John Leong, Carnegie-Mellon University, Pittsburgh, Pa.

CMU-ITC-85-057

Nuts-and-bolts guide to Ethernet installation and interconnection

A large university reports on some frequently overlooked problems faced when linking several departmental networks.

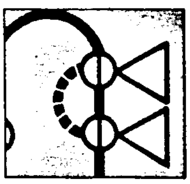

Currently, various departments within Carnegie-Mellon University run 12 separate 10-Mbit/s Ethernet networks. The networks vary in geographical size, ranging from small backbone networks that connect computing facilities inside a computer room to big department networks with 50 or more stations covering a broad area of the campus. The larger networks tend to be populated by a higher percentage of transient stations or microcomputers and terminals that leave and rejoin the network frequently.

There are, at present, three versions of the Ethernet specification:
- Ethernet specification version 1, September 1980.
- Ethernet specification version 2, November 1982.
- IEEE 802.3 carrier-sense multiple access with collision detection (CSMA/CD) specification, revision D, December 1982.

Version 1 is the original Ethernet specification jointly developed by Digital Equipment Corp. (DEC), Intel Corp., and Xerox Corp. The majority of Ethernet products available today are designed according to this specification. Yet in the process of making this specification an IEEE standard, some changes have been made. The main differences among the three versions lie in the connection from the data-terminal equipment (DTE) controller to the media-access unit (MAU) transceiver interface.

For instance, the line-idle state was changed from 0.7 volts idle (version 1, Fig. 1A) to zero volts idle (version 2, Fig. 1B). As a result, the first bit of the message preamble for version 2 starts with a voltage half step instead of a voltage full step. Differences between the versions also exist with the collision-present test as well as with the interface coupling specification.

Because of these differences, a controller designed according to one specification may not work with a transceiver designed according to another. In the past, this has not been much of a problem, since most of the equipment on the market conforms to the Ethernet version 1 specification. However, within the past few months, equipment conforming to the official IEEE 802.3 standard has started to appear in quantity. Compatibility problems arise from the Manchester encoder/decoder chips now on the market, which are designed for the IEEE 802.3 specification with Ethernet version 1 support only optional. (The Ethernet version 2 specification was introduced as an attempt to bring version 1 in line with IEEE 802.3, but it has not generated much interest.)

Down to earth

An often overlooked aspect of Ethernet installation is cable grounding. All the specifications stress the importance of cable grounding at a single point to ensure reliable communications. Still, networks are frequently installed with either multiple groundings or no ground connection at all. And though networks so installed work in most cases, stability problems can easily creep in as the sites grow.

For proper installation, cable connectors and terminators should be covered by insulating sleeves to prevent accidental grounding. During installation and maintenance, accidental grounding can be a potential safety hazard. Detailed safety issues are described in the IEEE 802.3 Ethernet specification.

The most common Ethernet cable is the standard 50-ohm coaxial cable covered with yellow polyvinyl chloride (PVC). Carnegie-Mellon uses Beldon 9880 cable, which costs roughly $0.80 a foot. (Similar cable is available from other manufacturers.) The basic PVC

1. Incompatibilities. *Ethernet version 1, Fig. 1A specifies an idle of 0.7 volts; version 2, Fig. 1B specifies an idle of zero volts. Version 2's preamble is a half step.*

