

An Introduction to Industrial Ethernet

When you talk about office and home networking, usually you are talking about Ethernet-based networks—computers, printers and other devices that contain Ethernet interfaces connected together via Ethernet hubs, switches and routers. In the industrial area the networking picture is more complex, but as time goes on Ethernet is becoming a bigger part of that picture. This article is an introduction to the basics of Ethernet, with a bit of added detail on how it is beginning to fit into the industrial networking picture.

Ethernet's Roots

Although Xerox's Bob Metcalfe sketched the original Ethernet concept on a napkin in 1973, its inspiration came even earlier. ALOHAnet was a wireless data network created to connect together several widely separated computer systems on Hawaiian college campuses (different islands). The challenge was to enable several independent data radio nodes to communicate on a peer-to-peer basis without interfering with each other. ALOHAnet's solution was a version of the carrier sense, multiple access with collision detection (CSMA/CD) concept. Metcalfe based his Ph.D. work on finding improvements to ALOHAnet. This led to his work on Ethernet.

Ethernet, which later became the basis for the IEEE 802.3 network standard, specifies physical and data link layers of network functionality. The physical layer specifies the types of electrical signals, signaling speeds, media and connector types and network topologies. The data link layer specifies how communications occurs over the media—using the CSMA/CD technique mentioned above—as well as the frame structure of messages transmitted and received.

Ethernet Physical Layer

In the early days Ethernet options were more limited than they are today. Two common options were 10Base2 and 10Base5 configurations. Both operated at 10 Mbps and used coaxial cable with nodes connected to the cable via Tee connectors, or through 'attachment unit interfaces' (AUI) in a multidrop bus configuration. 10Base2 networks allowed segment lengths of up to 185 feet using RG 58 coaxial cable (also called Thin Ethernet). 10Base5 offered greater distances between nodes but the thick coaxial cable and 'vampire tap' connections were bulky and difficult to work with. Later, another solution in this speed category was 10Base-FL, which uses fiber optic media and provides distances greater than 2000 feet.

Another early 10 Mbps physical layer option—10Base-T—quickly gained popularity because it was easier to install and used inexpensive unshielded twisted pair (UTP) Category 3 cable. Nodes (typically computers with network interface cards, or NICs) were connected in a star topology to a hub, which in turn was connected to other network segments. Each computer had to be less than 100 feet from the hub. Standard RJ-45 connectors were used.

In the mid-1990s 100 Mbps Ethernet equipment became available, increasing the data transfer rate significantly. NICs that would automatically adjust to operate at 10 Mbps or 100 Mbps made migration to the faster standard simple.

Today, virtually all computer network interface cards implement 100Base-TX. Category 5e UTP cable is the standard cable used with 100Base-TX and cable lengths are the same as for 10Base-T networks. Coaxial-based networks are increasingly being replaced with fiber optic media, especially for point-to-point links. For example, 100Base-FX uses two optical fibers and allows full duplex point-to-point communications up to 2000 feet. Gigabit Ethernet (1000 Mbps) options also are available using twisted pair and fiber optic media.

(For more information on Ethernet physical layer options see the Ethernet Basics article on the B&B Electronics website at: www.bb-electronics.com)

Data Link Layer

Ethernet's data link layer defines its media access method. Half-duplex links, such as those connected in bus or star topologies (10/100Base-T, 10Base2, 10Base5, etc), use carrier sense, multiple access with collision detection (CSMA/CD). This method allows multiple nodes to have equal access to the network, similar to early party-line telephone systems in which users listened for ongoing conversations and waited until the line was free before accessing the line. All nodes on an Ethernet network continuously monitor for transmissions on the media. If a node needs to transmit it waits until the network is idle, then begins transmission. While transmitting, each node monitors its own transmission and compares what it 'hears' with what it is trying to send. If two nodes begin transmitting at the same time, the signals will overlap, corrupting the originals. Both nodes will see a different signal to that which they are trying to send. This is recognized as a 'collision'. If there is a collision, each node stops transmitting and only attempts to re-transmit after a preset delay, which is different for each node.

This method of media access makes it simple to add to or remove nodes from a network. Simply connect another node and it begins to listen and transmit when the network is available. However, as the number of nodes grows and the volume of traffic from each node increases, opportunities to gain access to the network decrease. As utilization increases, the number of collisions increases exponentially and the probability of getting access within a given length of time decreases dramatically. This characteristic makes Ethernet a probabilistic network, as opposed to a deterministic network, in which access time can be reliably predicted. (Master/slave and token passing network schemes are deterministic.)

Full-duplex point-to-point Ethernet links (10Base-FL, 100Base-FX, etc) collisions are not an issue, since only two nodes are present and separate send and receive channels are available. Another advantage is that data can be sent in both directions simultaneously, effectively doubling the data transfer rate.

The Ethernet Frame

The Ethernet data link layer also defines the format of data messages sent on a network. The data message format, or frame, contains several fields of information in addition to the data to be transferred across the network.

