Guide to Ethernet

Charles Spurgeon

Networking Services
University of Texas at Austin

Document version 2.5

1.0 Introduction

This is a brief guide to the Ethernet system and the variety of media systems that can be used to build an Ethernet. The whole subject of designing, building, and operating a local area network (LAN) is a large one, and this guide makes no attempt to deal with all of the issues. Instead, this guide provides you with an brief introduction to the Ethernet system, media types, and configuration guidelines.

2.0 The Ethernet System

Ethernet is a LAN technology that transmits information between computers at 10 million bits per second (10Mbps). New Ethernet standards are currently under development that will provide for data rates of 100Mbps.

There are several LAN technologies in use today, but Ethernet is by far the most popular technology for departmental networks. The vast majority of computer vendors provide equipment with Ethernet attachments, making it possible to link all manner of computers with an Ethernet LAN. Because of this widespread use there is a large market for Ethernet equipment, which helps keep the technology competitively priced. The ability to link a wide range of computers using a vendor-neutral network technology is essential in a university environment. For these and other reasons the UT Networking Services group recommends Ethernet technology for use on UTnet.

From the time of the first Ethernet standard, the specifications and rights to Ethernet technology have been easily available to anyone who wished to build Ethernet equipment. This openness resulted in a large Ethernet market, and is one reason Ethernet is so widely implemented in the computer industry today. The specifications for Ethernet were first published in 1980 by a multi-vendor consortium that created the DEC-Intel-Xerox (DIX) standard. Ethernet technology was then adopted by the 802 committee of the Institute of Electrical and Electronics Engineers (IEEE).

The IEEE standard was published in 1985, and its formal title is "IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications." This standard provides an "Ethernet like" system based on the original DIX Ethernet technology. All Ethernet equipment since 1985 is built according to the IEEE 802.3 standard, which is pronounced "eight oh two dot three."

Ethernets can be linked together to form extended networks using devices called bridges and routers. Bridges can be used to link multiple Ethernets within a department to support more computers. Routers are used on UTnet to provide a campus-wide backbone network that spans multiple buildings. While individual Ethernets in a campus LAN system may only support dozens of computers, the total system of UTnet Ethernets linked with bridges or routers supports thousands of machines.

2.1 Operation of Ethernet

Each Ethernet-equipped computer, also known as a station, operates independently of all other stations on the network, and there is no central controller. All attached stations are connected to a shared media system. Signals are broadcast over the medium to every attached station. In order to send an Ethernet packet a station first listens to the medium, and when the medium is idle the station transmits its data.

Access to the shared medium is determined by the medium access control (MAC) mechanism embedded in each station interface. The media access control mechanism is based on CSMA/CD, and functions somewhat like a dinner party in a dark room. Everyone around the table must listen for a period of quiet before speaking (Carrier Sense). Once a space occurs everyone has an equal chance to say something (Multiple Access). If two people start talking at the same instant they detect that fact, and quit speaking (Collision Detection.) The CSMA/CD mechanism is invoked for every transmission on the network. The mechanism is designed to enforce fair access to the shared medium so that all stations get a chance to use the network.

If two stations happen to transmit at the same instant their signals collide, the stations are notified of the collision, and they reschedule their transmission. To avoid another collision, the stations involved each choose a random time interval to schedule the retransmission of the collided fame.

If repeated collisions occur for a given transmission attempt, then the stations begin backing off by expanding the interval from which the random retransmission time is chosen. Repeated collisions indicate a busy network. The backoff process, formally known as “truncated binary exponential backoff,” provides an automatic method for stations to adjust to traffic conditions on the network.

2.2 Elements of the Ethernet System

The Ethernet system consists of three basic elements: 1. the physical media used to carry Ethernet signals between computers, 2. a set of media access control rules embedded in each Ethernet interface that allow multiple computers to access the shared Ethernet channel, and 3. an Ethernet packet, or frame, that consists of a standardized set of fields used to carry data over the system.

Computers attached to an Ethernet send application data to one another using high-level protocol packets, which are carried in the data field of Ethernet frames. The system of high-level protocols and the Ethernet system are independent entities that cooperate to deliver application data between computers. A given Ethernet system can carry several different kinds of high-level protocol data. The Ethernet is simply a trucking system that carries packages of data between computers; it doesn’t care what is inside the packages.

Each computer on the LAN is equipped with an Ethernet interface which is connected to the media system. For the Ethernet media access control system to work properly, all computers must be able to respond to one another's signals within a specified amount of time. To ensure that every computer can hear the network signals within the specified time, the maximum round trip travel time of signals on the shared Ethernet channel must be limited.

The longer a segment, the more time it takes for a signal to propagate over it. To ensure that the round trip propagation timing limits are met, each media variety has maximum segment lengths defined in the standard. The configuration guidelines for Ethernet provide rules for combining these segments so that the correct signal timing is maintained for the entire network system.

If the specifications for media segment lengths and the configuration rules for combining segments are not followed, then computers attached to the Ethernet system may not hear one another's signals within the required time limit, and could end up interfering with one another.

Therefore, the correct operation of an Ethernet depends upon a media system that is built according to the rules for each media type. More complex systems with multiple segment types must be built according to the multi-segment configuration guidelines provided by the IEEE for combining segments. This guide describes the basic rules for each media type, and also contains the multi-segment configuration rules from the standard.

Now that we’ve taken a quick look at how Ethernet works, let’s see what the Ethernet media varieties look like. But first, we need to learn some jargon.

3.0 IEEE Acronyms

The computer world is full of acronyms and jargon, and Ethernet is no exception. Before making a survey of Ethernet media types, let's look at how the IEEE 802.3 standard defines the components used to connect a computer to a media system. This will also serve to introduce the set of acronyms and other jargon used in the standard.