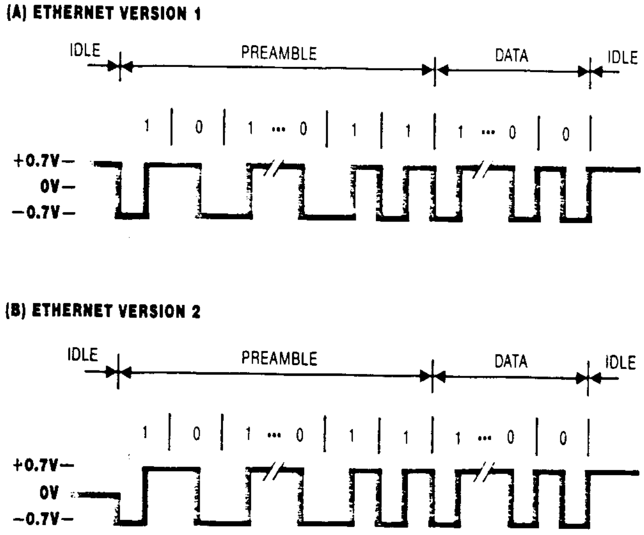

2. Thick and thin. *RG58 50-ohm cable is cheap and flexible and has high signal attenuation. An N-type adapter allows connection with Ethernet cable.*

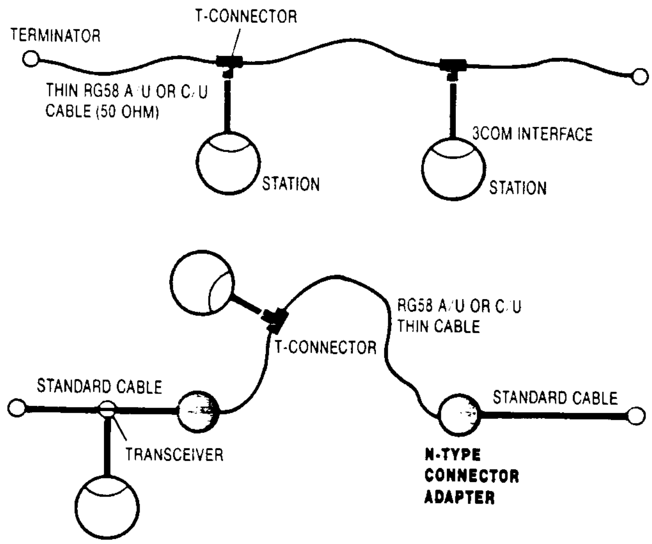

cable cannot be used for plenum (air space in the environment) installation. A Teflon-coated cable is required for these applications. Teflon is a much stiffer cable and has a large bending radius. Predictably, it is more difficult to install and costs up to four times as much as the PVC cable. But in most cases, the Teflon alternative is still cheaper than installing new conduits.

The maximum length of an unrepeatered Ethernet segment is 500 meters. Cable can be installed in short sections joined by connectors. According to the 802.3 specification, the preferable section lengths are 23.4, 70.2, and 117 meters. These lengths will produce minimum signal reflection. To ensure an even cleaner signal, the cable sections should be obtained from the same manufacturer and—even better—from the same manufacturer's lot.

According to the IEEE specification, transceivers must be placed at least 2.5 meters apart. It is therefore desirable to have the cable marked at those intervals. Some cable comes with marking already on the jacket. If not, most dealers will mark the cable at no additional cost, since they have to measure out the cable anyway. During installation, number the markings sequentially to make troubleshooting the network easier.

For operation and maintenance, it is desirable to have accurate and detailed documentation of the cable installation. This should include not only the location of the cable, but also descriptions of the cable trays, raceways, and conduits that the cable traverses. Such documentation is valuable for troubleshooting and future expansion (Data Communications, "How to avoid some common pitfalls of a local net installation," March 1984, p. 243).

Another popular type of cable is the 50-ohm RG58 A/U or C/U. This is a thin, flexible, inexpensive cable that costs approximately $0.09 per foot. Both Ungermann-Bass and 3Com support this type of cable—especially with their microcomputer products. When RG58 is used with the 3Com Etherlink IBM PC interface, no transceiver is required. This represents a savings of approximately $250 per connection. Figure 2 shows a configuration using such a cable. Note that intermixing the RG58 cable with the regular Ethernet cable is allowed.

