AIRBORNE NETWORK SWITCH WITH IEEE-1588 SUPPORT

John Hildin
Director of Network Systems
Sergio Arias
Principal Engineer
Teletronics Technology Corporation
Newtown, PA USA

ABSTRACT

Today's data acquisition systems are typically comprised of data collectors connected to multiplexers via serial, point-to-point links. Data flows upstream from the sensors or avionics buses to the data acquisition units, to the multiplexer and finally to the recorder or telemetry transmitter. In a networked data acquisition system, data is transported through the network "cloud". At the core of the network "cloud" is the network switch. The switch is responsible for distributing and directing data within the network. Network switches are commonplace in the commercial realm. Many businesses today could not function without them. A network-based data acquisition system, however, places additional burdens on the network switch. As in a commercial network, the switch in a data acquisition system must be able to distribute data packets within the network. In addition, it must be able to perform in a harsh environment, occupy a minimal amount of space, operate with limited or no external cooling, be configurable, and deal with the distribution of time information.

This paper describes the required features of a ruggedized network switch and the implementation challenges facing its design. As a core component of a network-based data acquisition system, an ideal switch must be capable of operating in a large number of configurations, transporting and aggregating data between data sources and data sinks, with a mixture of devices operating at rates ranging from a few thousand bits per second to several gigabits per second, over twisted pair or fiber optic links. To ensure time coherency, the switch must also facilitate a time distribution mechanism, e.g., IEEE-1588 Precision Time Protocol (PTP). The gigabit switch described here uses the PTP to implement an end-to-end clock synchronization, for distributed acquisition nodes, to within 300 nanoseconds.

KEYWORDS

Network, Gigabit Switch, IEEE-1588, PTP, Data Acquisition

INTRODUCTION

Whether used within a corporate network or an advanced airborne data acquisition system, the network switch is responsible for filtering and forwarding data packets throughout the network. Although an airborne switch shares many of the same features and functions as a commercial switch, there are several significant differences. One obvious difference is the increased ruggedization required for the switch to operate in an airborne environment. This includes more stringent requirements related to shock, vibration, temperature, altitude, resistance to contaminants, etc. than those imposed upon a commercial product. Another important difference is the requirement to support time coherency within the data acquisition network. Historically, commercial data networks have not had this requirement.

Although key switch features and implementation challenges are discussed, a treatment of fundamental network operation and protocols is beyond the scope of this paper. The reader is encouraged to refer to any of the numerous texts available regarding network operation.

SWITCH BASICS

The primary task of a network switch is to filter and forward packets. In the case of a switch designed for a data acquisition network, time distribution is also a key function. Switches connect various network segments to form larger heterogeneous networks. Acting as the traffic cop for network data, the switch carries and directs frames between segments connected to its ports. The simplest example of this is broadcasting arriving frames to all other ports on the switch. This is inefficient, however, as all the segments would carry this traffic, regardless of where the destination node were located. The most basic switch has the ability to manage the traffic by forwarding the frames directly to the destination node or to another switch that can forward them to the node.

To aid the switch in making forwarding decisions, the switch maintains addressing tables. Each frame contains a source and destination Media Access Control (MAC) address. The MAC address is the hardware address of a node on the network. The switch reads the source MAC address of frames entering the switch and stores them in the table. From this the switch can associate the physical port to MAC addresses. The switch also reads the destination MAC address to determine where the frame needs to go. The MAC address table is parsed to find the relationship between the MAC address and physical switch port. Once found, the frame is sent to the appropriate port. If no address match is found, the switch sends the frame out on all its ports, except the port it arrived on.

Multicast addressing allows a node to send frames to multiple destination nodes. A switch can maintain a list of node addresses that are members of a “multicast group”. Once registered to the group, nodes can receive multicast traffic bound for that group. Some switches allow the automatic detection of multicast groups by “snooping” the packets exchanged by nodes during the registration process. This helps the switch determine what physical ports to send multicast packets.

Switches can be placed into two categories, managed and unmanaged. A basic Ethernet switch can operate autonomously with no external configuration or interaction required. This type of switch is an example of an unmanaged switch. A managed switch can be configured or reconfigured during operation. In addition, a managed switch can offer additional benefits. Features including flow control, fault detection, QoS (Quality of Service), port trunking, performance statistics, and others are often available on managed switches.

CORE OF THE NETWORK-BASED DATA ACQUISITION SYSTEM

The switch is the main building block of an Ethernet network. All nodes on the network connect to a hub or switch port via a point-to-point link. Without the switch (or hub), the network would not have more than two nodes. The switch enables scalability of the network. As data sources and sinks are added to the network, the number of switch ports needed also increases. Switches are simply added to the network to support the growth of network nodes. Figure 1 below shows a block diagram of a network-based data acquisition system.