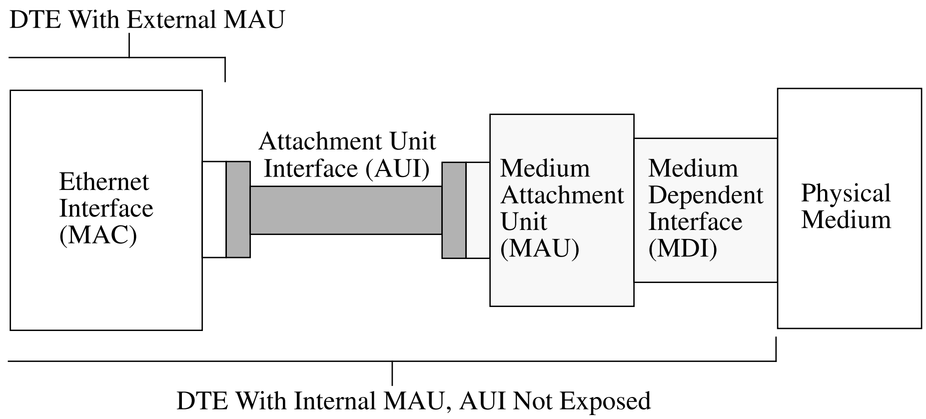

FIGURE 1. Block diagram of a network connection.

The figure shows the set of entities defined in the IEEE standard to describe an attachment to the Ethernet system. While this set of entities and their three-letter identifiers might seem like alphabet soup of interest to network engineers only, these identifiers describe real-world devices that you need to know about.

Starting at the right hand side of the figure we find the physical medium, which is used to carry Ethernet signals between computers. This could be any one of several Ethernet media types, including thick or thin coaxial cable, twisted-pair cable and fiber optic cable. Connected to the medium is the medium dependent interface, or MDI. This part of the standard describes the piece of hardware used for making a direct physical and electrical connection to the medium.

In the case of thick Ethernet, the most commonly used MDI is a type of clamp that is installed directly onto the coaxial cable. For twisted-pair Ethernet, the MDI is an RJ-45 telephone-style jack that provides a connection to the four twisted-pair wires used to carry network signals in the twisted-pair media system.

The next device is called the medium attachment unit, or MAU. This device is called a transceiver in the original DIX Ethernet standard, since it both TRANSmits and reCEIVEs signals on the medium. The medium dependent interface is part of the MAU, providing the MAU with a direct physical and electrical connection to the medium.

Following the MAU is the attachment unit interface or AUI. This is called a transceiver cable in the DIX standard. The AUI provides a path for signals and power carried between the Ethernet interface and the MAU. The AUI may be connected to the Ethernet interface in the computer with a 15-pin connector. The computer itself is defined as data terminal equipment (DTE) in the IEEE standard. Each DTE is equipped with an Ethernet interface that performs the medium access control (MAC) functions.

And there we have it: the DTE contains an Ethernet interface which forms up and sends Ethernet frames that carry the data between computers attached to the network. The Ethernet interface is attached to the media system using a set of equipment that includes an AUI and a MAU with its associated MDI.

The MAU and MDI are specifically designed for each media type used in Ethernet. Coaxial MAUs differ from twisted-pair MAUs, for example, both in the technology used for the actual connection to the media (MDI), as well as the method used for sending Ethernet signals over the media and for detecting collisions.

Notice that in the figure above there are two kinds of DTE configurations shown — one with an external MAU and one with an internal MAU. With an external configuration the DTE contains only an Ethernet interface, and the AUI and MAU are both located outside the DTE. This is how a DTE looks when connected to a thick coaxial system using an external AUI cable and MAU.

However, it’s also possible for the MAU and AUI to be part of the network interface inside the DTE, with the only exposed device being the MDI that connects directly to the network media. This is the type of connection made in the thin coax and twisted-pair media systems. In this case, the AUI is nothing more than a set of wires on the interface board that link the Ethernet chips together.

To help make more sense of this alphabet soup let's look next at the Ethernet media types. We will also show a computer connected to segments of each media type. It should be emphasized that this is just a brief survey, and the descriptions of each media type do not contain all the information you need to correctly build large media systems.

4.0 Thick Ethernet — Type 10BASE5

The identifier “10BASE5” is one of a set of media identifiers that have been defined by the IEEE. As new media systems are developed, the IEEE creates a shorthand identifier for each system. The “10” refers to the speed of the system, which is 10 megabits per second. “BASE” refers to the signalling method known as baseband. This is the method used for the majority of Ethernet media types, and simply means that the Ethernet signals are the only signals being carried by that particular medium. The “5” refers to the maximum segment length allowed when multiplied by 100. In the case of thick Ethernet, each segment may be up to 500 meters in length.

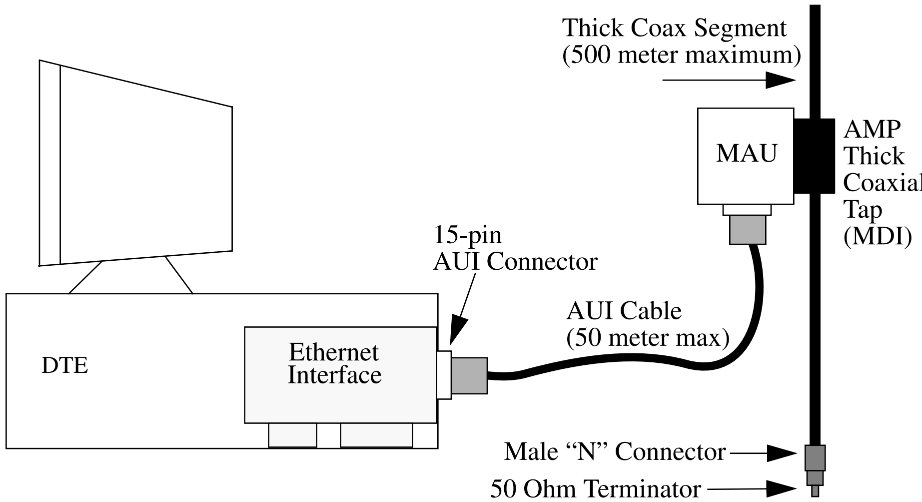

FIGURE 2. Connecting a computer to thick Ethernet.