Weighing against the RG58 cable's low cost and ease of installation is its exposure to damage. With the standard Ethernet cable, the main cable is typically hidden in the ceiling or under the floor, and only the drop cable is exposed. Damage to or disconnection of a drop cable will normally have no effect on the overall network. But with the RG58, there is no drop cable. This thin Ethernet cable is exposed and runs directly into the stations. Damage or disconnection (at the T-connector) can bring down the entire network. Another RG58 disadvantage is its high signal attenuation, which can reduce the maximum permissible segment length by a factor of three. In general, the RG58 thin cable is suitable for use in a small environment with microcomputers. Further, if this cable is selected, the authentic RG58 A/U or C/U should be purchased. The cheaper RG58-like cables are not as good and should be avoided.

Links

There are two popular types of connector—crimp-on and screw-on. The crimp-on type requires a tool for installation. Barrels and connectors cost approximately $5 apiece from Amphenol. Amphenol also sells crimp-on tools for $120. The older version of the screw-on connector comes with many pieces and is difficult to assemble.

A newer connector from Cambridge Products consists of only two parts: a contact and the connector itself. It requires no special tool and is relatively simple

3. Critical placement. *Special circuit-insertion relays: (A) shows Ethernet terminated in two discrete sections, (B) presents a thin-wire Ethernet for fault isolation.*

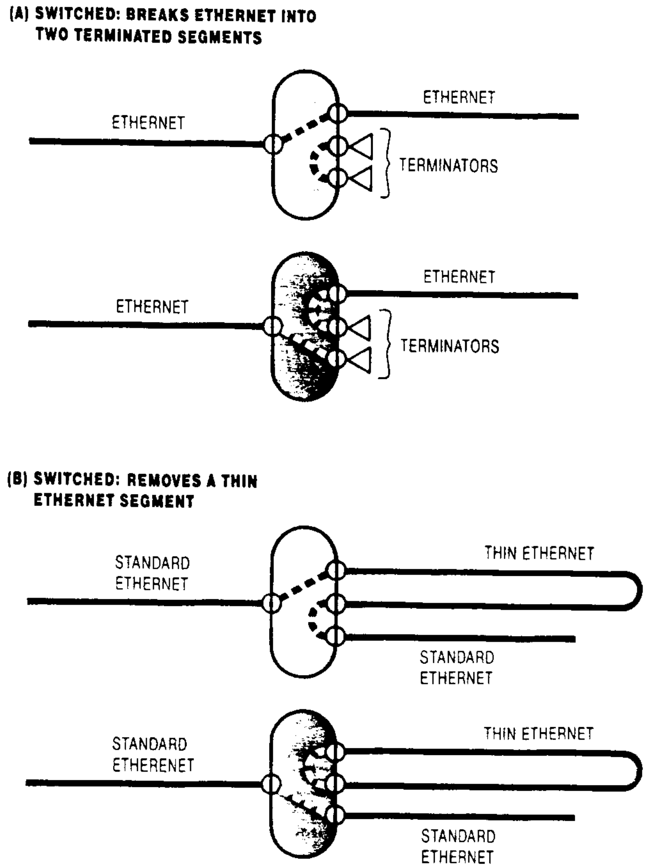

to install. This connector works well and costs from $3.50 to $4.75 apiece. Plastic boots or sleeves should be used with the connectors to provide ground isolation.

Placement of circuit insertion relays at strategic points of the Ethernet cable can be invaluable for fault isolation. Carnegie-Mellon has used N-series coaxial RF (radio frequency) relays from Amphenol. These cost about $100 and can be configured either for breaking the Ethernet cable into two discrete terminated segments (Fig. 3A), or for taking a loop out of troublesome cable (Fig. 3B). The loop configuration is particularly useful when using a mixture of thick and thin Ethernets. If the thin cable is damaged, it can be easily taken out of service.

Drop cable can either be homemade or bought off-the-shelf. Almost every company that sells transceivers also sells drop cables. If the cable must be threaded through narrow conduits or other tricky access paths, one should get cables with shells that can be taken apart. DEC offers two grades of cable: the standard drop cable and a thinner, more flexible variety. The more flexible grade is easier to install but has roughly four times the attenuation.

