Volume 7 • Issue 6
NOVEMBER-DECEMBER 2006

the EXTENSION
A Technical Supplement to Control Network

Incorporating Media Converters

By George Thomas, Contemporary Controls

Introduction

Media converters are devices that allow for interconnecting two different cabling technologies. With Ethernet there are three types of cables—coaxial, twisted-pair and fiber optic. With modern Ethernet systems, the most common need is to connect a twisted-pair (copper) segment together with a fiber optic segment and there are media converters available to do the job. Fiber optics offer greater distances, galvanic isolation, and immunity to high levels of electromagnetic interference and lightning strikes. Product selection would seem to be easy—just make sure that connectors are compatible and select the lowest price. However, not all media converters operate the same and supplying a media converter requires some planning since an improperly installed media converter will degrade a system or simply will not function. This document reviews the issues of applying media converters to an Ethernet system.

What is an Ideal Media Converter?

The ideal media converter would provide a transparent conversion from one medium to another without introducing data latency. However, conversion electronics within devices introduce data latency. Physical layer devices such as repeating hubs introduce little data latency while switching hubs introduce much more due to the store-and-forwarding of frames. A media converter is a physical-layer device and therefore should introduce little data latency. Another trait of a media converter is transparency. Ideally, all data generated on the copper side should appear on the fiber side. The same should be said of data generated on the fiber side—it should faithfully appear on the copper side. This applies to special signals such as Link Integrity bursts and support for the Auto-Negotiation protocol. Media converters need to support these special functions as well. As a transparent device, the data rate on one side must be the same as the data rate on the other side. In other words, there can be no buffering of data within an ideal media converter.

A media converter does more than simply provide a compatible physical connection between two types of cabling technologies. The signaling on the two media could be different so active electronics, and an associated power source, are needed to make this signaling conversion. This also means that media converters are unique to the networking technology involved. Ethernet, ARCNET® and Controller Area Network media converters are all different and, therefore, incompatible with one another. We will address the unique requirements of Ethernet media converters that convert twisted-pair cabling to that of fiber optic cabling since that is the most popular need. Twisted-pair to fiber optic conversion can occur in a repeating hub or switching hub, but operation within a hub differs from a stand-alone media converter—as will be discussed later. We will start by reviewing the twisted-pair and physical-layer standards for Ethernet. We will focus our discussion on 10 Mbps and 100 Mbps, and twisted-pair and fiber optic cabling.

10BASE-T

This is the original twisted-pair standard for Ethernet with operation at 10 Mbps using two twisted-pairs.

A media converter allows for a simple cnversion from one medium to another such as twisted-pair to fiber.

Two pairs allow for full-duplex operation when a switching hub is employed. With a repeating hub, only half-duplex operation is possible. Manchester encoding is used over the wires and cable length is limited to 100 m. Category 3 (CAT 3) or better cable must be used on 10 Mbps, Manchester-encoded systems. The 10BASE-T standard introduced the Link Integrity function to verify that a valid link exists. This is accomplished using a special signal called the Normal Link Pulse (NLP) that is sent in the absence of data. RJ-45 connectors are used along with a standardized pin-out.

100BASE-TX

This is the Fast Ethernet twisted-pair standard. When we say Fast Ethernet, we mean 100 Mbps operation. While the data rate increased ten-fold, the maximum cabling length remained at 100 m and the number of pairs remained at two. The RJ-45 connector and pin-out remained the same as 10BASE-T—allowing for the potential for a migration path from 10 Mbps to 100 Mbps while retaining the existing infrastructure.

This would be true if the cabling had adequate bandwidth for 100 Mbps operation.

Instead of CAT 3 cable, at least CAT 5 cable is required. To increase the data rate ten-fold while retaining the maximum cabling length, encoding became much more sophisticated. A multi-level signaling scheme using three logic levels (MLT-3) replaced the two-level scheme of 10BASE-T. Not only are the signal levels unlike, but the 4B/5B encoding scheme is completely different from that of Manchester encoding. The Fast Ethernet standard introduced the Auto-Negotiation protocol which defined a means for low-speed devices to communicate with high-speed devices upon initial connection—thus allowing support for legacy devices. A Fast Link Pulse (FLP) replaced the NLP.