4.1 Thick Ethernet Components

1. **Network medium**. The thick Ethernet system uses a thick (approx 0.4 inch diameter) and relatively inflexible coaxial cable. The cable may have plain PVC (yellow color) or Teflon^{TM} (orange-brown color) outer insulating jacket. Teflon is used for "plenum rated" cable, which is often required for installation in air handling spaces (also called plenums) to meet fire regulations. Thick Ethernet cable must be rated at 50 ohms characteristic impedance, and have a solid center conductor. Thick coaxial segments are equipped with male type "N" coaxial connectors at each end. Typical cable types are Belden 9880 (PVC) and 89880 (plenum rated).
2. **Terminator**. There must be a type-N 50 ohm terminator installed at each end of a thick coaxial cable segment.
3. **MAU (transceiver).** An Ethernet interface is attached to a thick Ethernet segment with an outboard MAU. There may be a maximum of 100 MAUs attached to a segment. The specifications state that each MAU connection to the thick coax must be separated by 2.5 meters of coaxial cable between it and the next MAU connection, and there are black bands printed on the thick coaxial cable to help maintain this spacing. The MAU spacing and the restriction on the number of MAUs are both designed to limit the amount of signal attenuation and distortion that can occur on a given cable segment.

   The most popular attachment mechanism (MDI) for a 10BASE-5 MAU is sold by AMP Corporation, and consists of a metal and plastic clamp that makes a direct physical and electrical connection to the coaxial cable. This clamp is also called a transceiver tap, since to install the clamp you must drill a hole into the thick coaxial cable in a process known as tapping the cable. Since this clamp may be installed while the network is active, it is also called a "non-intrusive" tap.

Another, much less popular, form of thick Ethernet MDI consists of a tap composed of two type-N coaxial cable connectors. Installing this tap requires cutting the thick coaxial cable, installing N connectors on each cable end, and then installing the tap as a sort of “barrel” connector in-line with the coaxial cable. Cutting the cable halts the operation of the network, earning this approach the label of “intrusive tap.”

Thick Ethernet MAUs are equipped with a male 15-pin connector to provide an attachment for the AUI cable. This connector has two locking posts, providing an attachment point for a sliding latch connector.

4. **AUI cable** (transceiver cable). An AUI cable is used to provide power to the MAU, and to carry signals between the MAU and the Ethernet interface. The AUI cable has a female 15-pin connector on one end that is equipped with a sliding latch; this is the end that is attached to the MAU. The other end of the AUI cable has a male 15-pin connector that is usually equipped with locking posts; this is the end that is attached to the Ethernet interface. Some 15-pin connectors on Ethernet interfaces are equipped with screw posts instead of the sliding latch fastener, requiring a specially equipped AUI cable.

   The standard AUI cable is a relatively thick wire (0.4 inch diameter) that may be up to 50 meters (164 feet) long. “Office grade” AUI cables are thinner (approximately 1/4 inch) and more flexible. Office grade AUI cables also have higher signal loss than standard AUI cables and consequently must be limited in length. One vendor of office grade cables rates them as having four times the amount of signal attenuation as standard cables, and only sells them in two and five meter lengths.
5. **Ethernet interface**. An Ethernet interface is a board that is installed in the DTE or is built into the DTE at the factory. Thick Ethernet interfaces have a female 15-pin connector equipped with a sliding latch for the AUI cable attachment.

These five components are sufficient to build a single thick Ethernet cable segment that can support up to 100 MAU attachments, and can be as much as 500 meters (1,640 feet) long.

The IEEE standard requires that individual segments be connected together with Ethernet repeaters. The repeater is a signal amplification device that keeps the system operating correctly by cleaning up and amplifying the signals that it repeats from one segment to the other. The repeater also has circuits that ensure that collisions that occur on any segment are propagated onto all other segments to which the repeater is attached. By doing this the repeater makes all segments function as though they were a single big segment, or what is known as a single Ethernet “collision domain.” This makes it possible for computers attached to any segment in a system of Ethernet segments linked with repeaters to hear the same signals and to operate as a single LAN.

A thick coaxial segment is known as a “mixing segment” in the multi-segment configuration guidelines. A mixing segment is formally defined as one which may have more than two MDI connections.

5.0 Thin Ethernet – Type 10BASE2

The thin Ethernet system uses a much more flexible cable that makes it possible to connect the coaxial cable directly to the Ethernet interface in the computer. In this connection scheme the AUI, MAU, and MDI are part of the network interface in the computer. This reduces the number of outboard components you need to purchase and install to connect a computer to the medium, thereby lowering the cost of an attachment to the network.