Carnegie-Mellon has used a variety of homemade drop cables. They range in cost from $0.30 to $0.60 a foot. Although all of the cables worked well electrically, some were better than others in terms of their thickness, flexibility, and the ease of installation. Carnegie-Mellon has had positive experience with both Phalo 2384 and Belden 9891 cables.

Before selecting a transceiver, one should make sure that it matches the controller. That is, a controller conforming to the Ethernet version 1 specification should be used with a version 1 transceiver. And an IEEE 802.3 controller should be used with a similar transceiver. The transceiver must supply the collision-present test if the controller requires it.

Taps

Transceiver installation can be tricky since the transceiver must be placed directly on the Ethernet cable; and often, the cable tray or conduit is difficult to access. Furthermore, most of the conduits and cable trays are often already congested and are not designed for placement of bulky objects like transceivers. Such forced transceiver installation into crowded conduits can cause shorting. This is the most common cause of Ethernet failure.

Carnegie-Mellon has used 3Com, TCL, DEC, and Interlan transceivers. Of the four, 3Com is the only invasive type. It must be inserted in series on the cable and requires cutting the cable during installation. All the other transceivers are of the "vampire" tap variety, clamping directly onto the cable with metal teeth. The 3Com transceiver does not supply a heartbeat, or collision-present test signal. Nevertheless, it works well without interfaces and can be placed less than 2.5 meters apart.

The Interlan and DEC transceivers have the same physical housing but different electronics. Both are physically bulky. They require special, and rather costly, installation tools. Removing these units from the cable is also difficult.

The TCL transceiver is smaller than the DEC and 3Com devices. It is easier to install and remove, though it does not fasten as securely to an Ethernet cable as do the others. The Interlan transceivers come in two varieties: the E-series and the I-series. The E-series conforms to Ethernet version 1, and the I-series conforms to IEEE 802.3.

Carnegie-Mellon has been experiencing compatibility problems with the new IEEE 802.3 interface boards because it has mainly installed Ethernet version 1 transceivers. All the transceivers cost from $200 to $250. Some are available with a substantial quantity discount. The university prefers the TCL model and has installed more than 100 of them.

Multiport transceivers

All multiport transceivers support two operating modes: They can be connected to a transceiver and provide attachments to multiple stations or used as standalones. When standalone, the multiport transceiver itself functions as an Ethernet network. By cascading devices (Fig. 4B), most models will support up to 64 stations. When used in-line with an Ethernet network transceiver (Fig. 4A), the maximum distance

4. Cascading. *Multiport Ethernet transceivers can act as separate local networks and can be cascaded to include as many as 64 stations.*

