Buses and Networks for Contemporary Avionics

By Mike Glass
Principal Marketing Engineer
Data Device Corporation

©Data Device Corporation

WHITE PAPER

November 2007

Buses and Networks for Contemporary Avionics

Abstract

MIL-STD-1553 has served the needs of military system integrators for over 30 years, particularly in the area of command and control applications. Nevertheless, contemporary applications such as high-speed digitized sensors, file transfers, processor clusters, and displays require much higher data rates than 1553's 1Mb/s. For some environments, particularly for legacy aircraft, the optional solution is to transmit faster data rates over existing 1553 buses. However, there are other applications that can accommodate and benefit by the deployment of gigabit or multi-gigabit copper or optical switched fabric networks. In addition to MIL-STD-1553, this paper presents and comments about several avionics networking technologies including High-Speed 1553, Fibre Channel, Gigabit Ethernet, and ARINC 664, a form of profiled Ethernet.

Introduction

For decades, MIL-STD-1553 has served as the workhorse networking standard for the integration of military/aerospace avionics platforms. However, modern avionics applications provide increasing demands for network bandwidth beyond 1 Mb/s. In addition to 1553's traditional command and control functions, these applications include processor and DSP clusters, digitized sensor interfacing, displays, file transfers, and data storage. This paper discusses the directions in which military avionics networks are now evolving.

One nascent trend is the use of modern telecommunication modulation techniques for providing order(s) of magnitude increase in data rate while leveraging the existing 1553 cable infrastructure. This eliminates the significant costs involved with rewiring. Specifically, this includes High Performance 1553, and the US Air Force's 1553B Notices 5 and 6. With these methods, it will be possible for some subsystems to operate at hundred(s) of Mb/s data rates with co-existing with legacy 1 Mb/s systems on the same physical data buses.

For avionics, Fibre Channel provides 1 and 2 Gb data rates, and is deployed on a number of military/aerospace platforms and programs, including F-18E/F, F-16, F-35, B-1B, B-2, E-2D, the Apache Longbow and MMH helicopters, and AESA Radar. Fibre Channel includes a variety of upper layer protocols, which are targeted to general and avionics networking, storage, sensors, displays, and weapons interfacing.

Ethernet, leveraging the ubiquitous TCP/IP protocol, is the dominant standard for commercial local area networking, with a vast ecosystem of operating system and application software. Gigabit Ethernet enables seamless internetworking between airborne computers, sensors and other embedded systems on air and ground vehicles, soldiers, wireless networks, satellites, along with local and remote command centers.

ARINC 664, or AFDX, is a profiled version of switched Ethernet and Gigabit Ethernet used for commercial avionics. AFDX's optimizations for avionics networking include provisions for guaranteed bandwidth and maximum latency, and in-order delivery of multicast messages. Additional features of AFDX include robust redundancy and network diagnosis capabilities.

MIL-STD-1553

History

Throughout the 1950s and 1960s, the integration of various avionics "black boxes" for military aircraft was achieved by means of thousands of feet of point-to-point discrete wiring. To the military services, this illuminated on a clear need to develop a networking standard. In addition to eliminating the burdens of the additional weight and installation effort associated with discrete wiring, there were other benefits to be realized by means of networking. These included ease of insertion, removal and maintenance of the various subsystems, diagnosibility, as well as the capability to implement upgrades by means of software changes.

In 1968, the U.S. Air Force, in conjunction with the Society of Automotive Engineers (SAE), formed the A2K committee for the purpose of developing a networking standard. In 1973, MIL-STD-1553, the Military Standard Aircraft Time Division Multiplexing Data Bus, was issued. In 1978, the standard was re-issued as MIL-STD-1553B, its current revision level. Over the past 30 years, 1553 has been widely implemented, principally for Air Force weapons platforms, but also for Army, Navy, and NASA applications. These include a list of fighter, attack, and bomber aircraft; helicopters, missiles, ground vehicles, ships, satellites, and the International Space Station.

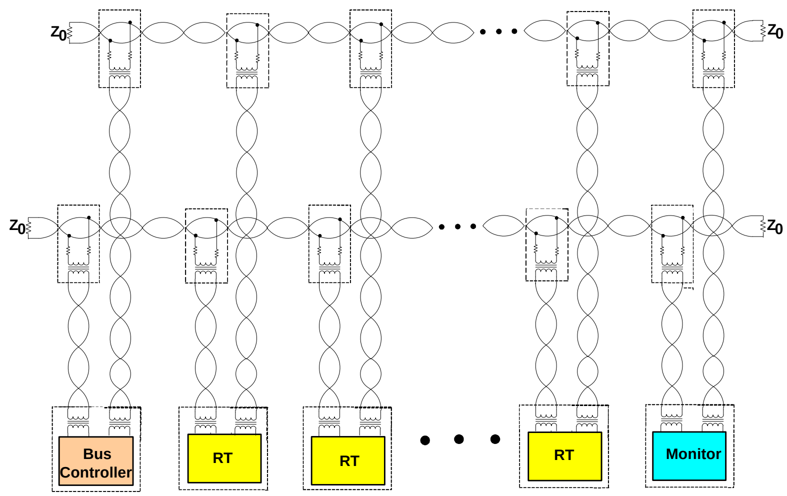

Figure 1. MIL-STD-1553 Dual Redundant Bus

Topology

Figure 1 illustrates the basic multi-drop topology of a dual redundant MIL-STD-1553 bus. For systems requiring high reliability, a dual redundant bus is generally used, but a single bus system is also possible. The media used for 1553 buses is twisted/shielded cable with a characteristic impedance in the range of 70 to 85 ohms.

As shown in Figure 1, 1553 specifies three types of terminals: bus controller (BC), remote terminal (RT) and monitor terminal (MT). At any given time, there is one BC and up to 31 RTs, and there may be one or more bus monitors on a data bus. MIL-STD-1553's media access scheme involves central control, by means of command/response time division multiplexing. With command/response operation, the bus controller operates as the bus master, initiating all message sequences, while the remote terminals respond as slave devices. A bus monitor is a passive (non-transmitting) device that receives and stores either all data or a subset of data received from the bus.

Data Encoding and Waveshaping

MIL-STD-1553's methods for data encoding and signal waveshaping are shown in Figure 2. MIL-STD-1553 specifies Manchester Biphase-L encoding. For a 1 Mb/s data rate, this entails the use of 1 µS symbols for data bits, and 3 µS word sync symbols used to indicate the start of each word. Relative to other encoding methods, Manchester offers a number of advantages including simplicity (only 4 symbols types): it is self-clocking; it has zero DC value, enabling the use of transformers and minimizing offset voltages; it requires relatively low encoding and decoding clock rates, reducing power consumption; and noise immunity as it is decoded on a half-bit basis, thereby providing undetected error rates on the order of 10^{-12} or lower. In addition, 1553's 3 µS word sync patterns minimize false start-of-message determinations and provide clear boundaries between individual word times.

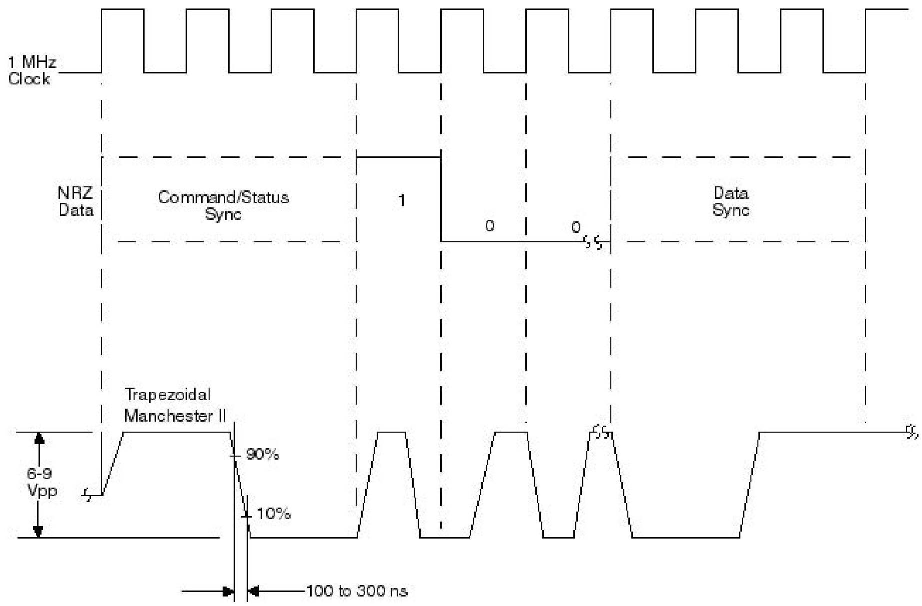

Figure 2. MIL-STD-1553 Data Encoding and Waveshaping

For transformer coupled terminals, as shown in Figure 1, the specified voltage across a transmitting terminal's stub in 18 to 27 volts, for driving a 70-ohm test load. As shown in Figure 2, this results in a voltage of approximately 6 to 9 volts across a 1553 bus. As shown, MIL-STD-1553 specifies trapezoidal waveshaping with a range of rise and fall times of 100 to 300 ns. Trapezoidal, rather than sinusoidal waveshaping results in simpler transmitter designs, including improved control over the critical parameter of zero crossover timing. The purpose of the lower limit on rise/fall times is to minimize the signal's harmonic content above 1 MHz. This serves to minimize crosstalk as well as transmission line reflections that can result in false zero crossings and hence bit errors. As a means of minimizing drive stage power dissipation, most transmitter designs tend toward the lower limit of the rise/fall time spec.