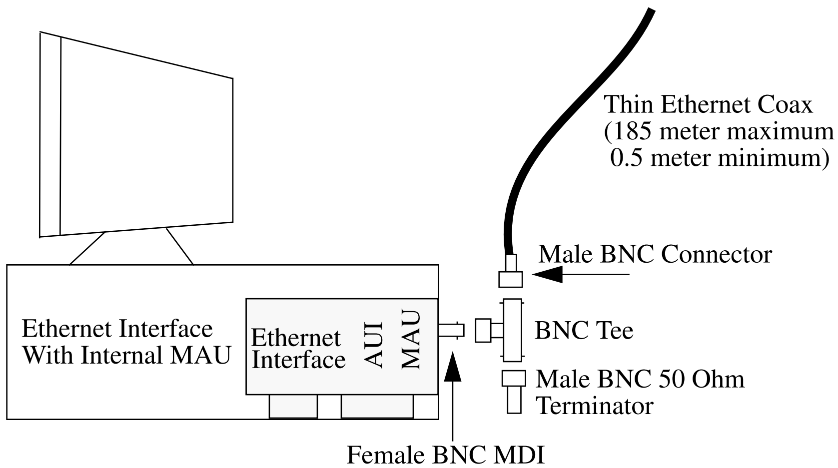

FIGURE 3. Connecting a computer to thin Ethernet.

5.1 Thin Ethernet Components

1. **Network Medium**. The thin Ethernet system is based on thinner (approx. 3/16th of an inch) coaxial cable that is more flexible and easier to deal with than the thick Ethernet variety. The cable must have a 50 ohm characteristic impedance rating and a stranded center conductor. These specifications may be met by cable types RG 58 A/U or RG 58 C/U, but cable vendors sometimes use these cable numbers to identify cables with different impedance ratings. It's up to you to make sure that the cable you purchase meets the specifications.

   Thin Ethernet cable sections must be equipped with male BNC connectors at each end. Segments may be a maximum of 185 meters in length, and not 200 meters as the rounded-up "2" in the shorthand identifier might lead you to believe. The standard requires that multiple segments be linked with repeaters.

2. **Terminators**. Each end of a complete thin Ethernet segment must be equipped with a 50 ohm terminating resistance. Multiport repeaters used for linking thin Ethernet segments often have internal 50 ohm termination provided on each port, which eases the task of terminating the end of the thin coax segment attached to the repeater. Some thin Ethernet repeaters have switchable termination that you can enable or disable, depending on your requirements. It is essential that no more than two 50 ohm

terminators be installed on a given segment, or the collision detection mechanism in the MAUs attached to the segment will not function correctly.

3. **Interface and MAU.** In the thin Ethernet system the MAU is built into the Ethernet interface, and therefore an external AUI cable is not required. The thin coax is flexible enough to be connected directly to the MDI on the interface. The thin Ethernet MDI is a female BNC connector. To make an attachment to a thin Ethernet segment this connector is attached to one end of a BNC Tee connector, so called because it looks like the letter "T." The other two ends of the BNC Tee make the physical and electrical connection to the thin Ethernet segment.

To help make the individual pieces clearer, the BNC connectors in the figure are shown unattached to one another. The thin Ethernet segment in the figure is drawn as terminating at this computer to show you how a thin Ethernet terminator is connected. However, a given thin Ethernet segment may also be connected to several computers in a topology known as "daisy chaining." In the daisy chain topology another piece of thin coax is connected to the BNC Tee, instead of a terminator, and this piece of coax is attached to the BNC Tee on the next computer in line. The BNC Tee at the very end of the segment is the only one that requires a terminator.

The thin Ethernet coaxial segment is defined as a mixing segment, since it can support more than two MDI connections. You may have up to 30 MAUs connected to each thin Ethernet segment. Each repeater connection requires a MAU, and must be counted toward the total of 30 MAU connections per segment. Since thin coaxial cable has higher signal attenuation than thick coax, the limit of 185 meters of cable helps ensure that signal losses are held to acceptable limits. The standard also recommends using high quality BNC connectors with low resistance gold plated center conductors.

The limit on the number of connections, and the recommendation of low resistance connectors is intended to help reduce the DC resistance caused by the coaxial connectors used in a thin Ethernet system. This, in turn, helps ensure that the total DC resistance of the segment is kept low enough so that the essential collision detect mechanism continues to work properly.

There are no special MAU spacing rules in the thin Ethernet media system. However the specifications state that the pieces of coaxial cable used to build a thin Ethernet segment may be no shorter than 0.5 meters (1.64 feet) in length. This effectively limits the minimum spacing between MAU connections to 0.5 meters.

Notice that the BNC Tee is connected directly to the BNC MDI on the interface. The standard notes that the length of the "stub" from the BNC MDI on the interface to the coaxial cable should be no longer than four centimeters, to prevent the occurrence of signal reflections which can cause frame errors.

6.0 Twisted-pair Ethernet – Type 10BASE-T

The twisted-pair Ethernet system operates over two pairs of wires, one pair used for receive data signals and the other pair used for transmit data signals. The two wires in each pair must be twisted together for the entire length of the segment, which is a standard technique used to improve the signal carrying characteristics of a wire pair.

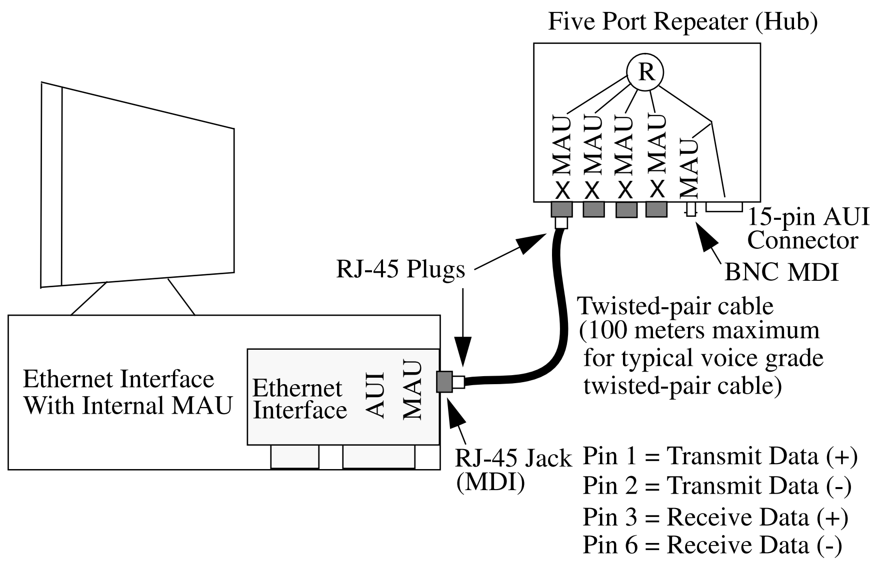

FIGURE 4. Connecting a computer to twisted-pair Ethernet.