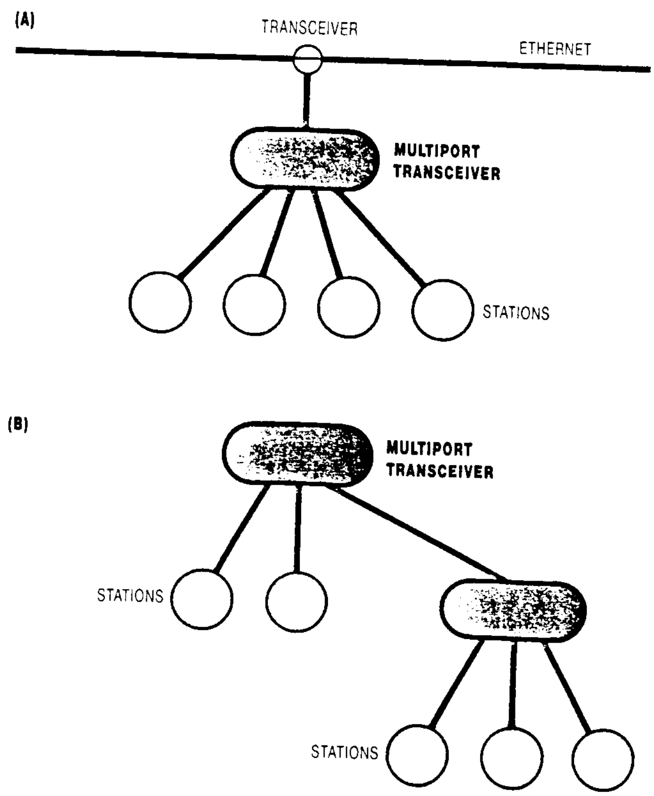

between the Ethernet and the station is typically reduced from 50 meters to 40 meters. This is due to the added delay incurred in the multiport transceiver.

Not only can a multiport transceiver reduce the per-connection cost, but it can also provide a degree of centralized control for problem determination and isolation. In addition, this type of transceiver eliminates the 2.5-meter separation requirement between devices and is particularly useful in a cluster situation. In one department at Carnegie-Mellon, minicomputers are connected to one multiport transceiver.

Repeaters

Repeaters connect multiple Ethernet segments. They are either local or remote. Local repeaters can be used if the two Ethernet segments are less than 100 meters apart. Otherwise a remote repeater must be used. Ungermann-Bass and DEC sell both types of repeaters.

A local repeater is typically housed in one cabinet with two transceiver-cable connectors, one for each Ethernet segment. A remote repeater consists of a pair of boxes containing half-repeaters, one for each segment of the Ethernet. Two half-repeaters are connected by point-to-point links of up to 1 kilometer. Since remote repeaters are mainly used for interbuilding connections, and taps are rarely required between the two repeaters, optical fiber is often used for this link. Fiber-optic cable also avoids the problems that lightning and other electrical interference can cause.

The fiber-optic cable used by AT&T and other Bell operating companies (BOCs) is the 50-micron or 62.5-micron variety. Although this type of cable is good for long-distance connections, it is not as suitable for local network environments, where connector loss is a significant factor. Furthermore, since the 50-micron fiber-optic cable has 25 percent of the surface area of the 100-micron light source, the 75 percent light loss—approximately 8 decibels (dB)—can be damaging. As a result, the university has had difficulty using the DEC remote repeater, which has a flux budget (total amount of repeater power) of 10 dB. The Ungermann-Bass repeater fares better with a flux budget of 16 dB.

While remote repeaters from both DEC and Ungermann-Bass specify 100-micron fiber-optic cables, the university has used them over its 50-micron cable plant with no serious problem. Repeaters from both manufacturers have status lights for each segment. These are useful aids for indicating incoming data and collision detection.

If one segment to the DEC repeater goes down, the device will isolate that segment from the network, rather than drag the other side down. In addition, it is possible to deploy a spare DEC repeater in parallel with the main unit and have the spare set for standby mode. When the main unit fails to repeat eight full packets, the spare repeater will automatically take over.

Other Ethernets

Both Codenol and Ungermann-Bass market fiber-optic Ethernets. These are star-shaped networks with optical defusers at the hubs. A station is connected to the hub through a pair of fiber-optic cables (Fig. 5). The hub ranges in price from $900 to $4,000, depending on the number of ports (5, 8, 16, or 32 ports).

At present there are two basic problems with this technology. In the first instance, the flux loss by the hub is relatively high—especially with 50-micron cable. The second problem is with collision detection in an asymmetrical network: If the network consists of stations a long way from the hub, as well as some that are close to the hub, a station close to the hub may not be able to detect collision with a station that is far away. Ungermann-Bass overcomes this problem by placing an optional collision-detection unit at the hub, but the unit is quite expensive ($12,000 for seven ports and $20,000 for 14 ports).