MIL-STD-1553 does not specify a maximum cable length. Instead, it specifies a maximum cable attenuation of 1.5 dB/100 ft. at the specified signaling rate of 1 MHz. Further, a 1553 physical bus must be designed so as to provide a signal level of 1.0 volt peak-to-peak to the stub of every terminal on the bus.

The receivers for all 1553 terminals must respond to stub voltages within the range from .86 to 14 volts and must *not* respond to signals less than 0.2 volts. Other receiver parameters include input impedance, tolerance to zero-crossing distortion (±250 nS minimum) and common mode rejection. 1553 also includes criteria for receiver noise rejection. For this, terminals mush exhibit a maximum word error rate of 10^{-7}, while receiving 140 mV rms of 1 kHz to 4 MHz white Gaussian noise superimposed on a 2.1 volt peak-to-peak signal.

Word and Message Formats

MIL-STD-1553 specifies three types of 16-bit words that are transmitted over the 1553 bus: Command Words, Status Words and Data Words. All three word types are 20 µS long, consisting of a 3 µS sync field, 16 data bits, and a single parity bit. Bus Controllers transmit Command Words and Data Words, while Remote Terminals transmit Status Words and Data Words.

BC Command Words are transmitted at the start of each message by the bus controller. Command words include the 5-bit address of the target RT; along with a "T/R" data direction bit; a 5-bit Subaddress field, which generally identifies a system-specific function, along with a 5-bit Word Count/Mode Code field, indicating the number of Data Words to be transmitted or received. RT Status words provide acknowledgement by including the RT's 5 address bits, along with 8 additional bits indicating about the status and validity of messages, along with health and status information for the RT and its attached system. Data words contain 16 data bits; other than stipulating that the MSB be transmitted first, MIL-STD-1553 includes no specifications about Data Words.

Figure 3 illustrates the MIL-STD-1553's basic message formats: BC-to-RT transfer, RT-to –BC transfer, RT-to-RT transfer, and mode code messages.

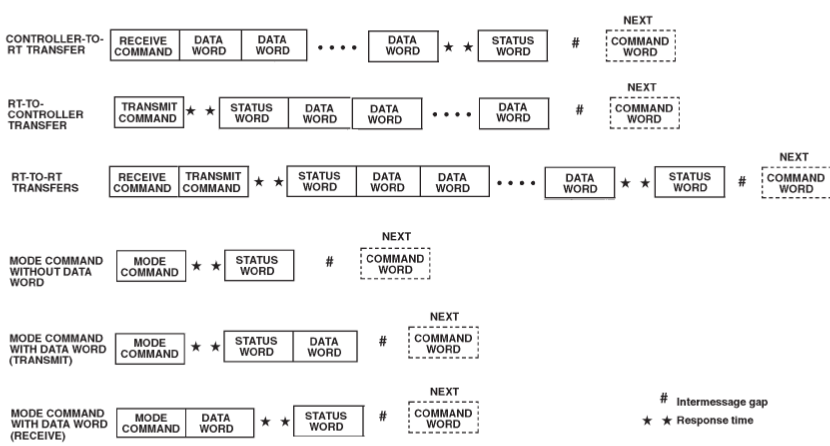

Figure 3. MIL-STD-1553 Basic Message Formats

For all these message formats, the BC initiates the message transaction by transmitting one or two Command Words, or a Command Word followed by one or more Data Words. MIL-STD-1553B requires an RT to respond with either a Status Word, or Status Word followed by one or more Data Words. After receipt of a valid transmission, RTs must respond within 12 µS. BCs must wait a minimum of 14 µS before determining that an RT is non-responsive. RTs receiving transmissions that are determined to be invalid must *not* respond.

BC-to-RT transfers are used for sending data from a BC to an RT, while RT-to-BC transfers enable RTs to transmit to the BC. In an RT-to-RT transfer, the BC transmits two contiguous Command Words, one to the receiving RT and the second to a transmitting RT. This directs one RT to transmit to the other. In all three cases, the number of data words must be between 1 and 32.

In general, mode code transfers are used to provide bus management functions, rather than to convey system data. These include checking RT status, resetting an RT, initiating RT self-test or polling self-test results, shutting down or re-enabling faulty (usually "babbling") transmitters, synchronizing time or other parameters, inhibiting or re-enabling the "Terminal flag" (RT fault indicator) bit, re-transmitting the last Command Word, or transmitting the system-defined "Vector word" for the system attached to an RT.

MIL-STD-1553 also includes broadcast message formats. These enable either the BC or an RT to send data to all other terminals on a bus. In addition, the BC can also broadcast certain mode codes.

MIL-STD-1553 continues to see widespread use. It is widely deployed and well understood, with a large installed base along with a broad industry ecosystem of components, boards, software, transformers, and cables. Its salient features continue to be its rugged physical layer, highly deterministic command/response protocol, and robust error checking. As a result, it continues to be deployed in systems such as aircraft, weapons and space platforms requiring the most reliable levels of operation.

High-Speed 1553

For new military manned aircraft and UAVs, high-speed networks such as Gigabit Ethernet and Fibre Channel are often deployed for achieving higher rates and wider connectivity than is provided by MIL-STD-1553 buses. However, for the U.S. DoD's large fleet of legacy aircraft, migration to newer networking standards can necessitate extensive re-wiring of network cabling. This can entail prohibitively high costs, not only for the new cabling but also for taking aircraft out of service.

While some systems will benefit greatly by migrating to higher data rates, there will also by many other boxes on existing 1553 buses that do not require higher bandwidth. These systems will need to continue using 1 Mb/s MIL-STD-1553 signaling, while needing to co-exist with the higher speed traffic on the same buses. As a result, there is a need to be able to transmit higher data rates over the existing MIL-STD-1553 bus infrastructures, while also allowing concurrent operation with 1 Mb/s 1553 traffic for legacy systems.

The Challenge

The theoretical maximum capacity for a communication channel is given by the Shannon equation:

$$C = W \times Log_2\left(1 + \frac{S}{N}\right)$$

where:

C = the channel's maximum capacity, in bits per second
S = the signal level at the receiver, in RMS volts
N = the level of white Gaussian noise at the receiver, in RMS volts
W = the bandwidth of the signal, in hertz

For the case where the signal and noise levels vary as a function of frequency over a frequency range f_1 to f_2, the modified form of Shannon's capacity equation is given by this equation:

$$C = \int_{f1}^{f2} Log_2\left(1 + \frac{S(f)}{N(f)}\right) df$$

Figure 4 illustrates the modeled and measured channel response of a MIL-STD-1553 bus, from DC to 80 MHz. The indicated high frequency attenuation is attributable to bus cable, coupling transformers, and stub loading. As a result, for the purpose of maximizing channel data rate, the "prime real estate" area of frequency spectrum is in the lower portion of the range.

Figure 5 illustrates the frequency spectrum for the signaling of a MIL-STD-1553 message, shown for the case of a 32-word message with random data. As shown, for 1553's 1 Mb/s Manchester Biphase-L (2 Mbaud) encoded data, the peak amplitude is in the range of the fundamental frequencies 500 KHz and 1 MHz. Nevertheless, there is still significant energy contained in the odd harmonics above 1 MHz.

Referring again to Figure 4, it can be seen that in terms of candidate frequency spectrum for High-Speed 1553, the band with the lowest channel attenuation is the region below 20 MHz. Unfortunately, this is also the band containing the highest levels of harmonic energy from the 1 Mb/s MIL-STD-1553 signal. This encroachment by the legacy 1553 signal significantly reduces the amount of usable spectrum for High-Speed 1553 signaling.