6.1 Twisted-pair Components

1. **Network Medium.** The twisted-pair Ethernet system was designed to allow segments of up to 100 meters in length when using modern “voice grade” twisted-pair telephone wiring that meets the EIA/TIA Category Three wire specifications and follows the correct wiring scheme. The maximum segment length at your site may be shorter or longer than this depending on the quality of the twisted-pair cabling in your system.

   There are twisted-pair Ethernet cable testers available that allow you to check the electrical characteristics of the cable you use, to see if it meets the important electrical specifications. These specifications include signal crosstalk, which is the amount of signal that crosses over between the receive and transmit pairs, and signal attenuation, which is the amount of signal loss encountered on the segment.

   The transmit and receive data signals on a twisted-pair segment are polarized, with one wire of each signal pair carrying the positive (+) signal, and the other carrying the negative (-) signal. When connecting a twisted-pair interface to a repeater hub this polarity must be preserved, so that the positive terminal on one end of the segment is connected to the positive terminal on the other end.

2. **MAU.** When connecting two twisted-pair MAUs together over a segment, the transmit data pins of one MDI must be wired to the receive data pins of the other MDI, and vice versa. For a single segment connecting only two computers you can do this by building a special "crossover" cable, with the transmit pins on the RJ-45 plug at one end of the cable wired to the receive data pins on the RJ-45 plug at the other end of the crossover cable.

   However, when you are wiring multiple segments in a building it's much easier to wire the cable connectors "straight through" and not worry about whether the wires in the jumper cables or other twisted-pair cables in your building have been correctly crossed over. The way to accomplish this is to do the crossover wiring inside the multiport repeater hub. The twisted-pair Ethernet standard recommends this approach, and states that each port of the hub that is crossed over internally should be marked with an "X."

   Twisted-pair MAUs send a special link pulse to one another over the twisted-pair segment when the segment is idle. Vendors can provide a link light on the MAU and if the link lights on both MAUs are lit when you connect a segment, then you have an indication that the segment is working correctly.

Twisted-pair Ethernet segments are defined as link segments in the Ethernet specifications. A link segment is formally defined as a point-to-point full duplex medium that connects two and only two MDIs. The phrase "full duplex" means that there are separate signal paths for sending and receiving data. The smallest network built with a link segment would consist of two computers, one at each end of the link segment.

The more typical installation uses multiport repeaters, also called "hubs" or "concentrators," to provide a repeater connection between a larger number of link segments. You connect the MAU in the Ethernet interface in your computer to one end of the link segment, and the other end of the link segment is connected to the MAU in the repeater hub. That way you can attach as many link segments with their associated computers as you have hub ports, and the computers all communicate via the repeater hub.

In any twisted-pair Ethernet system with more than two computers, you need a multiport repeater hub to connect the individual segments together, and a five-port hub is shown in the figure. Four of the ports are equipped with twisted-pair MAUs and twisted-pair RJ-45 jacks as MDIs. The fifth port may be connected either to a thin Ethernet segment, or to an outboard MAU using the 15-pin AUI connector.

A common error when connecting a computer to a twisted-pair segment is to use the widely available "silver satin" patch cable typically used to connect telephones to the telephone jack on the office wall. The problem is that the silver satin patch cable for telephones does not have twisted wire pairs in it, and the lack of twisted pairs results in excessive signal crosstalk and "phantom collisions." This occurs because collisions are detected in twisted-pair Ethernet by the simultaneous occurrence of signals on the transmit and receive wire pairs, and excessive crosstalk can trigger the collision detect circuit. This problem can be avoided by using only twisted-pair patch cables rated for use in twisted-pair Ethernet systems to make a connection between the MAU in the computer or the hub and the rest of segment.

7.0 Fiber Optic Ethernet – Types FOIRL and 10BASE-F

The fiber optic media system use pulses of light instead of electrical currents to send signals, which provides electrical isolation for equipment at each end of a fiber link. The electrical isolation provides immunity from the effect of lightning strikes and the different ground currents found in separate buildings. This is essential when segments must travel outside a building to link separate buildings.