American Photonics has a fiber-optic drop cable called an Ethernet expander—model number RL5000. The expander allows stations to connect to a transceiver over an optical-fiber pair at distances up to 1 km. While this violates the Ethernet drop-cable distance specification, the device works if the overall network topology does not exceed the Ethernet-specified slot-time limitation. Careful network planning is required when using this product. The RL5000 is listed for $3,400.

Ungermann-Bass also offers a broadband 5-Mbit/s network over a 6-MHz channel. The network uses CSMA without CD and is therefore subject to a higher probability of packet loss. IBM's recently announced PC Network uses CSMA/CD over broadband. The

5. Stars. *Fiber-optic Ethernets are usually configured as star networks radiating from a single hub. A station is connected to a hub by a pair of fiber cables.*

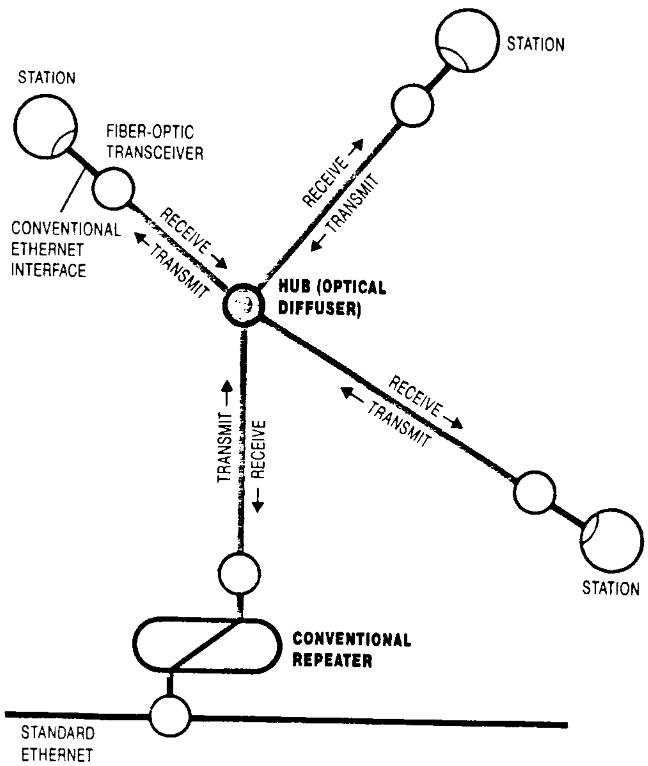

interface board contains a simple fixed-frequency RF modem, an Intel 82586 CSMA/CD chip, an Intel 8088 processor, and 128 kbytes of RAM. Although the IBM network uses CSMA/CD, it is not Ethernet compatible in the way it implements its network parameters. Furthermore, it operates at sub-Ethernet speeds (2 Mbit/s) over a 6-MHz channel.

Monitoring tools

Excelan makes an Ethernet monitor called the Nutcracker. The product is similar to serial-line monitors commonly used on conventional data communications networks. The Nutcracker allows capture of all network traffic with or without user-specifiable filters. The device, however, has no special provision for protocol interpretation and display. At $50,000, the Nutcracker is an expensive diagnostic tool.

For general protocol development and network administration not requiring the capture of continuous back-to-back packets, there are a number of diagnostic software packages available for popular machines such as the IBM PC. These are often free.

Problem determination

An operation logbook can also be a valuable problem-solving tool. Lots of problems are caused when installing or changing connectors or transceivers. Often, matching a problem's symptom with the most recent logbook entry can help users locate the problem.

Another important tool is a set of cabling documents for the plant. These help tremendously in locating suspected faulty components.

The availability of cable status indicators can also be useful. These indicators—the most common and useful of which are the receive-data and collision-detection lights—are available on most Ethernet repeaters. With these, it is possible to determine the network load and the frequency of collisions. If there are no repeaters on the network, signal indicators can be purchased for approximately $150.