Therefore, in order to maximize the available High-Speed 1553 frequency spectrum and thus the data rate during concurrent operation, it is necessary to deal with the effects of the 1 Mb/s 1553 signal. In a High-Speed 1553 receiver, the amplitudes of the 1 Mb/s 1553 signal harmonics can be significantly reduced by high-pass filtering the composite (1553 + High Performance 1553) signal.

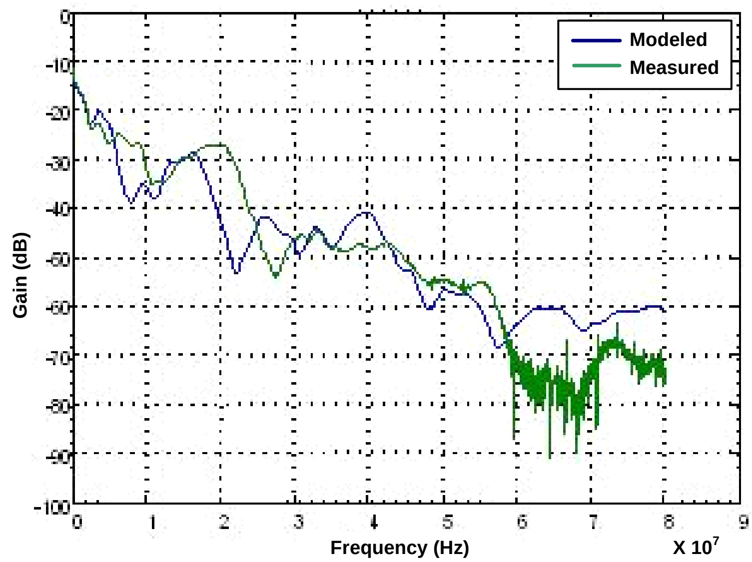

Figure 4. Modeled and Measured Response of 1553 Bus Channel

For transmitting higher data rates over legacy MIL-STD-1553 cable plants, there are a number of issues that need to be addressed. These include channel bandwidth capacity, EMI constraints, noise, and the presence of 1 Mb/s MIL-STD-1553 signals. Channel capacity is the result of bandwidth limitations of legacy 1553 cables, stubs, and couplers. Signal level limitations are the result of MIL-STD-461 EMI constraints; the presence of 1 Mb/s 1553 signals compounds this by further reducing the allowable High Performance 1553 signal level. Sources of noise include transmitters and receivers connected to a multi-drop 1553 bus, as well as conducted and radiate transients from external sources.

HY-PER 1553

Modern telecommunication modulation techniques provide the means for transmitting much higher data rates over existing 1553 buses. For example, in 2005, Data Device Corporation (DDC) designed, built and demonstrated its “Hy-Per 1553” technology. This demo was developed prior to the Air Force’s controlled release of MIL-STD-1553B Notices 5 and 6. For this project, DDC’s implementation consisted of a custom analog front end and FPGA logic. In December 2005, these cards were flown in the EMI environment of an F-15 aircraft. For this test, DDC was able to demonstrate 40 Mb/s operation with concurrent 1 Mb/s 1553 traffic on the same bus cable, and 80 Mb/s without concurrent 1 Mb/s traffic.

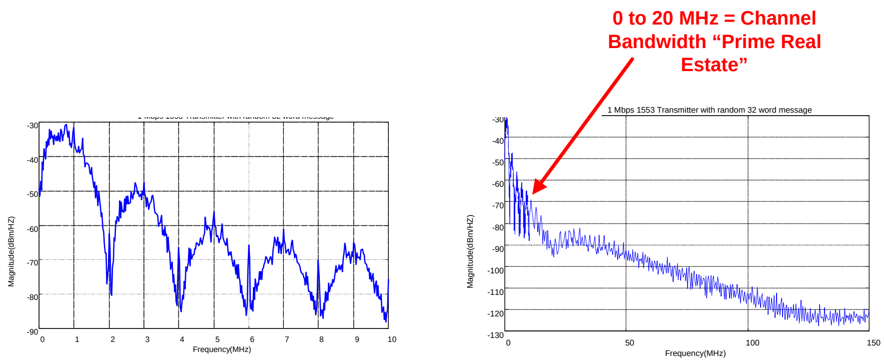

Figure 5. Frequency Spectrum of 1 Mb/s MIL-STD-1553 Signal (32-word message with random data)

For this demonstration, the cards operated for several hours with no errors recorded, thereby meeting the industry-recognized bit error rate benchmark of less that 10^{-12}. In lab tests subsequent to this flight, DDC has demonstrated Hy-Per 1553 with data rates up to 150 Mb/s with concurrent 1 Mb/s 1553, and up to 200 Mb/s without concurrent 1 Mb/s traffic, all with bit error rates of less than 10^{-12}.

Subsequently, the U.S. Air Force has issued controlled distributions of MIL-STD-1553B Notices 5 and 6. This defines a technology for achieving higher speeds over 1553 cable that is in some ways similar, but in other ways different than DDC’s Hy-Per 1553.

Hy-Per 1553 and Notices 5 and 6 are candidate technologies for avionics upgrades and weapons interfaces. For weapons, these technologies can increase data rates, while eliminating the need to run new cables to weapon stations. Moreover, the use of Hy-Per 1553 and/or Notice 5 and 6 interfaces can allow the same weapons to be used on both old and new platforms.

Fibre Channel

Fibre Channel is a very high-performance data transfer standard that is widely used in high-end commercial SAN (storage area networks) data storage applications. SAN networks enable access to storage resources, which are shared among multiple servers. The capacity of SANs is typically in the range of tens or hundreds of terabytes. Key functions of SANs include the storage and retrieval of critical data files, with most storage on magnetic media. Other functions include real-time data mirroring to ensure reliability, data backup, work redistribution in case of failures, and disaster recovery.

In addition to commercial SANs, Fibre Channel is deployed on a number of military/aerospace platforms and programs. These include F/A-18E/F, F-16, F-35, B1-B, B-2, E-2D, the Apache Longbow and MMH helicopters, and AESA

Radar. Applications for Fibre Channel include mission computers, processor and DSP clusters; data storage; video processing and displays; sensors such as radar, FLIR, video, and IFF; and serial backplanes.

Data Rates and Media

Fibre Channel operates at multiple data rates, including 1.0625 and 2.125 Gb, with 4.25 Gb/s starting to emerge, and 8.5 Gb and 10.3125 Gb/s on the near-term horizon. Fibre Channel provides auto-speed negotiation, which enables connected ports on a network to operate at the highest common data rate. Fibre Channel also offers a choice of copper or optical media, with options for twisted/shielded or coax copper cable, along with options for multimode or single mode fiber optics.

For either Fibre Channel or Gigabit Ethernet, depending on the application, there are some factors favoring copper media, and other factors favoring the use of fiber optics. For a 1 or 2 Gb network entailing limited distances, the use of copper provides a lower cost solution. Twisted/shielded cable can support 1 Gb links up to 30 meters, and 2 Gb links up to 10 meters. In addition, relative to optical networks, copper interfaces generally have lower initial installation, maintenance, and training costs.

On the other hand, the use of fiber optics provides lower weight and eliminates EMI emissions and susceptibility. In addition, optical media supports longer link distances and provide an upgrade path to higher data rates. For 2 Gb operation, use of 50 µm multimode fiber can support link distances up to 300 meters, while 1310 µm single mode fiber can support 2 Gb links up to 10,000 meters or farther.

Topologies, Addressing and Routing

Referring to Figure 6, Fibre Channel supports multiple topologies. For avionics, point-to-point links are used for high-speed sensor interfaces, and for connecting between compute nodes and storage devices. Arbitrated loops provide low-cost networks, in which it is possible to establish temporary full-duplex connections between two nodes and a loop. It is also possible, using selective replicate multicast, for one node to transmit to multiple nodes on a loop.