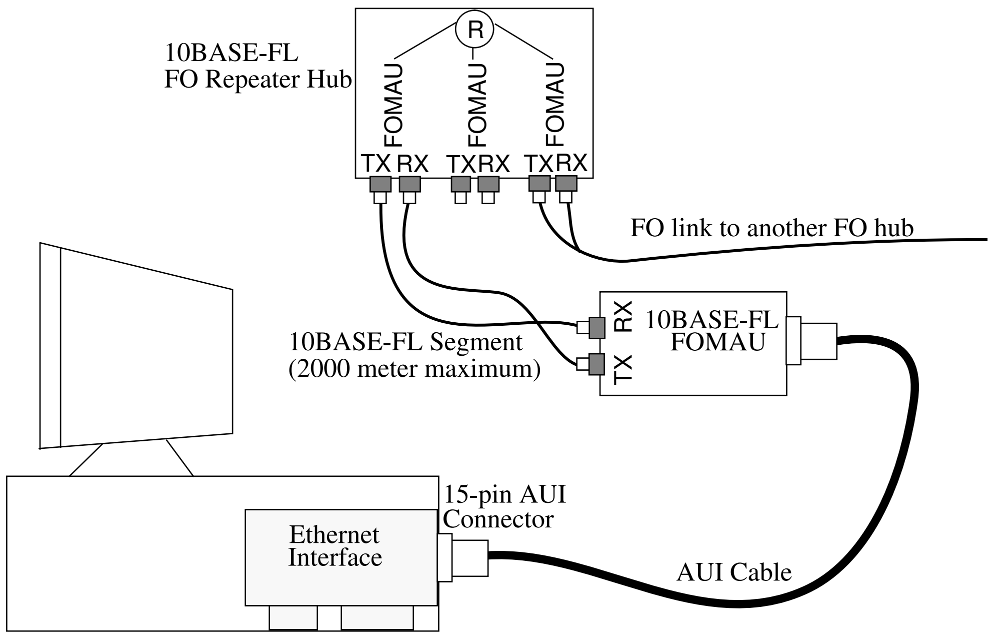

FIGURE 5. Connecting a computer to a 10BASE-FL segment.

The most commonly used fiber optic medium type is the link segment. There are two fiber optic link segments in use today, the original Fiber Optic Inter-Repeater Link (FOIRL) segment, and the newer 10BASE-FL segment. The original FOIRL specification provided a link segment of up to 1000 meters between two repeaters only. As the cost of repeaters dropped and more and more multiport repeater hubs were used, it became cost-effective to link individual computers to a fiber optic port on a repeater hub. Vendors created outboard FOIRL MAUs to allow this, although a repeater-to-DTE fiber connection was not specifically described in the FOIRL standard.

To deal with this and other aspects of fiber optic Ethernet, a new set of fiber optic media standards, called 10BASE-F, was developed. This new set of standards includes revised specifications for a fiber optic link segment that allow direct attachments to computers. The 10BASE-F specifications include the following three segment types.

- **10BASE-FL** The new Fiber Link specifications. The 10BASE-FL standard replaces the older FOIRL specifications, and is designed to interoperate with existing FOIRL-based equipment. 10BASE-FL provides for a full duplex fiber optic link segment that may be up to 2000 meters long providing that only 10BASE-FL equipment is used in the segment. If 10BASE-FL equipment is mixed with FOIRL equipment, then the maximum segment length may be 1000 meters. A 10BASE-FL segment may be attached between two computers, or two repeaters, or between a computer and a repeater port. Because of the widespread use of fiber links, 10BASE-FL is the most widely used portion of the new 10BASE-F fiber optic specifications.

- **10BASE-FB** The Fiber Backbone link segment system. The 10BASE-FB specifications describe a special synchronous signaling backbone approach that allows the limit on the number of repeaters that may be used in a given Ethernet system to be exceeded. 10BASE-FB links do not attach to computers or end nodes, and are only used to link special 10BASE-FB repeater hubs together in a large repeated backbone system. 10BASE-FB links may be up to 2000 meters in length.

- **10BASE-FP** The Fiber Passive system. This provides a set of specifications for a fiber optic mixing segment that links multiple computers on a fiber optic media system without using repeaters. 10BASE-FP segments may be up to 500 meters long, and a 10BASE-FP fiber optic passive star coupler typically links up to 33 computers.

In the figure we show a computer linked to a repeater hub with a 10BASE-FL segment. The computer is equipped with an Ethernet interface that has a 15-pin AUI connector. This connector allows us to make a connection to an outboard fiber optic MAU (FOMAU), using a standard AUI cable. The FOMAU, in turn, is connected to the repeater hub with two strands of fiber optic cable. Another port on the repeater is shown connecting to another FO repeater hub, which could be located some distance away. A major advantage of the fiber optic link segment is the long distances that it can cover.

8.0 Universal 15-pin Connector

Note that the 15-pin AUI connector on an interface is a “universal” connector that makes it possible to connect the interface to any Ethernet segment type, for the small additional cost of an outboard MAU. In the last figure we saw a 15-pin AUI connection to an outboard 10BASE-FL MAU. You can make the same sort of connection to a thin or twisted-pair segment.

You can attach a 15-pin AUI connector to a thin Ethernet segment, for example, by using an external MAU equipped with a thin Ethernet BNC MDI. The MAU with its BNC MDI is attached directly to a BNC Tee on the thin Ethernet coax, and the 15-pin AUI connector on the MAU is connected to the 15-pin AUI connector on the Ethernet interface with an AUI cable. With a small enough MAU you can even eliminate the AUI cable, and connect the 15-pin connector of the MAU directly to the 15-pin connector on the Ethernet interface of the computer.

Now that we've seen what the Ethernet media varieties look like, let’s look at the guidelines used for building a multi-segment Ethernet system with these varieties. The next part of this guide describes one of the models provided by the IEEE for multi-segment configuration.

9.0 Configuration Rules: Transmission System Model 1[1]

Section 13 of the IEEE 802.3 standard provides two models for verifying the configuration of multi-segment Ethernets. The version of section 13 shown in this guide was published in November, 1993, along with the 10BASE-F specifications. The configuration model shown here is called Transmission System Model 1. It consists of a set of basic configuration rules that can be applied to most Ethernets. The second model provides a set of calculations that you can use to verify more complex Ethernet topologies, and is described in the **Ethernet Configuration Guide.**

In the rule-based configuration model shown here, a set of multi-segment configuration rules are provided for combining Ethernet segments based on conservative calculations for the components involved. You shouldn't let the fact that these configuration rules are based on conservative calculations lead you to believe that you can bend the rules and always get away with it. There isn't a lot of engineering margin left in maximum-sized Ethernets, despite the allowances made in the standards for manufacturing tolerances and equipment variances. If you want guaranteed performance and reliability, then you need to stick to the published guidelines.

