

Past issues of the Extension are available. If you would like a copy, please send your request to info@ccontrols.com

the EXTENSION

A Technical Supplement to control NETWORK

© 2002 Contemporary Control Systems, Inc.

Hubs Versus Switches—Understand the Tradeoffs

Introduction

In previous Extension articles, we discussed traditional Ethernet and its dependency on repeater sets to expand the length of these networks. As Ethernet evolved to incorporate twisted-pair cabling and star topology, a repeating hub was necessary in order to connect the various link segments together. With the introduction of switching hubs as a replacement for repeating hubs, network performance was enhanced by breaking up one collision domain into several collision domains. However, is it necessary for modern Ethernet networks to only incorporate switches? The answer is no. Repeating hubs have their place, but you must understand the tradeoffs when not selecting switch technology.

Repeating Hubs

Since modern Ethernet is wired in a star topology, a hub is required to expand a network beyond two stations. A repeating hub is intended for shared Ethernet or half-duplex Ethernet where only one station can transmit at a time, otherwise collisions happen. Collisions must be sensed by all stations on the network in order to discard the transmitted frame that has been corrupted and to notify all stations that bus arbitration is occurring. Based upon a back-off algorithm, one station will win bus arbitration thus gaining sole access to the network and, therefore, allowing it the ability to transmit without error. A repeating hub must be transparent to the network and must, therefore, reinforce collisions so that stations connected to repeating hub ports can continue to participate in network arbitration as if the repeating hub were not present. Repeating hubs must perform in other ways in accordance to IEEE 802.3 requirements for repeater sets. They must:

- Restore the amplitude of the signal
- Restore the symmetry of the signal
- Retime the signal
- Rebuild the preamble
- Enforce collisions on all segments
- Extend fragments

Repeaters and repeating hubs are physical layer devices operating on symbols sent over the wire. Repeaters do not understand Ethernet frames or protocols. Their role is to simply extend distance and to facilitate star topology.

Switching Hubs

A much more sophisticated hub is the switching hub. A switching hub or “switch” is not a repeater but actually a bridge which acts upon the content of the Ethernet frame it receives and forwards the frame to the appropriate port on the switch. For example, if station A is connected to switch port 1 and station B and C are each connected respectively to ports 2 and 3, a message from A to B should only require the passing of the data from ports 1 to 2 but not to port 3. The switch could do this if it knew all the station locations. The switch learns these locations by observing the source addresses of Ethernet frames sent to its ports. It builds up a table of station addresses for each port and studies this table before forwarding the frame. If the switch does not know the location of a particular station address, it floods the frame to all ports except the port on which the frame was received. As soon as the unknown station announces its location by initiating a frame, the switch will update its table accordingly. Like a repeater, a switching hub is also effective in expanding a network beyond two stations. However, unlike a repeater, a port on a switch becomes an end point to the link segment attached to its port. Each switch port functions like any other Ethernet station except it passes along the data and station addresses generated by originating stations. Therefore, each switch port must abide by the same arbitration rules for collision detection and access resolution, but it need not pass collision information to other switch ports because these ports reside in separate collision domains. Switching hubs are fundamentally different from repeating hubs and, therefore, their performance can differ as well.

Increasing Distance

Twisted-pair segment lengths cannot exceed 100 m and, therefore, the need to extend network distances is frequently encountered. If there is a need to extend an Ethernet network, a switch can provide an advantage over a repeating hub. Repeating hubs are considered part of the collision domain and for reliable operation of a shared Ethernet network, the network diameter must not exceed the collision domain (see Figure 1). This distance limit is based upon the round-trip propagation delay between the two furthest stations on

the network, and it cannot be such that collisions are not sensed by all stations within a prescribed time. Although repeating hubs can be cascaded for greater network distance, there are restrictions that usually limit the number of repeating hubs to four. The rules for cascading repeating hubs were explained in previous Extension articles. The simplified rule is the 5-4-3 rule that states there cannot be more than five segments, four repeaters and three mixing segments. A mixing segment is a bus segment such as thick-wire and thin-wire coaxial cable. Since modern Ethernet only uses link segments consisting of either twisted-pair or fiber-optic cabling, the mixing rule can be dropped. Therefore, having five 100 m twisted-pair segments and four repeaters is fine. Since fiber optic segments can be up to 2 km in length, we need to impose limits to this rule. Instead of having four repeaters, we can have three with two interconnecting fiber optic segments, each only 1 km in length. The expansion rules for repeating hubs are indeed confusing, but they are necessary in order that the network diameter not exceed the collision domain. What is even more confusing is that these rules only apply to 10 Mbps operation.




Figure 1. With shared Ethernet, all devices and associated cabling must reside in a single collision domain.