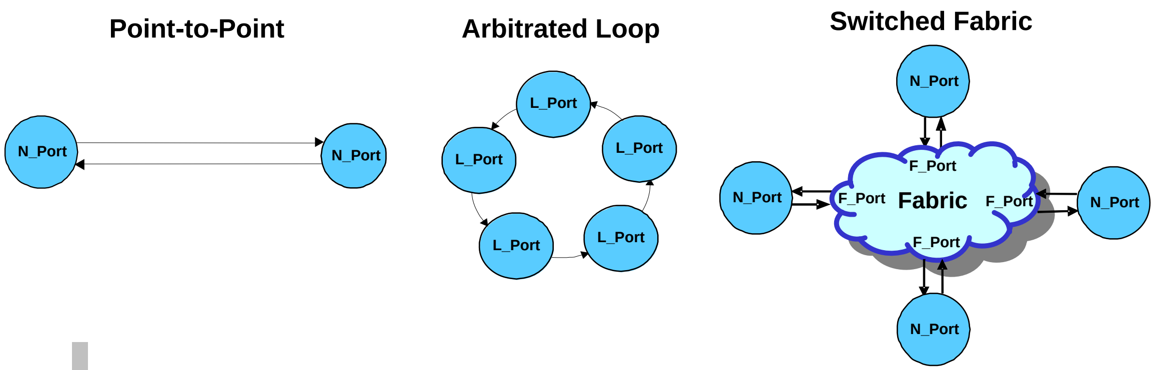

Figure 6. Fibre Channel Topologies: Pont-to-Point, Arbitrated Loop, and Switched Fabric

For Fibre Channel, switched fabric networks provide the highest level of performance. Fabric networks allow any number of nodes to be transmitting and receiving simultaneously. Fibre Channel supports an addressing space of up to 2^{24}, or over 16 million ports. To leverage this, it is possible to scale the size of the fabric by cascading multiple switches together.

Fibre Channel supports multiple modes of addressing and routing. These include unicast, multicast, broadcast, and hunt groups. Unicast entails the transmission between a pair of ports on a network. Multicast is a mechanism for one-to-many messaging, for example, for routing the output of a sensor to multiple processing nodes simultaneously. While multicast capability is generally not provided by commercial Fibre Channel switches, it is commonly implemented on switches designed for avionics use.

Broadcast addressing enables one node to be able to transmit to *all* other nodes on a network. Hunt groups allow a port to transmit to a *choice* of one of multiple ports: to Port A *or* Port B, etc. This improves performance for a multi-switch fabric by allowing traffic to be routed over multiple paths, depending on availability, and/or to whichever port on a redundant end point is currently available. Further, for the case of failed cables, switch ports, or node ports, hunt groups provide an inherent method of rapid self-healing of a network.

Classes of Service and Flow Control

Fibre Channel provides multiple classes of services. Class 3 provides an unacknowledged packet-switched service, in which each Fibre Channel frame is routed independent of other frames involving the same or other source and destination addresses. Class 2 also provides a packet-switched service, but with the added feature of end-to-end flow control, as illustrated in Figure 7.

The payload for each Fibre Channel Frame contains up to 2112 bytes of payload. Multiple successive frames sent by the same port comprise Fibre Channel Sequences, while Fibre Channel Exchanges consist of series of Sequences transmitted by either the same end node or alternating between a pair of end nodes. Upper layer protocol mappings define the series of Sequences comprising individual Exchanges.

Fibre Channel also includes options for circuit-switched service. With Class 1, a virtual, full-duplex, dedicated circuit is established between two endpoints. Class 6 is similar to Class 1, except that it supports reliable (acknowledged) multicast, by establishing dedicated paths from a sending port to multiple destination ports.

The most commonly used service class for both commercial and military Fibre Channel networks is Class 3. In addition to its simplicity, Class 3 enables the highest throughput performance by eliminating the need to pass acknowledgements across a network. Since commonly used protocols such as SCSI and TCP/IP include their own acknowledgement mechanisms, in these cases, end-to-end flow control becomes superfluous.

Figure 7 illustrates one of Fibre Channel's key features, flow control. As shown, buffer-to-buffer flow control (R_RDY(s)) provides a mechanism to throttle the flow of frames between connected ports. This prevents data frames transmitted by source N_Ports (nodes) to switches from overflowing the switch's input frame buffers, and similarly prevents the switch from overflowing the destination N_Port's frame buffers. For all Classes except 3, end-to-end flow control (ACK frames) provides a means for a receiving end node (N_Port) to acknowledge back to the transmitting end node. End-to-end flow control, which can be applied on either the Frame or Sequence level, provides closed loop acknowledgement, and signals to transmitting ports that they are able to free up memory resources.

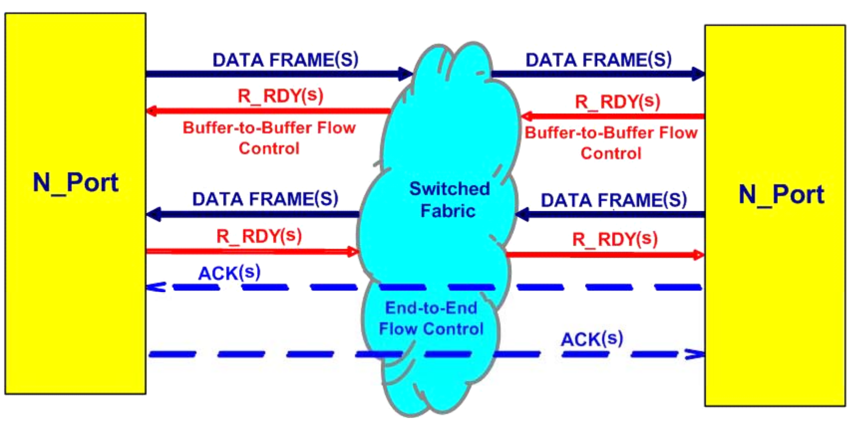

Figure 7. Fibre Channel Buffer-to-Buffer and End-to-End Flow Control

To support the transmission of large data structures such as program files, Fibre Channel provides segmentation and reassembly services at a low level. Segmentation divides a large payload into smaller pieces (Frames) for transmission at the sending end. Reassembly consolidates and if necessary re-orders received Frames into contiguous structures at the receiving end. This includes capability to support out-of-order Frame delivery, which can occur on multi-switch fabrics.

Priorities

Fibre Channel includes a provision for message prioritization. This allows any of 127 priority levels to be assigned to individual messages. By doing so, a fabric switch is able to route high priority messages ahead of low priority traffic. In a real time network, it is advantageous to assign higher urgency to certain messages, either because they are scheduled more frequently – this helps to enforce deterministic scheduling – and/or they are highly imminent in nature; e.g., “the missile is approaching”. Prioritization, which is supported on most military Fibre Channel switches, provides a mechanism to ensure that messages that are transmitted closer to their “deadline” times are routed through a network with minimal delays.

Login

Login provides a means for initializing switches and nodes on a network. This enables node and switch ports to operate using a common set of capabilities, or “login parameters”. Port login enables two end ports to negotiate their parameters, while fabric login enables an end port to negotiate its parameters with the switch F_Port to which it is connected.

There are two forms of login, explicit and implicit. Explicit login, which is generally used for commercial Fibre Channel networks, enables “plug-and-play” operation by allowing nodes and switch ports on a network to “learn” each others’ capabilities during initialization time. On the other hand, with implicit login, node and switch ports “know” each others’ capabilities a priori (enforced by a system spec), thereby allowing them to skip the login process. Since implicit login accelerates network initialization – and possibly more important, re-initialization time following in-flight power cycling – it is the preferred choice for most avionics networks.

Clock Sync

Fibre Channel Clock Sync provides a means for supplying a highly accurate common time base across multiple systems. This service supports resolution of up to 64 bits, with an LSB value of 73 pS., and a rollover time of 43.1 years. In avionics, clock sync may be used for the scheduling of messages, and/or for synchronizing of system events.

Upper Layer Protocols

The suite of ANSI Fibre Channel standards are sometimes described as a “toolbox”, in that they allow individual implementations to be highly optimized or “profiled”. That is, they allow users to choose from among multiple physical layers, data rates, topologies, classes of service (packet vs. circuit-switched, unacknowledged vs. acknowledged, etc.), various login and error handling policies, etc. One of the significant degrees of freedom involves the choice of upper layer protocols, or ULPs.

The primary ULP for commercial Fibre Channel applications is SCSI. SCSI is also used for military Fibre Channel applications, for interfacing with on-board storage devices. For example, the NATO STANAG-4575 standard specifies Fibre Channel with SCSI protocol for use at reconnaissance ground stations for accessing imagery data from solid-state memories or disk arrays.

Another ULP used in avionics Fibre Channel networks is TCP/IP. TCP/IP is a widely deployed protocol which provides network addressing, acknowledged operation, bridging to other types of networks, and operates seamlessly in conjunction with many commonly used operating systems, services and software programs. For operations such as real time video, which do not require acknowledgement, a related “lighter weight” protocol, UDIP/IP, is commonly used.