The multi-segment configuration rules are as follows:

1. **Repeater sets are required for all segment interconnection.** The repeaters used must comply with all IEEE specifications in section 9 of the 802.3 standard, and do signal retiming and reshaping, preamble regeneration, etc. If you do not use true IEEE 802.3 repeaters for all segment interconnections, then your Ethernet system cannot be verified using either configuration model.
2. **MAUs that are part of repeater sets count toward the maximum number of MAUs on a segment.** Thick Ethernet repeaters typically use an outboard MAU to connect to the thick Ethernet coax. Thin coax and twisted-pair repeater hubs use internal MAUs located on each repeater port.
3. **The transmission path permitted between any two DTEs may consist of up to five segments, four repeater sets (including optional AUIs), two MAUs, and two AUIs**. The repeater sets are assumed to have their own MAUs, which are not counted in this rule.
4. **AUI cables for 10BASE-FP and 10BASE-FL shall not exceed 25 m. (Since two MAUs per segment are required, 25 m per MAU results in a total AUI cable length of 50 m per segment).**
5. **When a transmission path consists of four repeaters and five segments, up to three of the segments may be mixing and the remainder must be link segments. When five segments are present, each fiber optic link segment (FOIRL,**

1. Portions of the information contained herein are copyrighted information of the IEEE extracted from IEEE Std 802.3j-1993, IEEE Standard for Local and Metropolitan Area Networks – Fiber Optic Active and Passive Star-Based Elements, Type 10BASE-F, copyright © 1993 by the Institute of Electrical and Electronics Engineers, Inc. The IEEE takes no responsibility for damages resulting from the reader's misinterpretation of said information resulting from the placement and context in this publication. Information is reproduced with the permission of the IEEE.

10BASE-FB, 10BASE-FL) shall not exceed 500 m, and each 10BASE-FP segment shall not exceed 300 m.

6. **When a transmission path consists of three repeater sets and four segments, the following restrictions apply**:

   a. The maximum allowable length of any inter-repeater fiber segment shall not exceed 1000 m for FOIRL, 10BASE-FB, and 10BASE-FL segments and shall not exceed 700 m for 10BASE-FP segments.

   b. The maximum allowable length of any repeater to DTE fiber segment shall not exceed 400 m for 10BASE-FL segments and shall not exceed 300 m for 10BASE-FP segments and 400 m for segments terminated in a 10BASE-FL MAU.

   c. There is no restriction on the number of mixing segments in this case. In other words, when using three repeater sets and four segments, all segments may be mixing segments if desired.

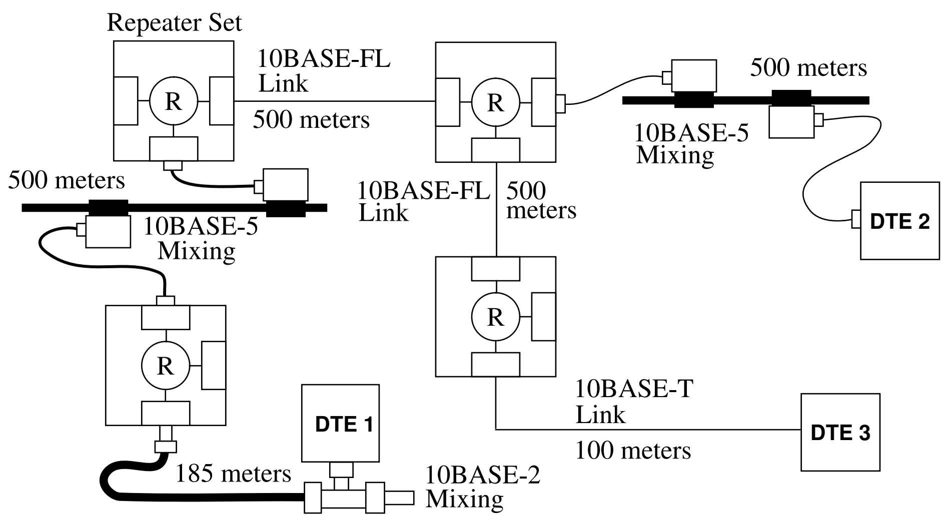

FIGURE 6. One possible maximum configuration.

The figure shows an example of a maximum Ethernet configuration that meets the rules. The maximum packet transmission path in this system is between DTE1 and DTE3, since there are four repeaters and five segments in that path. Two of the segments are mixing segments, and the other three are link segments.

While the configuration guidelines emphasize the maximum limits of the system, you should beware of stretching things as far as they can go. Ethernets, like many other systems, work best when they are not being pushed to their limits.

10.0 Non-standard Ethernet Equipment

The Ethernet standard describes the minimum set of devices needed to make the system work. There's a small set of basic components whose rules of operation and electrical characteristics are completely specified in the standard. The components include the Ethernet interface, AUI cable, MAU, repeater, and the physical medium.

By building and selling standard Ethernet equipment, vendors can be assured that their devices will operate correctly when attached to any properly configured Ethernet system. As long as your Ethernet system is built using devices and media systems that fully comply with the standards you can use the configuration models to verify that your Ethernet system will operate correctly.

However, there are some non-standard devices designed by vendors to be used in the packet transmission path of an Ethernet media system. By being part of the packet transmission path these devices end up being part of the signal timing that is essential for correct Ethernet operation.

Since these devices are not in the standard and are not covered by any configuration guidelines developed by the IEEE, it's difficult to state exactly what the impact of using such a device will be on your network. It may work, or it may not, depending on the total size of your network, the number of computers attached to each network segment, etc. The performance of non-standard devices varies, and each vendor typically has their own special rules for the operation of their non-standard device. The most commonly used non-standard devices are the multiport transceiver and the media converter.