For a network of any significant size, users should purchase a time domain reflectometer (TDR). These devices help determine the position of cable shorts, breakages, bad connectors, sharp cable kinks, and other abnormalities. TDRs have been used with the telephony and cable industry for almost a decade. Essentially, the unit sends out a sample pulse. Depending on the resultant return pulse, abnormalities on the cable can be observed on the device. These devices sell for approximately $5,000.

Once the distance between the test point and the suspected problem location is known, it is necessary to find out exactly where that section of cable is on the network. If the cable is marked at 2.5-meter intervals, and each marker is sequentially labeled, it is easy. For example, if a problem is 250 meters away from a test point near marker 5, the nearest marker to the problem area should be number 105. With a cable-plant diagram, it is easy to get an idea of the area in the building where marker 105 should be.

Common problems

In general, Ethernet and related problems produce one of three symptoms: no network communications, a station malfunction, or network performance degradation.

■ *No network communications.* This symptom is probably the result of a cable break, a short, or a jabber problem (in which a device sends data out of control). In all three instances, the problems can be detected with an ohm meter. In the case of jabber, the cable impedance will be normal, whereas impedance will be exceptionally high with a cable break and low with a shorting. A TDR will also allow quick detection of shorting and cable breaks.

If the jabber-control logic on all the transceivers is working properly, this problem should not occur. If not, it is sometimes possible to determine which station is jabbering by monitoring the packets and collision fragments in the network. If this is not possible, and if the network is small, stations may be powered-off, one by one, until the offending station is found. A potential problem with this approach is that the stations are often located inside locked rooms. In this case, it helps if the stations are connected to transceivers that are accessible. Stations can then be disconnected from the transceiver without entering the room.

If the network is large, the problem must first be isolated to a small manageable section of cable. This is applicable to all three classes of problems. Cable segmentation can be done in the following order:

6. Easy access. *It is desirable to configure large networks as a series of star networks connected by a bus. The stars are multiport transceivers, the bus is the Ethernet. But with this configuration it is essential that maintenance personnel have complete access to all transceivers on the network.*