FC-AE-ASM, or Anonymous Subscriber Messaging is a “military-specific” ULP that defines “lightweight”, unacknowledged operation. In this context, “lightweight” refers to ASM’s minimal header size (16 bytes) and relatively low level of protocol functionality. This protocol enables near “line rate” sustained data rates (e.g., 206 MB/s at a 2 Gb/s signaling rate) and low memory-to-memory latencies. ASM features a 32-bit MESSAGE_ID field, used to associate messages with specific system-defined functions, constructs for security, and message sizes up to 2^{24} bytes. ASM, which uses implicit login, operates in an environment in which multiple “publisher” nodes transmit individual messages to groups of one or more “subscriber” nodes. The “anonymous” aspect of ASM is that “publisher” applications are not “aware” of what nodes are receiving data that it sends, while receiving nodes are not “aware” of what nodes sent specific messages.

Another avionics-specific protocol is Raw Mode. Raw Mode is the ultimate “lightweight” protocol, as it involves the transmission and reception of Fibre Channel Frames and/or Sequences with *no* additional header information. Raw Mode does however leverage all of Fibre Channel’s physical and encoding layers, framing, flow control, error checking, and segmentation and reassembly features. With Raw Mode, it is possible to transmit “line rate” data rates with memory-to-memory latencies on the order of 10 µS. As a result, Raw Mode is well suited for applications such as sensor-to-processor interfacing and multi-processor networks.

FC-AE-1553 provides an extension of the familiar MIL-STD-1553 command/response protocol. It supports all MIL-STD-1553 constructs, including command and status, subaddresses, mode codes, RT-to-RT transfers, extensive error checking, and broadcast, thereby providing a path for the reuse of existing MIL-STD-1553 and MIL-STD-1760 (a weapons interface bus based on 1553) commands and legacy software. In addition, FC-AE-1553 includes extensions and optimizations supporting RDMA, which allows direct access into the memory space of remote systems over a Fibre Channel network; performing file transfers of up to 4.3 GB; and mechanisms for bridging to MIL-STD-1553 buses.

FC-AV (audio-video) supports containerization of multiple video and audio streams, and provides standardized methods for identifying pixel characteristics, lines, frames, frame rates, and color information. As FC-AV provides a great deal of flexibility in terms of the video formats it can handle, it is able to support a wide range of sensors and display heads. By providing a flexible standard for formatting video, the advent of FC-AV dovetails with the current trend to de-couple mission processing, display processing and displays. For example, the ARINC-818 standard for avionics displays leverages FC-AV as its base protocol.

The SAE High-Speed 1760 standard for weapons interfacing was completed in June 2006. High-Speed 1760 specifies use of FC-AE-1553 for command and control messages and file transfers, and FC-AV for video. In addition, High-Speed 1760 defines a physical layer based on 75-ohm coax cable, and also use of FFI, or Fast Fabric Initialization. FFI ensures fast, deterministic initialization and addressing for a multi-switch Fibre Channel network.

Gigabit Ethernet

Gigabit Ethernet derives from 10 and 100 Mb/s Ethernet. Over the years, Ethernet has evolved and maintained its position as the dominant standard for commercial local area networking. As a result, there is a vast base of operating system and application software written for use with Ethernet-based networks.

Military/aerospace applications for Gigabit Ethernet include general networking, computers, data storage, sensors, displays, and bridging to IP-based wireless links and satellites for manned and unmanned air and ground vehicles, and soldiers. In addition, Gigabit Ethernet is used for some “inside-the-box” serial backplane applications. In the future, the latter may migrate to 10 Gigabit Ethernet.

Gigabit Ethernet is being deployed or planned to be deployed on a significant number of military programs, including the Army’s FCS (Future Combat Systems) program, C-17A, C-130 AMP, C-130J, EA-6B, Apache AH-64, Firescout UAV, Future Lynx helicopter, EFV (AAAV) land vehicle, Trident submarine, B-1B, B-2, P-8A MMA, Blackhawk UH-60, B-52 CONECT, and many others.

Ethernet MAC Layer

Ethernet, as standardized by IEEE 802.3, defines the two lowest layers of the OSI stack, the physical layer and the MAC (Media Access Control) portion of the data link layer. As shown in Figure 8, Ethernet frames consist of a start delimiter; 48-bit destination and source MAC addresses; a "length/type" field, which usually contains a value of 2048 ('08 00'h), designating IP protocol; a user payload of 46 to 1500 bytes; and a 4-byte frame check sequence.

Preamble (7)	Start-of-Frame Delimiter (1)	Destination Address (6)	Source Address (6)	Length / Type (2)	Payload Data (46 to 1500)	Frame Check Sequence (4)

Figure 8. Ethernet Frame Structure

A major ingredient in Ethernet's success is its simplicity, which enables low cost implementations for the lower layers' hardware, with higher-level functions typically mechanized by processor software. As the result of its success, Ethernet is used in offices, factories, personal computers, and medical applications. This enables Ethernet users to be able to leverage a vast COTS ecosystem of network interfaces and switches, cables, connectors, operating system stacks, application software, and test equipment.

Early Ethernet networks relied on the collision-oriented CSMA-CD (Carrier Sense Multiple Access – Collision Detect) protocol. While CSMA-CD provided adequate performance at low levels of overall network throughput, at higher throughput levels, the collision protocol was inefficient, due to increased access contention times. However, contemporary Ethernet is based on a switched fabric topology. In addition to eliminating collisions, switched networks enable all nodes to be able to send and receive simultaneously, and as a result provides great improvements not only in determinism, but also overall throughput and latency performance.

Physical Layer

Ethernet provides multiple speeds and physical layer alternatives. 10/100 Mb/s Ethernet includes physical layer options for 10BASE-T, which uses Manchester encoding over two twisted wire pairs; 100BASE-TX, which is based on 4B/5B coding, scrambling to shape the frequency spectra, and Multi-Level Transition – 3-Level Encoding (MLT-3) line coding; and 100BASE-FX, which also uses 4B/5B coding over fiber optic cable, and includes options for 850 nm 1300 nm.

For Gigabit Ethernet, the most commonly used copper option is 1000BASE-T. This physical layer, which operates over four wire pairs, uses 4D-PAM5 (Pulse Amplitude Modulation) encoding to transmit in both directions simultaneously. As shown in Figure 9, 4D-PAM5 encoding transmits 8 bits of data simultaneously at any given time.

Separate symbols are transmitted simultaneously over each of the four wire pairs, with each symbol assuming one of 5 values: +2, +1, 0, -1, -2. There are therefore 625 possible symbol patterns, leaving 512 patterns to encode data, including extra bits for forward error correction; along with 113 that are used as control codes such as Idle (to maintain synchronization), start of packet, and end of packet. In this way, each wire pair transmits at 125 Mbaud, thereby transferring an aggregate data rate of 1Gb/s.

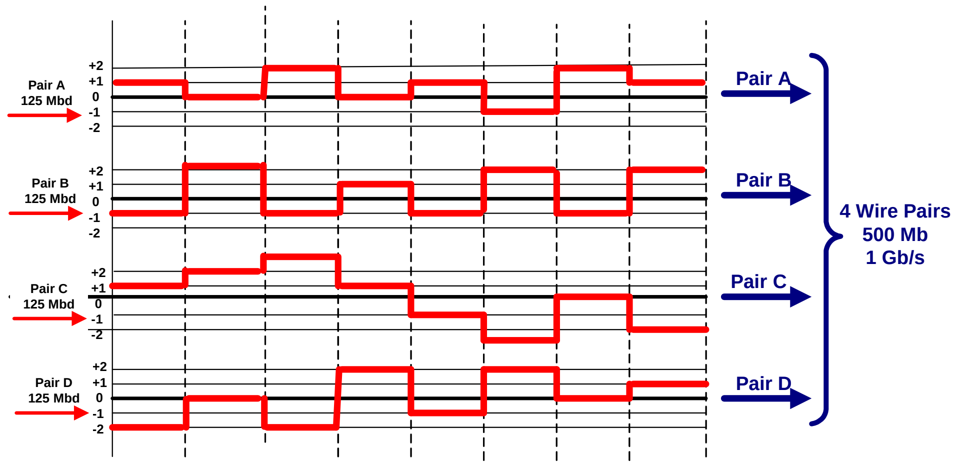

Figure 9. 1000BASE-T Gigabit Ethernet 4D-PAM5 Encoding

Advantages of 1000BASE-T derive from the fact that it operates over standard Cat 5 UTP (unshielded twisted pair) cable, which is widely and economically available. Further, since the symbol rate of 4D-PAM5 PAM is only 125 Mb, this enables operation over longer distances, up to 100 meters or more.