One way of simplifying the expansion rules is by avoiding the collision domain restrictions. Switching hubs are not part of the collision domain since they are end devices on a network (see Figure 2). Therefore, adding one switch to a network without switches can effectively allow the network diameter to double without exceeding the collision domain. Additional switches can be cascaded without the limit imposed on repeating hubs. In other words, if there is no collision domain issue to begin with, (the network diameter does not exceed the collision domain) then no collision domain issue will be introduced by adding a switch. In fact, switches can continue to be added beyond the four limit of repeating hubs regardless of data rate. This makes expansion rules much easier. However, the collision domain rules for each link segment on a particular port must be followed since it is possible that repeating hubs are attached to switch ports. The collision domain restrictions remain, but they are not aggravated by adding switches. It is also

possible to regain the 2 km segment length in fiber optic ports if the fiber segment is between two switch ports. This is true at 10 Mbps and possible at 100 Mbps. To achieve large network diameters, especially at higher speeds, switches may be the only option.




Figure 2. Because switches break the network into multiple collision domains, the physical size of the network is virtually unlimited.

Data Latency

Data latency is the time delay experienced when data is sent from one point to another. A contribution to data latency is the time it takes for an electrical signal to propagate down a wire. Although electricity travels fast, its speed is still finite and over a wire it is slower than in a vacuum. For twisted-pair wiring it is about 5.5 ns/meter. Since the same wiring is used to attach a repeating hub or switching hub, this delay is the same for both implementations. However, the biggest contributor to data latency is the hub itself and the amount of delay depends on if we are using a switch or a repeating hub. One requirement of a repeating hub is to rebuild the preamble in case a frame is received with less than the required number of preamble bits. At 10 Mbps, it takes 100 ns to send one bit so the repeating hub must be able to queue at least a few bits worth in case it needs to stuff more preamble bits. Now the repeating hub does not wait for the full 64 preamble bits to be received before it starts forwarding the frame, however, there is a delay. From our own measurements at Contemporary Controls, the delay through a repeating hub can vary from 625 ns to 840 ns. This delay is negligible to the overall performance of the network.

As mentioned before, a repeating hub operates at the physical layer and handles the symbols on the wire. A symbol is the waveshape on the wire that distinguishes a logical 1 from a logical 0. A frame is a collection of symbols representing one transmission unit sent over the wire. An Ethernet frame contains fields beginning with a preamble and ending with a check frame sequence as shown in Figure 3. A valid Ethernet frame must contain all fields including preamble, start of frame, destination address, source address, length, data and frame check sequence. Each field has a fixed length except for the data field which can vary in length. Since the Ethernet protocol requires a minimum length frame size, the data field cannot be

less than 46 bytes. The largest data field can have up to 1500 bytes. Since the other fields are fixed for a total of 18 bytes, the minimum Ethernet frame is 64 bytes and the maximum Ethernet frame is 1518 bytes. The preamble is excluded in this calculation. This information is needed to calculate the delay through a switch.

| IEEE 802.3 Frame | | | | | | | |
|------------------|----------------|------------------------------|---|---------------------|--|-----------------|----------------------|
| 56 bits | 8 bits 48 bits | | 48 bits | 16 bits | 368 to 1200 bits (46 to 1500 bytes) | 32 bits | |
| Preamble | SFD | Individual/Group Address Bit | Globally/Locally Administered Address Bit | Destination Address | Source Address | Length LLC/Data | Frame Check Sequence |
| | | | | | | | |

Figure 3. A store-and-forward switch must read in the complete Ethernet frame before forwarding.

Switch Operation

For sake of discussion, let us assume we have a two-port switch. Some individuals would call this a bridge. On each side of the switch is a transceiver. Each Ethernet transceiver provides a physical end to an Ethernet wiring segment and the actual end of the network diameter. Therefore, one two-port Ethernet switch links two separate Ethernet networks. Since there is no distinction between one device on one network versus another device on the other network, the two networks function as one larger network. In this case a switching hub and repeating hub function in a similar manner. However, there is one big difference. A switch port stores the complete frame before it passes it to the other port. It does this because it needs to know the destination of the received frame and to verify that a valid frame was received by observing the frame check sequence. If the frame is invalid, it should be discarded instead of forwarding a faulty frame. Since the switch must store one complete frame before forwarding, the delay in observing the frame on the other port is at a minimum one frame. Since frame sizes can vary, the delay can vary. At 10 Mbps and the shortest frame, the delay is 51.2 μ s but with the longest frame the delay is 1.21 ms. Is this a problem? Not necessarily.

Assume we are going to send a message consisting of 1000 Ethernet frames, and we were fortunate to be able to send them back-to-back with the minimum interframe gap. What would be the data latency due to the switch? The answer is still only one frame's worth. Therefore, by either sending out one frame or 1000 frames, the switch only queue's one frame's worth of data under normal conditions. Therefore, it would appear that switch latency is not an issue. Let us study the situation closer.

Let us assume we have one controller functioning as a master and one input/output (I/O) device functioning as a slave. The slave only responds to requests by the master and never initiates a request

itself. Further assume that the two devices are separated by a switch. If the master initiates a one-frame message there will, of course, be a one-frame delay before the slave receives the message. The slave acts upon the request and initiates a one-frame response which also incurs a one-frame delay. Therefore, with a single command/response session, there is a two-frame delay introduced into the process simply by adding a switch. If both frames were long, a total of 2.4 ms can be added to the process. Now if we substitute a hub for the switch, we would not suffer the 2.4 ms delay, demonstrating that under certain conditions a repeating hub can outperform a switching hub.