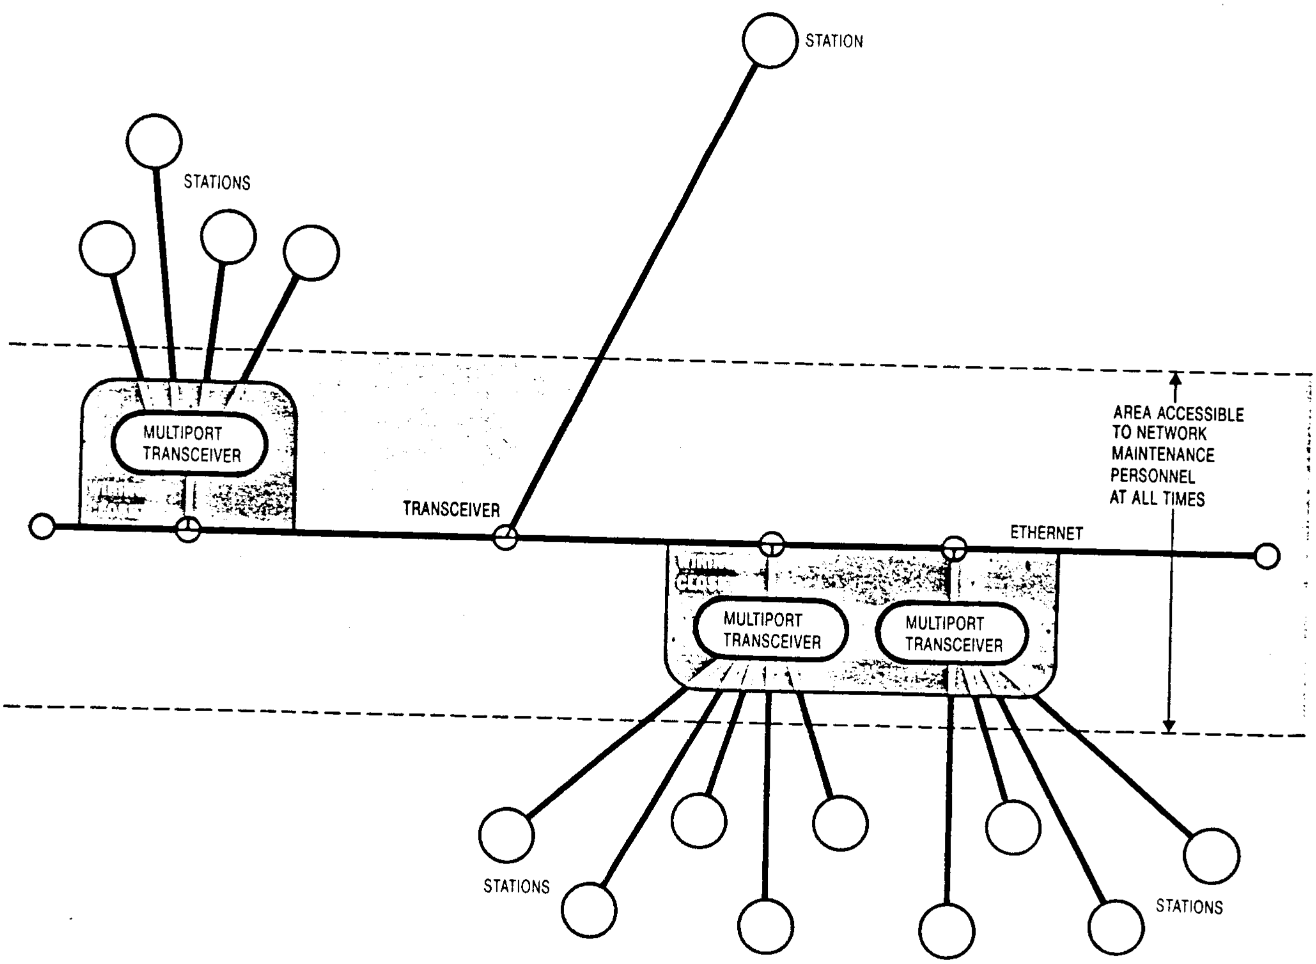

power-off repeaters, switch RF relays, disconnect the cable where the invasive transceivers are located, and power-off multiport transceivers.

From an operational point of view, it is highly desirable to configure the network as a series of stars or stars connected by a bus (Fig. 6). The stars are multiport transceivers in strategically located wiring closets with drop cables radiating from them to stations. The bus is the Ethernet. Although this topology calls for longer drop cables, it provides network operations staff with focal points for troubleshooting and maintenance.

■ *Station malfunction.* This is most likely due to a problem with the station hardware, the Ethernet interface, station software, the transceiver unit, and/or the drop cable. If it works to replace the suspected station with a good station, the problem is likely to be with the suspected station, not with the transceiver. For this type of testing, it would be ideal if all the stations implemented the configuration-testing protocol defined in the Ethernet version 2 specification. But few machines offer it. Some networks use a higher-level protocol that provides signal echo. Although echo provides less comprehensive testing than the configuration testing protocol, it is useful as a go/no-go test.

■ *Network performance degradation.* It is possible for a station to fail in such a way that it will continuously send out maximum-size packets. Although this will not completely overload the network, the overall performance can be significantly degraded. This problem can often be identified with a network monitor, especially one that can show the network's heavy users.

Such problems are not frequent unless the cable is often being worked on or network management has been careless about the cable. At Carnegie-Mellon the Ethernet networks have been more than satisfactory, even under relatively heavy loading. The university has supported a variety of uses—from low traffic to high-response requirement applications, such as terminal emulation or paging across networks to disk servers—and has not experienced any performance problems due to the Ethernet networks. ■

John Leong received a B. S. and an M. S. in computer science from the University of Manchester, Manchester, England.