The most common optical physical layer for local area Ethernet use is 1000BASE-SX, which is based on the Fibre Channel physical layer, using 8B/10B coding and operating over 850 nm multimode fiber. Another optical Ethernet physical layer is 1000BASE-BX. 1000BASE-BX provides full-duplex operation over a single fiber using WDM (wavelength division multiplexing). This involves 1310 nm light transmitted in one direction, and 1490 nm light transmitted in the opposite direction over single mode fiber.

IP

Nearly all implementations of Ethernet leverage IP, or Internet Protocol, as their Layer 3. IP is a protocol providing routing services for Ethernet and other local wider area packet-switched networks. Since IP is a connectionless protocol, it does not require any setup procedures; i.e., it is always ready to use.

In operation, the IP header is encapsulated into the first portion of the Ethernet payload. IP accepts data packets from upper layer protocols such as TCP or UDP, and if necessary fragments the data into smaller packets and routes it across a network. In order to route packets to nodes that are on the local network, IP leverages ARP (Address Resolution Protocol), which enables mapping IP address to the MAC addresses nodes on a local sub-network.

The primary function of IP is routing over local area networks or larger networks, including the Internet. Most of the work in packet routing is performed by routers, rather than end points. In part, IP performs routing by dividing larger networks into sub-networks. IP also allows bridging between different types of networks. For example, it is possible to bridge from Fibre Channel to Gigabit Ethernet at the IP level.

IP provides what is called a “best effort” service. This means there are *no* guarantees that data will not be lost due to corruption, or be discarded or lost by the network, or delivered out of order. The benefit of this architecture is that it simplifies and reduces the cost of IP routers, and reduces overheads. The one validation that IP does perform is a checksum verification, and discards packets that fail this check.

Version 4 of IP, Ipv4, provides 32-bit addresses, while the more recent version, Ipv6, supports 128-bit IP addresses. In addition to expanding the size of the IP address field to 128 bits, IPv6 provides improvements in the areas of security, multicasting, and the size of packets that can be sent. The US DoD, which is mandating the use of IPv6 for backbone networks starting in 2008, is looking to greatly expand the number of networking devices in use.

TCP and UDP

For Ethernet networks, the two primary upper layers protocols used to provide end-to-end delivery service are TCP (Transmission Control Protocol) and UDP (User Datagram Protocol). These provide the functions of the Transport layer, which is layer 4 of the OSI stack. TCP is able to support multiple simultaneous connections from different applications. Just as the IP header is encapsulated into Ethernet frames, the TCP header is the first portion of the payload of an IP packet.

TCP accepts data from user applications and passes it to the IP layer. Likewise, data received from the IP layer is passed up to applications. For sending, TCP’s segmentation process divides user payloads into smaller pieces. On the receiving side, TCP reassembles segments into aggregate payloads and passes these up to the receiving application. TCP also assigns sequence numbers to all packets to ensure that none are lost and to enable re-ordering by the receiving end point. In addition, TCP senders formulate a checksum, which is then verified by the receiving end point.

Also on the receive side, TCP issues acknowledgement packets following successful reception. As a means of avoiding congestion on the network, TCP uses “sliding window” flow control. This allows it to transmit a specified number of bytes before receiving an acknowledgement. TCP also includes a round-trip timer. If this timer times out awaiting an acknowledgement, TCP will retransmit data.

Unlike TCP, UDP is a connectionless protocol. It is typically used to send short messages called datagrams. In comparison to TCP, UDP is a “lightweight” protocol. As a result, it does not provide the guarantees that TCP provides, such as error-free in-order delivery, data re-transmission, and flow control. Therefore, with UDP, it is possible for data to be delayed, discarded, or lost when being passed across a network.

In addition, UDP does not perform segmentation or reassembly, and use of its checksum is optional. Conversely, UDP provides a lower cost protocol in terms of hardware gates and/or software computational resources, plus provides higher throughput performance.

For avionics, TCP is an appropriate protocol for applications such as file transfers where data corruption cannot be tolerated, and for critical messages requiring acknowledgement. On the other hand, for applications such as continuously streaming data from sensors or transmissions to some video displays where high throughput, low latency, and low CPU resources are considered important but occasional errors can be tolerated, UDP is often the better choice.

There are a number of available, commonly used applications that run over TCP/IP or UDP/IP. A few of these are:

- *NFS, Network File Server.* NFS, which can operate over TCP/IP or UDP/IP, enables applications running on one computer to be able to access files over a networks as if the files were stored directly on the local machine. As a means for providing access security, this service includes support for user authentication using Data Encryption Standard (DES) encryption and public key cryptography.

- *FTP, File Transfer Protocol.* FTP is another protocol used for file transfers. With FTP, which operates over TCP/IP, one computer operates as an FTP server, while one or more remote computers operate as FTP clients. FTP allows clients to perform such functions as downloading files, saving files, or re-naming files on a server.

- *NTP, Network Time Protocol.* NTP provides a means for synchronizing the clocks of remote computers using UDP port 132. NTP operates using a client-server model, under which server nodes synchronize the time for clients. NTP provides 64-bit resolution, divided into a 32-bit seconds portion, and a 32-bit fractional seconds portion, therefore providing a range of 136 years, with resolution of 233 pS.

TCP/IP can be run over Fibre Channel as well as Gigabit Ethernet. However, in most instances, Gigabit Ethernet will be the better choice of base network for TCP/IP. The reason is that TCP duplicates many of the functions such as segmentation, reassembly, flow control, and acknowledgement that are provided by the Fibre Channel lower layers. In general, running TCP/IP over Fibre Channel rather than Gigabit Ethernet will lead to a penalty in host utilization, and throughput and latency performance.

Exceptions to this would be situations where Fibre Channel is required for accessing Fibre Channel storage devices that are on a common network with workstations running TCP/IP and/or UDP/IP. For this scenario, Fibre Channel would be better suited. In addition, for applications requiring higher speed than Gigabit Ethernet’s 1 Gb/s data rate (but not needing 10 Gb/s), Fibre Channel offers options for 2 or 4 Gb/s. Further, for applications where throughput and latency performance are more important considerations than the need to leverage COTS software, the use of Fibre Channel with lightweight protocols such as ASM or Raw Mode will provide better solutions than a network using TCP/IP.

Quality of Service

For systems requiring performance guarantees, Ethernet includes options for providing variable quality of service (QoS). Two of these options are IntServ and DiffServ. IntServ, or integrated services, provides guaranteed maximum delays based on source and destination IP addresses. IntServ is a fine-grained service, which may be used, for example, to ensure that a video stream reaches a display with a guaranteed maximum delay time. IntServ entails what is called a Traffic SPECification or TSPEC, which is used to configure the parameters for a “token bucket” algorithm. With a token bucket, a (conceptual) bucket or “reservoir” is filled with tokens at an agreed-to, controlled rate. The current “level” of tokens in the bucket dictates the maximum amount of data that may be sent at a given point in time. Once the bucket empties, the sender must wait for more tokens to enter the bucket before transmitting further data. This algorithm supports a level, albeit limited, of data “burstiness”, with the TSPEC parameters being the token rate and bucket capacity.

A second IntSErv parameter, RSPEC, affects the data flow rate through a network. There are three possible settings for this: (1) “Best Effort”, which requires no reservation and provides no guarantees for packet delivery or latency; (2) “Controlled Load”, which provides a moderate level of improvement, in which delay is maintained at a fairly constant level, and packets usually reach their destination; and (3) “Guaranteed”, where delivery is guaranteed and latencies are strictly bounded. One downside of IntServ is that it adds a significant level of complexity to routers.

DiffServ also provides QoS, but at a higher level of granularity than IntServ. DiffServ operates by assigning one of 64 (26) classification levels to individual traffic flows. Packets are classified by network nodes, while routers give preference to higher priority classes of packets. While DiffServ is less complex than IntServ, it cannot provide IntServ's levels of guaranteed performance.

Gigabit Ethernet TOE Adaptors

Most implementations of Ethernet and Gigabit Ethernet are based on hardware implementations of the physical and MAC layers, with the IP and TCP layers implemented by software. For the latter, this usually implies use of operating systems' TCP/IP stacks. While software TCP/IP stacks are adequate for 10/100 Mb/s Ethernet and non-time critical Gigabit Ethernet applications, they can be problematic for the types of real time applications encountered in avionics networks.

The issue is computational resources, principally host loading for implementing the TCP protocol. It is estimated that it requires one GHz of processor bandwidth to support a Gb/s of TCP/IP Ethernet network traffic. In addition to "stealing" valuable computational resources, the use of software TCP stacks add layers of complexity to system software development efforts, such as the need to test and document additional "corner cases".