10.1 Multiport Transceivers

Multiport transceivers were developed when thick Ethernet was the only media type available, and network designers needed a way to concentrate a set of machines in a small space. That’s why you will sometimes see these devices referred to as “transceiver concentrators.”

The problem arises because the thick Ethernet standard requires that each MAU attachment be separated by 2.5 meters of cable from the next MAU attachment. This meant that when you needed to connect a number of machines located in the same space to the network, you had to coil up enough thick Ethernet coax in a wiring closet or under a machine room floor to provide cable to meet the 2.5 meter MAU spacing requirement. By providing several (usually eight) 15-pin AUI connectors in a single multiport transceiver, vendors made it easier to connect groups of computers to thick Ethernet.

However, each multiport transceiver adds a certain amount of delay and other effects to the signals that pass through it, and these effects may vary depending upon which vendor built the multiport transceiver. Since the multiport transceiver is not defined in the IEEE 802.3 standard the extra bit times of delay and other effects it may cause are not included in the configuration guidelines, and your system cannot be verified using the IEEE guidelines. If you use multiport transceivers you should read the vendor’s configuration guidelines and follow them carefully. Even then you may find that multiport transceivers may not perform well in large networks, or when attached to stations with maximum-length AUI cables.

10.2 Media Converters

The IEEE standard states that repeaters must be used to link all segment types. It used to be the case that buying repeaters could add a significant cost to a network design, and some vendors attempted to deal with this by offering a lower cost device called a media converter. These devices are also called media adapters and media translators, depending on the vendor.

Media converters are designed to link segments together inexpensively without using a repeater. While they provide some of the signal amplification functions of a repeater, they do not contain the more expensive circuits used by a repeater to retime signals, rebuild the preamble on the Ethernet packet, partition (isolate) the segment in case of errors, and so on.

The lack of these more expensive circuits explains why media converters were a lower-cost approach to linking segments than repeaters. However, the cost differential between media converters and true 802.3 repeaters has been dropping ever since low-cost repeater chips became more widely available due to the popularity of the twisted-pair Ethernet system.

Therefore, there is little economic reason to use a media converter to link segments. This is especially true when you consider that any Ethernet system that includes media converters cannot be evaluated using either configuration model, since the media converter is not part of the standard set of equipment defined in the Ethernet specifications and included in the configuration rules.

To make sure that your network meets the specifications in the standards and to make it possible to evaluate your network using the configuration rules, you must use true IEEE 802.3 repeaters for all segment interconnections. If in doubt when buying a device that links segments together, ask the vendor to verify that what they are selling is a true IEEE 802.3 repeater, and that it meets all of the specifications in section 9 of the 802.3 standard.

11.0 Network Design Guidelines

This document is written to provide a quick guide to the baseband Ethernet media types and a description of the rule-based configuration model provided by the IEEE. Network design and implementation issues are not covered in any detail.

The media descriptions shown in this guide can provide a useful overview, but they cannot provide the detailed information required for larger networks. Installing a small Ethernet can be as simple as buying a twisted-pair hub and some patch cables, for example, and connecting all of your computers to the hub. But in larger systems, the issues of structured cabling systems, what kind of media to use, and which media system will best provide for future growth, can be much more complex.

With regard to the configuration guidelines shown in this guide, note that while the guidelines describe how far you can stretch things, this should not be taken to mean that a good network design should push things to their limits. The design of a multi-lane highway can be a useful analogy, since a highway is somewhat like a LAN in that it is a multiple access system with a “shared channel” whose traffic increases and decreases over a 24 hour period.

When highway engineers design a multi-lane limited access highway they calculate the maximum number of vehicles that can be accommodated given the number of lanes and entrance and exit ramps, topology of the roadway, and so on. Like a LAN, a highway system has a theoretical maximum performance, but you do not want to push the system to its limits. After all, a highway loaded with bumper-to-bumper traffic during the homeward-bound commute is still operating within design limits, but no one is very happy with it.

In much the same way a LAN stretched to its limits with a large number of computers on maximum-length segments can end up loaded to its capacity and still be within the specifications and working properly, but no one will be happy with it. When you design a LAN system, you should not focus on what you can get away with or how far you can push the system. Instead, you should consider how many machines you need to support, how much traffic they will generate, and then build a system that will be able to accommodate the load without serious congestion.

Each network design is a special case, since every group has a different mix of computing equipment, and different computing requirements. In general, the UTnet Networking Services group recommends that conservative design practices be used to help deal with network management issues and network traffic growth. Your network designs should emphasize modular cabling systems and network topologies that can be easily reconfigured and upgraded when traffic growth demands more bandwidth.

We typically recommend that campus groups use twisted-pair Ethernet for new installations, based on twisted-pair hubs centralized in one or more wiring closets. Locating equipment in a wiring closet reduces the number of places you must visit when tracking down a problem. Also, as new network equipment becomes available, you can upgrade your centralized hub equipment as required to improve the capabilities of your network system. Installing high quality twisted-pair wiring to each desktop is another way to provide a network system that is reliable and easily managed, and that can be upgraded to higher speeds in the future.

We further recommend that campus groups use only standard IEEE 802.3 media types and equipment when building Ethernets, to make sure that their networks meet the configuration guidelines and to provide interoperability. With respect to standards you should beware of vendor claims. Ethernet has become a commodity market, and there are a number of vendors selling their own inventions which are not described in any IEEE specification. This issue is often further obscured by vendor claims that their proprietary technologies are “compatible” with IEEE equipment. When in doubt, ask the vendor to provide the exact IEEE standards and specifications that apply to the equipment they are selling.