net/One — supports fiber optics, but uses star coupler which becomes a critical component
Fourth: Token/Net — does not support fiber optics

Adherence to international standards

Best: Token/Net — compliant with 802.4
Second: IBM Token Ring — compliant with 802.5
Third: ProNET-4 — this one version of ProNET is compliant with 802.5
Fourth: Net/One — not IEEE standard

Cost

Recall that the cost figures in section 6.3 were based on the minimum equipment necessary to support 5 PCs and 10 RS-232 devices. Unless absolutely necessary for network operation, network monitors and such were not included. Software and cabling costs were not included.

Best: ProNET-4 (least expensive)
Second: ProNET-10
Third: IBM Token Ring
Fourth: Token/Net
Fifth: Net/One

Performance

The criterion of network performance was not explicitly included in section 3.7 because it was known that network loading was very low (see section 7). All four candidate networks would carry the projected network offered load with no trouble. Still, in terms of *rated* capacity and *expected* capacity (i.e. the capacity deliverable to the end user), this is their ranking.

Ranking	Network	Rated capacity	Expected capacity
Best:	ProNET-10	10 Mbps	9 Mbps
Second:	Token/Net	5 Mbps	4 Mbps
Third:	Net/One	10 Mbps	4 Mbps
Fourth:	ProNET-4	4 Mbps	2 Mbps
Fourth (tie)	IBM Token Ring	4 Mbps	2 Mbps