The use of TOEs, or TCP/IP offload engines, alleviates the "crunch" on host resources by offloading the various TCP operations. These include segmentation, reassembly, TCP checksum, flow control, and acknowledgement. Figure 10 shows an example of a Gigabit Ethernet TOE adaptor, in which the lower "agent" processor offloads the upper host processor by implementing its own on-board TCP/IP stack.

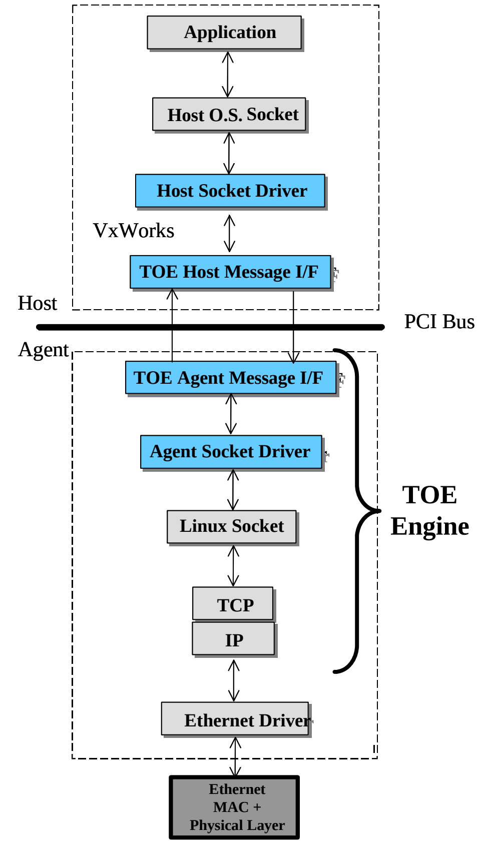

Figure 10. Stack Diagram for Gigabit Ethernet TOE Adaptor

ARINC 664 (AFDX)

ARINC 664, or Avionics Full Duplex Switched Ethernet (AFDX), is a profiled version of 10/100 and Gigabit switched Ethernet, or IEEE Standard 802.3. It is deployed on commercial aircraft, including the Airbus A380 and Boeing 787, and is planned for use on a number of additional commercial aircraft. In addition, there are plans for its deployment on multiple military programs.

AFDX is defined by the ARINC 664 series of standards. This consists of seven parts. In particular, Part 2 profiles the Ethernet Physical and Data Link layers, while Part 7 specifies the requirements of AFDX networks. While currently installed, AFDX networks are all 10/100 Ethernet; future networks will include support of gigabit-speed, at least for some network nodes.

Most of the common Ethernet physical layer options may be used for ARINC 664 networks. These include 10BASE-T, 100BASE-FX, 1000BASE-T, and 1000BASE-SX. However, because of the limited performance characteristics of copper cabling (speed, distance), ARINC 664 Part 2 recommends the use of fiber optic, rather than copper media, for speeds over 100 Mb/s.

One physical layer in particular that is being considered for Gigabit-speed AFDX is 1000BASE-BX. 1000BASE-BX enables reductions in weight and installation cost by providing a full-duplex physical layer over a single fiber using WDM (wavelength division multiplexing). Specifically, this involves

1310 nm light transmitted in one direction, and 1490 nm light transmitted in the opposite direction over single mode fiber.

ARINC 664 defines a highly deterministic network, which is well suited for the types of video, display, sensor, and control applications encountered in commercial avionics networking. AFDX, as illustrated in Figure 13, defines a cascaded star, or switched topology consisting of End Systems (ESs) and switches. End systems interface between individual avionic systems and AFDX switched networks.

For AFDX networks, there are specified maximum latencies for ES transmission and reception, and delays through switches. As a means of ensuring guaranteed delivery, AFDX specifies a highly robust "valid first wins" concept for redundancy, as shown in Figure 11. With this method, an ES always transmits the same frame out of both of its physical ports. These frames are routed through different switches to the two ports of the destination ES. For any frame, the destination ES accepts the first valid frame received, and discards the second one. This enables seamless network operation despite the occurrence of intermittent errors; or failed end systems, cables, and/or switches.

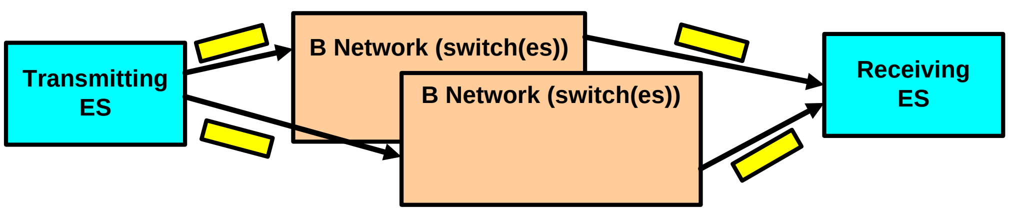

Figure 11. AFDX Redundancy

AFDX defines the concept of "virtual links", or VLs, between senders and receivers. Each VL provides a dedicated path for sending one of more flows of application data from a source ES to one or more destination ESs. AFDX leverages Layer 2 (MAC address) Ethernet routing, where each VL has an assigned destination MAC address. Virtual Links are used to segregate data flows over the network, for the purpose of guaranteeing transmission rates and providing fault isolation.

As AFDX is based on multicast, the MAC for each VL maps to one or more destination ESs. In addition to a source and destination MAC address, for each VL there is an associated minimum and maximum frame length, a value for bandwidth allocation gap (BAG), and a priority level, which is designated as a high or low. The priority feature provides a means to insure rapid routing for time critical messages. For each VL, an AFDX network guarantees in-order delivery, constant bandwidth, and maximum delay times. All of these characteristics are critical in order to ensure the type of deterministic behavior required for safety-critical applications.

All VLs for a given ES are managed independently. As illustrated in Figure 12, for each of an ES's VLs, the BAG value specifies the minimum time between frames that may be transmitted by the ES. As a minimum, all AFDX ESs must support BAG values of 1, 2, 4, 8, 16, 32, 64, and 128 mS. As the BAG specifies the minimum time between frames, it implies a maximum average frame rate of 1/BAG for each VL.

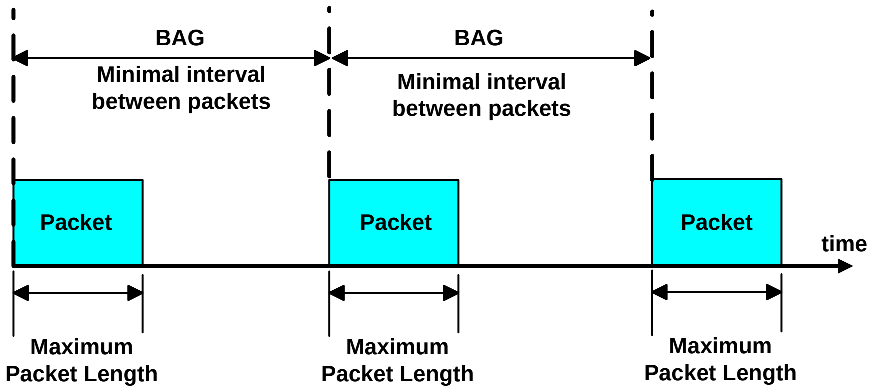

Figure 12. AFDX BAG (Bandwith Allocation Gap)

For operations such as file transfers involving bursty traffic, it is necessary for ESs to perform traffic shaping operations in order to operate within the BAG constraint. This involves implementation of what is called a "leaky bucket" algorithm. With a leaky bucket, frames to be transmitted for a particular VL are in effect queued into a "bucket". While a leaky bucket may be filled by incoming frames at a variable rate, which might be greater than 1/BAG, the bucket imposes a strict limit on the average rate that frames (packets) *exit* the bucket; i.e., the rate that the frames are transmitted over the network. In the case of an AFDX virtual link, this maximum rate is 1/BAG.

AFDX requires the use of UDP/IP, while TCP/IP is optional. It defines two types of UDP ports, AFDX Communication Ports and SAP (Service Access Point) Ports. It is possible to communicate within an AFDX network using either type of port. In addition, SAP ports may communicate through gateways or bridges with nodes on non-AFDX networks. SAP ports enable internetworking between closed AFDX networks and nodes on Layer 3 (IP addressed) networks such as the Internet.

AFDX Communication Ports Support Sampling and Queuing Operations, both of which are defined by the ARINC 653 standard for Avionics Application Standard Software Interface. The sampling service provides compatibility with legacy ARINC 429 (point-to-point(s)) buses. The ARINC 664 sampling service provides a unidirectional link with no acknowledgement, and no hardware flow control. Further, it does not support fragmentation, thereby limiting message sizes to a single frame.