10BASE-FL

Although not the first fiber optic standard for Ethernet, this standard became the most popular for 10 Mbps Ethernet. A single pair of 62.5/125 µm multimode fiber optic cabling was specified using ST connectors—thus allowing full-duplex operation when using switching hubs. Like 10BASE-T, Manchester encoding was employed, but cabling length over fiber can extend to 2 km (demonstrating one of the advantages of using fiber). At each end of the link are separate transmitter (TX) and receiver (RX) ports. To implement the crossover function, cable connections are reversed at either end of the link. Unlike 10BASE-T, there is no NLP. Instead, a 1 MHz carrier is transmitted during idle conditions.

100BASE-FX

This standard was created to match the newer 100 Mbps operation of 100BASE-TX but over fiber optics. Although the same single pair of 62.5/125 µm multimode fiber optic cabling was used, a new connector was introduced called the SC connector. To maintain the same maximum cabling distance of 2 km, a critical decision was made. Instead of maintaining the lower cost 850 nm (short wave) transceivers, the 100BASE-FX standard specified 1300 nm (long wave) transceivers because of the lower signal attenuation at this frequency. Transceivers operating at different frequencies cannot communicate with one another. This decision eliminated the possibility of a 100BASE-FX device communicating with a legacy 10BASE-FL device. Instead of Manchester encoding, non-return to zero inverted (NRZI) encoding was used in combination with the same 4B/5B encoding scheme used on 100BASE-TX.

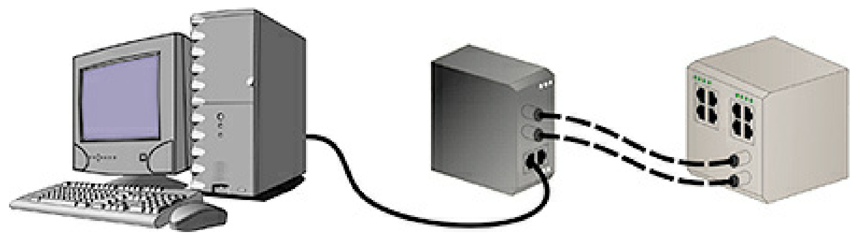

Media converters usually fit between an end station and a port on a hub. However, they can also fit between to hub ports.

100BASE-SX

The "fiber to the desktop" movement embraced the latest 100 Mbps fiber optic standard that is based upon short wave fiber optic technology. This standard is basically the same as the 100BASE-FX except that 850 nm transceivers and ST connectors were specified providing a possibility of connecting new 100 Mbps fiber optic equipment to legacy 10 Mbps fiber optic links. Instead of achieving a 2 km distance, 100BASE-SX is limited to 300 m. For desktop applications, this should be adequate. Like 100BASE-TX, the new 100BASE-SX standard supports auto-negotiation. The intent was to achieve the same degree of communication with fiber legacy devices as with copper legacy ones.

Auto-Negotiation Protocol

With equipment running at both 10 Mbps and 100 Mbps, a means was sought that would allow end stations and switching hubs to automatically adapt to one another. The result was the Auto-Negotiation protocol. With auto-negotiation, not only is communication speed determined, but duplex mode is determined as well. Each device that participates in the Auto-Negotiation protocol knows its own capabilities and "advertises" these attributes to its link partner upon initial connection. When connecting two devices over a twisted-pair link, each is to advertise its available speeds and its duplex capability (half or full). A duplex connection allows communication in both directions. For full-duplex links, bi-directional communication can occur simultaneously. For half-duplex, bi-directional communication can occur but not at the same time. Full-duplex links offer higher throughput but cannot reinforce collisions and therefore can only be used with switching hubs. Half-duplex links can work with either switching or repeating hubs. For the range of devices we are discussing here, there are four possible conditions of interest:

100 Mbps, full-duplex

100 Mbps, half-duplex

10 Mbps, full-duplex

10 Mbps, half-duplex

The four possible conditions have been ranked with the highest performance at the top of the list with the lowest at the bottom. Notice that preference has been given to full-duplex links since they offer higher performance. Of course, 100 Mbps offers higher performance over a 10 Mbps link.