Obviously, at the heart of the message is the actual data that is to be sent. This is called the 'data unit'. Ethernet data units can contain between 46 and 1500 eight-bit bytes of binary information. The actual length of the data unit is determined and included in the message as a field to tell the receiver how to determine which part of the message is data. Each message must include source and destination addresses so that other nodes can determine where the message is coming from and going to. These six-byte binary numbers are called MAC addresses. Every Ethernet node has a unique MAC address permanently stored in its hardware memory. The Ethernet frame also contains a four-byte 'frame check sequence' (FCS) field which is a binary number generated by the sending node that allows high reliability cyclic redundancy checking (CRC) error checking to be done by the receiving node.

Hubs and Switches

Ethernet hubs are simple physical layer devices used with 10/100Base-T(X) networks to repeat and split Ethernet signals. Nodes connect to ports on the hub as branches to create a physical star topology. Hubs receive data from the connected nodes, regenerate it and send it out on all other ports. By regenerating the data the maximum segment distance can be extended. All transmissions go to all the connected nodes, the same as on a bus network. Nodes respond to transmissions based on the destination address contained in the message frame. Hubs allow all wiring to connect to a central location making it easy to isolate problem nodes and make changes to the network.

Switches are similar to hubs except that they divide the network into segments. An internal table is maintained of the destination addresses of the nodes connected to the switch. When an Ethernet packet is received at one of the switch's ports the destination address in the packet is read, a connection is made to the appropriate port and the packet is sent to that node. This isolates the message traffic from the other nodes, decreasing the utilization on the overall network. Ethernet switches can be managed or unmanaged. Unmanaged switches operate as described above. Managed switches allow advanced control of the network. They include software to configure the network and diagnostic ports to monitor network traffic.

Higher Level Network Functions

To facilitate reliable communications across multiple, and in some cases dissimilar networks, other higher-level protocols are used on top of Ethernet's data link layer. The most common of these today, especially when connecting an Ethernet network to the Internet, is TCP/IP. IP, or internet protocol, ensures packets are moved across the network based on their IP address. TCP, or transport control protocol makes sure data is delivered completely and error-free. Two or more Ethernet networks may be connected together via a router, a device that maintains a list of IP addresses on each network connected to it. The router monitors the IP addresses on packets received at its ports and routes them to the port connected to the appropriate network.

Ethernet and Industrial Systems

Ethernet's simple and effective design has made it the most popular networking solution at the physical and data link levels, especially for general purpose office networks. With high speed options and a variety of media types to choose from Ethernet is efficient and flexible. Using inexpensive UTP cable and star topology, and CSMA/CD media access, Ethernet networks are easily designed and built. Nodes can be added or removed simply and troubleshooting is relatively easy to do. As Ethernet and related technologies have become prevalent in the general networking arena a large base of trained personnel has become available.

These factors, and the low cost of Ethernet hardware, have made Ethernet an attractive option for industrial networking applications. Also, the opportunity to use open protocols such as TCP/IP over Ethernet networks offers the possibility of a level of standardization and interoperability that has until now remained elusive in the industrial field.

However, the probabilistic nature of Ethernet is one characteristic that is a drawback for some industrial network applications. Historically, time critical networking applications have been handled using deterministic networks (using master/slave or token passing schemes). Utilization levels on industrial Ethernet networks must be carefully controlled as levels greater than 10% often result in inadequate performance. Still, as the overall cost/benefits of Ethernet have increased, industrial users have found ways to enhance Ethernet's data transfer performance. One method is to segment networks using switches and routers to minimize unwanted network traffic and reduce utilization. Another is

to use newer, higher level protocols that incorporate prioritization, synchronization and other techniques to ensure timely delivery of messages.

The result has been an ongoing shift toward the use of Ethernet for industrial control and automation applications. Ethernet is increasingly replacing proprietary communications at the plant floor level and in some cases moving downward into the cell and field levels.

Most major control system manufacturers now incorporate versions of Ethernet networks and higher-level Ethernet-related protocols into their product offerings. Often, several manufacturers and/or industry stakeholders have entered into cooperative efforts to develop Ethernet-related standards and products. Several other these now exist, though interoperability between them continues to be elusive.

EtherCAT (Ethernet for Control Automation Technology) is an open real-time Ethernet network developed by Beckhoff. It provides real-time performance, features twisted pair and fiber optic media and supports various topologies. It is supported by the EtherCAT Technology Group, which has 168 member companies.

Ethernet Powerlink is a real-time Ethernet protocol that combines the CANopen concept with Ethernet technology. The Ethernet Powerlink Standardization Group (EPSP) is an open association of industry vendors, research institutes and end-users in the field of deterministic real-time Ethernet.

EtherNet/IP is an industrial networking standard that takes advantage of commercial off-the-shelf Ethernet communications chips and physical media. The IP stands for 'industrial protocol'. ControlNet International (CI), the Industrial Ethernet Association (IEA) and the Open DeviceNet Vendor Association (ODVA) support it.

Modbus-TCP, supported by Schneider Automation, allows the well-proven Modbus protocol to be carried over standard Ethernet networks on TCP/IP.

PROFINET is Profibus' Ethernet-based communication system, currently under development by Siemens and the Profibus User Organization (PNO).

The ongoing level of interest, activity and new product introductions of Ethernet-based equipment suggests industrial use of Ethernet will continue to grow for the foreseeable future.

#####