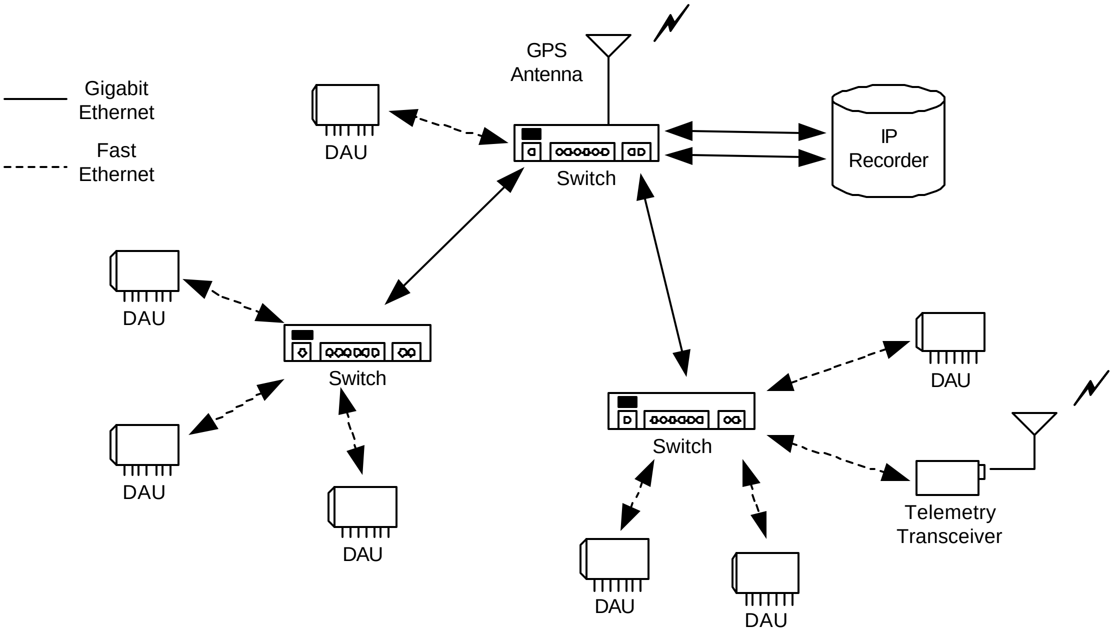

Figure 1. Network-Based Data Acquisition System

Time coherency is an important characteristic of a data acquisition system. In traditional systems, time is distributed on a separate network (out-of-band) with its own physical wiring. In a network-based system, time can be transported on the same network that carries the data traffic (in-band), and using the same physical wiring. To facilitate time distribution, the switch must be designed to terminate the time protocol and propagate it on its ports. This is one of the key elements of a network data acquisition switch that separates it from its commercial counterpart.

Often, system devices can generate from kilobits per second of data to gigabits per second of data, depending on the device (e.g., sensor, multiplexer, etc.). In order to maintain flexibility and protect the customer's investment, the switch should be able to support the connection of devices with varied data rate requirements. For Ethernet, standards exist for 10 Mbps, 100 Mbps, 1 Gbps and 10 Gbps data rates. The physical layers defined include both twisted pair copper wiring and fiber optic cabling.

In some cases, a device may require bandwidth greater than what a single switch connection can provide. For example, one switch may need to send 150 Mbps of traffic to another switch in the network. This would exceed the bandwidth of a single 100Base-T port on either switch. In this case, two 100 Mbps ports could be aggregated together, forming a virtual single 200 Mbps link between the switches. This illustrates the feature of port "trunking". Port trunking gives the user the ability to aggregate a number of ports on the switch to form a larger capacity link. In a similar fashion, two or more gigabit ports could be trunked together to form a single high-capacity link. In these cases, the user needs to decide whether to combine lower speed ports together or use a single, higher speed port. In some cases, using a higher speed port increases cost and underutilizes the available bandwidth.

Below is a list of important features to consider when including an Ethernet switch in an airborne data acquisition network.

- Port count/density
- Port speeds (10/100/1000 Mbps)
- Media support (twisted-pair copper or fiber)
- Support for IEEE-1588 PTP
- Support for 1588 Timing Grand Master
- Support for port trunking
- Managed or unmanaged switch
- Aggregate bandwidth ((number of ports) x (maximum throughput per port))
- Switching capacity (number of packets per second that can be forwarded)
- Number of multicast groups supported

NETWORK PACKET-BASED TIME DISTRIBUTION

Any accurate timing distribution system must consider the propagation delay of the information to the endpoint. With the delay known, it can be adjusted for at the node and virtually eliminated. The main issue with time distribution on network-based systems is the variation in the propagation delay of the packets. The delay can vary from hundreds to thousands of microseconds.

There are several factors affecting the delay of the information crossing the network at the application layer. The most important ones are software latency and network latency. Software latency is the time it takes the driver software to respond to an event. The variation of this delay is affected by the response time of the software running at each node and the loading at each node.

Network latency is associated with the time it takes a frame to move from the MAC of the sending node to the MAC of the receiving node. The network switches have the most effect on the variation of this delay. The time it takes an Ethernet frame to move from the incoming port of a switch to its destination port is not bound and can typically be from hundreds of nanoseconds to hundreds of microseconds. How well this time can be bound will determine the delay variation on the data being transmitted between two nodes. The variation of the delay through a switch results in jitter. The jitter must be removed or minimized to maintain an acceptable level of timing accuracy between the distributed nodes in the network. This suggests that the switch must support a mechanism to characterize its contribution to the overall timing uncertainty.