During negotiation, the two partners advertise their possible levels of performance and agree upon the highest common denominator for optimum communication. If one device advertised 10 Mbps, full-duplex and the other 100 Mbps, full-duplex, they would settle on 10 Mbps, full-duplex, and begin communicating accordingly. Devices that participate in auto-negotiation are end stations and switch ports. Repeating hubs do not participate. What happens if one of the two devices was manufactured before the Auto-Negotiation protocol was invented and does not participate in the advertising phase? The Auto-Negotiation protocol would anticipate such an occurrence and invoke "parallel detection." With parallel detection, the more intelligent device will recognize either an NLP or an FLP and adjust its speed accordingly to either 10 Mbps (in the case of an NLP) or 100 Mbps (in the case of an FLP). The duplex mode will default to half. The reason that auto-negotiation works on twisted-pair links is that connectors, pin-outs, distances, and cabling remain the same so communication is possible once speed and duplex are determined. The same potential does not exist over 100BASE-FX fiber optics but it does with 100BASE-SX fiber optics. However, most fiber optic switch ports only support 100BASE-FX.

Connecting to Switch Ports

Most 100 Mbps switching hubs have the capability of operating their twisted-pair ports at either 10 Mbps or 100 Mbps. Usually, auto-negotiation is enabled but some switches offer fixed settings of duplex and speed. Fiber optic ports are another situation. A 100BASE-FX port offers the greatest fiber optic length so you tend to see this technology versus 100BASE-SX. Since auto-negotiation on 1300 nm fiber is not possible, the speed of the fiber port is usually fixed at 100 Mbps. A half-duplex, 100 Mbps fiber link is limited to 430 m—so to achieve the maximum fiber distance, the duplex must be set to full. That is the reality when dealing with 100 Mbps fiber optic ports on switches. They are usually fixed to 100 Mbps full-duplex with no auto-negotiation. For copper, you can usually assume auto-negotiation and 10/100 Mbps operation. This should be kept in mind with applying a media converter between end stations and switches or between two switches.

Applying Media Converters

We are going to study some media converter applications assuming we are using one of the more sophisticated devices that are capable of operating in different modes depending upon configuration.

Forced 10 Mbps Mode

Most media converter applications require operation at only one speed. When inserting a media converter between a 10BASE-T port and a 10BASE-FL port, the media converter must operate at 10 Mbps and have 850 nm transceivers. For this situation, the media converter could be forced to operate at 10 Mbps. Auto-negotiation is disabled. This situation represents the simplest of applications.

Forced 100 Mbps Mode

This application is similar to the above but this time we want to connect at 100 Mbps speed. On the copper side is a 100BASE-TX port but with its speed set to 100 Mbps. On the fiber side is a 100BASE-FX switch port with its data rate set for 100 Mbps and its duplex set to either full- or half-duplex. Therefore, 1300 nm transceivers must be installed on the media converter. Auto-negotiation is disabled. (The switched fiber port does not generate auto-negotiation signals anyway.) Modern media converters usually support full-duplex transmissions as their inherent behavior. Therefore, it is not necessary to set the media converter to full-duplex. Half- or full-duplex communication is possible over a full-duplex link so in this case the media converter is transparent for duplex issues. The Forced 100 Mbps Mode will work as long as the copper side device is fixed to 100 Mbps.

Non-Transparent Mode

Let's assume the above situation except that the copper port is set for auto-negotiation. If we left the media converter in the Forced 100 Mbps Mode, which does not auto-negotiate, the device on the copper side will not increase its speed to 100 Mbps from its default 10 Mbps speed. With one side at 10 Mbps and the other at 100 Mbps, communication is not possible. If we do not want to put restrictions on the copper device, we need to change the operation of the media converter itself. Ideally, we would like the media converter to facilitate auto-negotiation, but we cannot do that when the fiber side does not auto-negotiate—we do not know speed or duplex.

Actually, we do know speed on the fiber side. We could parallel detect the speed by observing the valid link condition which differs between 10 Mbps and 100 Mbps. Once we know the fiber speed, we could communicate this to the copper side. We could participate in the auto-negotiation process on the copper side as if we were the link partner. We still do not know the duplex and we need some manual intervention (usually through a switch setting on the media converter) to select duplex. Remember that the media converter inherently functions in full-duplex mode so duplex selection is only "advertised" to the copper side as part of the auto-negotiation process. Thus, nothing different happens within the media converter. In our example, we know we have a full-duplex port so we would set the media converter to "advertise" full-duplex. This mode of operation is called Non-Transparent since auto-negotiation is possible on the copper side but not on the fiber side.