Now someone might take issue with my example by saying that by using a repeating hub, I could potentially introduce collisions on the network that would reduce performance. That is true, but let us study the protocol. Once the master senses a quiet line, it initiates a transmission and then waits for a response from the slave. The only slave to respond was the one polled. The master consumes the response and then initiates another command to another slave. This activity continues with no collisions since collisions are avoided by the rhythmic commands and responses of the master/slave protocol. Your most popular industrial protocols such as Modbus and Optomux operate this way. Therefore, my example is valid.

Now let us change the protocol slightly. Instead of the master making a single command to a slave, it makes multiple commands each to individual slaves without waiting for slave responses. Eventually each slave will act on its command and generate a response. Depending upon the complexity of the command and sophistication of the slave device, responses will begin to occur approximately at the same time and the possibility of collisions increases that will reduce throughput. Will a switch operate better in this situation? A switch will eliminate collisions, but it does not mean it will not drop packets. One of the switch ports connects to the master, and each slave has its own switch port. All the traffic will be directed to the master port due to the numerous responses, which could flood the output buffer of the port connected to the master. If there is a buffer overrun, then packets will simply be lost.

Reducing Data Latency

One way of reducing data latency in a switch is to operate at 100 Mbps instead of 10 Mbps. This could reduce latency by a factor of ten. Another approach is to use cut-through operation. With cut-through operation, a switch does not wait for receipt of the

complete frame before forwarding. A switch only needs to know the destination address before forwarding, and that address is available near the beginning of the frame. Therefore, there is no need to wait for the complete frame to be received. The problem with this approach is that the frame could have a failed frame check sequence (FCS) or it could be a runt frame. These frames should not be forwarded. The runt frame problem can be resolved by waiting a bit longer before forwarding, but the failed FCS problem is not solvable with cut-through operation.

100 Mbps Operation

If stations have adequate processing power, operating at 100 Mbps should improve network performance. However, the problem at 100 Mbps is that the collision domain diameter decreases by about a factor of ten. Using repeating hubs at this speed is not very practical and certainly not very popular because the reduction in network diameter is about 205 m. Therefore, the use of switches at 100 Mbps instead of repeating hubs is a clear choice.

Full-Duplex Operation

The original Ethernet technology, operated at half-duplex, is called shared Ethernet. Only one station can transmit at a time. With full-duplex and switch technology, a station can both receive and transmit at the same time over a link segment such as twisted-pair or fiber optics. Since the simultaneous transmitting would normally cause a collision, the collision detection circuitry is disabled, eliminating the collision domain. With no collision domain, the only limit to segment lengths is attenuation of the signal over the segment. Therefore, in full-duplex mode, a 2 km fiber optic segment length, or longer, is possible between two switches operating at 100 Mbps whereas it is limited to 412 m in half-duplex mode. Full-duplex further simplifies the expansion rules and could improve performance on some networks. However, with traditional master/slave industrial protocols, you will probably not witness any performance improvement.

Broadcasts and Multicasts

A broadcast is a transmission to all stations and a multicast is a transmission to a selected group of stations. Since the switch does not know all the potential broadcast and multicast stations, it must flood all ports on the switch with these frames. In this situation, a switch functions just like a hub. Modern

producer/consumer protocols utilize broadcasts and multicasts and, therefore, network performance may not necessarily improve with the use of switches.

Using Network Analyzers

A network analyzer is a helpful tool when troubleshooting networks. These tools observe all the traffic on the network and capture portions of traffic for analysis purposes. The network analyzer connects to an unused port on one of the switches or repeating hubs. With a repeating hub there is no problem since all traffic can be viewed at any repeating hub port. The problem exists with switches. Since switches filter traffic to its various ports, the network analyzer will not observe the same traffic on all ports. To remedy this problem, the switch could provide a feature called port mirroring where all the switch traffic would be directed to a particular port. This feature is usually only found in higher-cost managed switches.

When to Use Hubs and Switches

| <u>Issue</u> | <u>Recommendation</u> |
|-----------------------|-----------------------|
| Low cost | Hubs |
| Reduce data latency | Hubs |
| Simple application | Hubs |
| 100 Mbps operation | Switches |
| Full-duplex operation | Switches |
| Auto-negotiation | Switches |
| Large distances | Switches |
| Master/slave protocol | Hubs or Switches |
| Peer-to-Peer protocol | Switches |
| Network analyzer | Hubs |

Summary

As you can see, there are both advantages and disadvantages in using switches. For simple systems, a repeating hub can be quite effective so their use should not be immediately ruled out.

REFERENCES

The Switch Book, Rich Seifert, 2000, Wiley Computer Publishing

Ethernet The Definite Guide, Charles E. Spurgeon, 2000, O'Reilly & Association, Inc.

International Standard ISO/IEC 8802-3 ANSI/IEEE Std 802.3, 1996, The Institute of Electrical and Electronic Engineers, Inc.