There are various protocols that exist to allow network nodes to synchronize their time to a master clock source. One popular protocol is the Network Time Protocol or NTP. NTP is an application layer protocol whose performance is related to the particular software implementation. NTP can provide clock synchronization to within hundreds of microseconds on a controlled, local network and to within tens of milliseconds on larger networks (e.g., the Internet). Generally, NTP does not provide the degree of synchronization required in an airborne data acquisition system.

The IEEE-1588 Precision Time Protocol (PTP) provides a standard way to synchronize devices over a packet-based network with sub-microsecond precision. As opposed to NTP, PTP operates at the Physical, Data Link and Network layers and is isolated from the larger time uncertainties of the upper software layers. It should be noted that PTP could be implemented on a switch or other device as software only. When this is done, the precision of the timing is reduced due to the variations in the software execution, driver software performance, etc. These variations can result in time differences greater than several milliseconds. To ensure maximum accuracy, the PTP must be terminated as close to the physical interface as possible. Doing so requires specialized hardware. Implementing specialized hardware allows time packets to be read between the physical layer device (PHY) and the Media Access Controller (MAC). Figure 2 below shows a high-level block diagram of the IEEE-1588 processing logic in relation to the Ethernet MAC and PHY devices. Using the hardware approach, accuracy can be improved to the 200 nanoseconds range. So the designer must make the trade-off between accuracy and product cost/complexity when deciding whether to implement the IEEE-1588 protocol in software versus hardware.

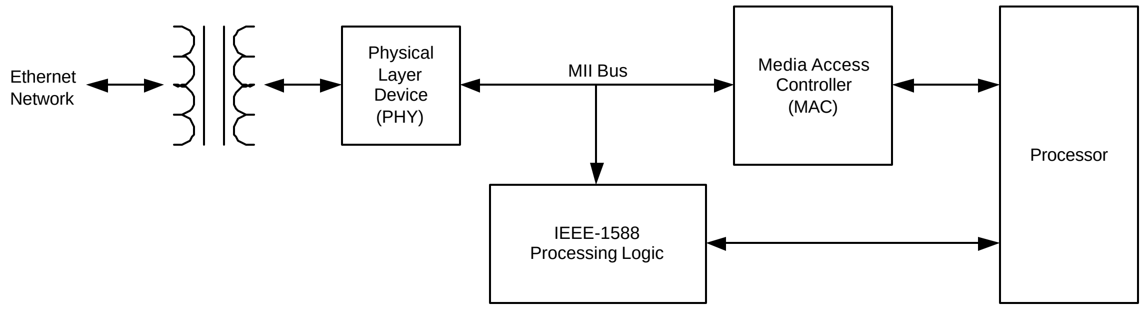

Figure 2. IEEE-1588 Circuit Block Diagram

APPLICATIONS

The switch is primarily designed for an airborne data acquisition network. The airborne application may require installation within the test article or outside, but within the aerodynamic structure of the test article. The switch must meet harsh environmental conditions for use in the airborne environment. These conditions are found in other applications suitable for a ruggedized switch that requires time distribution capability. The applications may include:

- Network with time distribution in a test range
- Naval Shipboard network switch with time distribution
- Network based telemetry ground test vehicle
- Wind tunnel network based instrumentation

CONCLUSION

This paper has attempted to describe the switch as a key element in a network-based data acquisition system. Aside from the role of traffic cop, the switch must faithfully distribute time through the network for use by the nodes. In some cases, the switch may also act as the timing grand master in the network. Offering flexibility relative to the port types and speeds allows a customer to change the topology of the data acquisition system without replacing the switch. This flexibility has a price in terms of the complexity of circuits needed to terminate the IEEE-1588 protocol for any given port speed, especially higher rate ports like gigabit Ethernet.

The IEEE-1588 protocol allows for accurate time coherency within a packet-based network. Using the protocol enables time information to flow in-band with the data traffic as opposed to out-of-band like traditional TDM-based (Time Division Multiplexed-based) systems. The switch design must provide for a protocol that supports time distribution (IEEE-1588 or other) to ensure accurate time at the nodes.

Lastly, building management capability into the switch gives the user the ability to characterize performance, diagnose data traffic issues and monitor the operation of the switch itself. This feature becomes more important as the size and complexity of the data acquisition system increases.

REFERENCES

Institute of Electrical and Electronics Engineers, Inc., IEEE Std. 1588 - 2002 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurements and Control Systems, Nov. 8, 2002.

Berdugo Albert, and Roach John, “Migrating from PCM to Network Data Acquisition System”, Proceedings of the 26th European Telemetry Conference (ETC 2006) Germany, May 02-05, 2006.

Roach, John A., Hildin, John “IP-Based Recording System”, Proceedings of the International Telemetering Conference (ITC 2006), San Diego, CA USA, October 23-26, 2006.