When the media converter is operating in non-transparent mode, the media converter's copper output port is disabled but the output of the fiber side is not. This allows the connected fiber device to "hear" the equivalent auto-negotiation signals sent by the copper device. Once the media converter recognizes that signals are being sensed on both the media converterís copper and fiber optic inputs, it enables its copper port output in order to participate in the auto-negotiation process. If the media converter receives FLP bursts on the copper side, it will send out its own FLP to the copper partner advertising the speed and duplex of the fiber side. In this way it is operating as a proxy for the fiber side.

Once auto-negotiation is complete, communication begins at the required 100 Mbps data rate.

Transparent Mode

The transparent mode is used when both copper side and fiber side devices are capable of auto-negotiation. In this mode, the media converter functions exactly like the non-transparent mode except the media converter does not proxy the information from the fiber side. Instead, information from both link partners is transferred as if the media converter were not present. In other words, the media converter is transparent. In this mode, full support of 100BASE-SX is possible as long as 850 nm transceivers are used. However, transparent mode is not restricted to just 100BASE-SX.

For example, assume two back-to-back media converters operating in transparent mode but using 1300 nm transceivers. It is desired to extend the distance between two copper switch ports so we use media converters at each end with the fiber sides connected together. The two copper switch ports would auto-negotiate, and the media converters would simply pass the information. In this situation, something interesting could happen. What if the two copper switch ports negotiated 10 Mbps? Would communication over fiber using 1300 nm transceivers be possible? No matter what transceivers are used, as long as they are the same, communication is possible regardless of data rate. Although 1300 nm transceivers operating at 10 Mbps is not covered in the IEEE standards, it will still work. The resulting link is just termed proprietary.

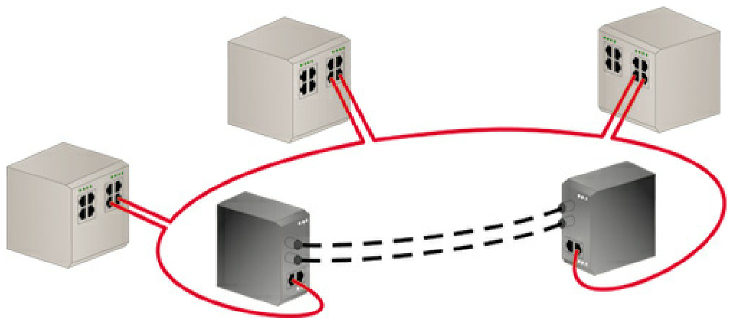

Media converters operating in transparency mode will faithfully carry critical link status in cable redundancy applications.

Why Not Just Use A Switch?

Could not a five-port switch with one fiber port operate just as well as a media converter? Switches operate at the data link layer and introduce data latency because switches must first receive frames before they are resent. That might not be a significant issue for most installations, but there is another potential issue. When using redundancy schemes such as proprietary ring or Rapid Spanning Tree Protocol, all intervening devices must be of similar technology or transparent to the network. This is because it is important to pass accurate link status between devices supporting the redundancy protocol. Assume we have two back-to-back media converters in a proprietary ring topology. Media converters in transparent mode would pass the link status of the connected copper partners within a ring. If the fiber link was broken, the link status on the copper side of the media converters would indicate loss of link. If we had substituted plug-and-play switches in place of media converters, and we broke the fiber interconnect, there will be no loss of link indicated since the fiber and copper ports on a switch are independent for link indication. The ring status would be communicated as intact when in fact it is broken. Switches are not good substitutes for transparent media converters.

Conclusion

Media converters are viewed as simple devices and are frequently used to extend distances of links that are too short or will travel into hostile environments. Because of the Auto-Negotiation protocol, applying media converters can be tricky. It is necessary to fully understand the capability of the media converters and the network in which they will be applied to ensure a successful installation.

References

IEEE 802.3-2005 Standard for Information Technology-Telecommunications and Information Exchange between systems—Local and Metropolitan Area Networks—Specific Requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications

Micro Linear DS6652-F-02, ML6652 10/100 Mbps Ethernet Fiber and Copper Media Converter with Auto-Negotiation

CONTEMPORARY CONTROLS
www.ccontrols.com