The AFDX queuing service is also unidirectional with no acknowledgement. However, unlike the sampling service, queuing leverages IP fragmentation and reassembly such that it guarantees in-order delivery for sending and receiving multi-frame messages containing up to 8K bytes of application data. To provide reliability, the queuing service includes mechanisms for the ES to signal to its application(s) about any overflow conditions that occur.

In AFDX, UDP SAP ports may be used for file transfers. File transfers are performed using TFTP, Trivial File Transfer Protocol. TFTP provides acknowledged file transfers operating over unicast (point-to-point) virtual links.

In order to avoid switch congestion, and to diagnose problems on AFDX networks, switches are required to implement specified "policing" functions (see Figure 13). These include discarding frames which violate BAG timing constraints, or include addressing errors, checksum failures, or frame length errors. Frame lengths, which are assigned minimum and maximum values on a per-VL basis, have absolute minimum and maximum values of 64 and 1518 bytes respectively. AFDX switches are required to maintain a database of all monitored errors and other network statistics on a per-VL basis. This database is accessible to a management computer by means of the network.

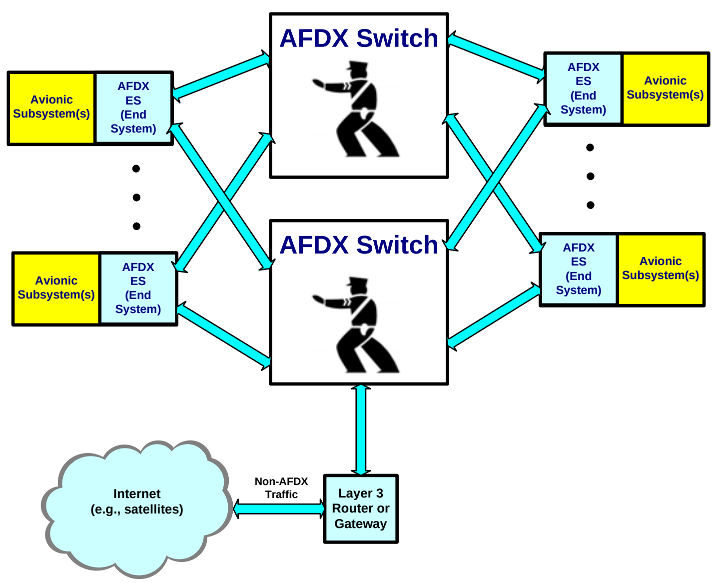

Figure 13. AFDX Network

Since ARINC 664 operates under the assumption of pre-budgeted traffic levels, it is necessary for these to be monitored and enforced on a continuous basis. The intent of switch BAG policing is to eliminate the possibility of traffic congestion in network switches. For example, if as the result of a fault in an ES or upstream switch, one of more VLs are transmitted continuously, it is imperative that the switch discard excessive traffic for that VL, rather than continue to forward it to its destination ES(s). If this discarding operation is not performed, it is possible that other necessary traffic will not be able to reach its destination end systems(s).

ARINC 664 BAG policing uses a "leaky bucket" algorithm, and provides two options for the policing of individual VL BAGs: byte-based and frame-based. AFDX switches must implement either or both of these options, and perform policing with a maximum resolution of 100 µS.

In both cases, ADFX switches maintain an "account variable", AC_i, for each virtual link, VL_i. Referencing Figure 14, each time a frame is received for a given VL_i, the switch determines whether the value of AC_i is above a pre-determined threshold level. If the value of AC_i is above this threshold, the frame *will* be forwarded and the value of AC_i is immediately reduced. For frame-based policing, the value of AC_i is reduced by a fixed amount. For byte-based policing, the value of AC_i is reduced by an amount proportional to the number of bytes in the forwarded frame. After passing the frame, the switch then proceeds to increase the value of AC_i linearly as a function of time. As shown in Figure 14, when AC_i reaches its saturation value, it then levels off.

However, upon receipt of a frame, if the value of AC_i is below threshold, then this indicates a BAG violation; that is, this frame was received too early. In this case, the switch will discard the offending frame, and the value of an MIB (Management Information Base) parameter will be incremented. Similarly, received frames not meeting other policing criteria will result in the incrementing of other MIB parameters for the respective VL_i.

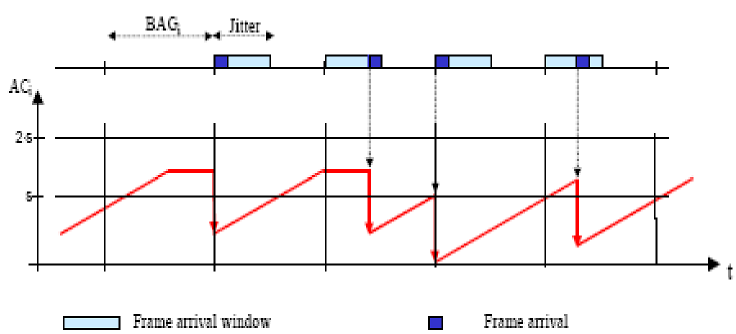

Figure 14. AFDX BAG Policing

In effect, the "leaky bucket" BAG policing monitors for violations in aggregate traffic for each VL. That is, if one frame for a particular VL arrives at a switch a bit later than its nominal time, the algorithm will then allow the subsequent frame for that VL to arrive "early" (i.e., with an inter-frame gap less than the BAG value) without having to discard this subsequent frame.

AFDX switches include an internal end system (ES). This enables external management computers to be able to communicate with the switch. All communication to this ES is performed by "in-band" communication over the

AFDX network. This includes for downloading switch firmware and configuration tables. Configuration tables are used to store data associated with individual VLs, including such variables as MAC addresses, priority, minimum and maximum frame length, and BAG values. In addition to downloading, the ES facilitates communication between the switch's internal monitoring function and management computers. The monitoring function observes all switch functions and logs events such as frames received, BAG and CRC errors, and creates statistics.

The same characteristics that allow AFDX to provide highly reliable operation for commercial aircraft flight control also make it a strong candidate for military applications. As its baseline standard, Ethernet is ubiquitous, available in different speed grades and media options, well understood, and supports a wide base of available application software.

Moreover, AFDX's VL, BAG, and redundancy constructs provide the type of highly deterministic and reliable performance that is demanded for the sensor, control system, computer cluster, and display applications that are common in military avionics. This includes its multicasting capability, along with its guarantees of bandwidth, maximum delay times, and in-order delivery. Moreover, its priority feature allows for the expediting of highly urgent messages, which is a common scenario. Finally, AFDX's capability to bridge to non-AFDX Layer 3 (IP) networks such as satellites, the Internet and the Global Information Grind (GIG) make it an ideal choice for use in network-centric warfare.

Conclusion

For modern avionics networking, it is clear that there is no "one size fits all" solution. Selection of the optimal networking technology is dependent upon a number of factors. MIL-STD-1553 continues to be used for military command and control applications requiring high levels of determinism and reliability, and where 1 Mb/s provides adequate speed. For applications on legacy aircraft where re-wiring is cost prohibitive and/or otherwise impractical, and bandwidths of hundreds of Mb/s are adequate, different forms of High-Speed 1553, including Hy-Per 1553 and MIL-STD-1553B Notices 5 and 6 are beginning to emerge.

For applications where gigabit speed is required and installing new cabling is cost effective, there are multiple alternatives based on Fibre Channel and Gigabit Ethernet COTS technologies. Fibre Channel, particularly when leveraging "lightweight" protocols such as ASM or Raw Mode, provides a robust solution for applications requiring high data throughput, low latency, and a low level of host computational resources. In addition, Fibre Channel, running the SCSI upper layer protocol, provides a very high performance interface to magnetic or solid-state storage systems.

For applications where the foremost consideration is the capability to leverage the wide body of applications that run over standard TCP/IP sockets, Gigabit Ethernet is the best choice. Further, Gigabit Ethernet provides a high degree of inherent compatibility for interfacing to IP-based wireless links and satellite networks. ARINC 664, or AFDX, is a profiled version of Ethernet and Gigabit Ethernet. AFDX, which is deployed on multiple commercial aircraft and starting to see use on military platforms, provides deterministic performance; guarantees of bandwidth, maximum delay times, and in-order delivery; and a highly robust form of redundancy.

Mike Glass
Principal Marketing Engineer
Data Device Corporation

For more information, contact Mike Glass at 631-567-5600 ext. 7409 or Glass@ddc-web.com. Visit DDC on the web: www.ddc-web.com.

Data Device Corporation is recognized as an international leading supplier of high-reliability data interface products for military and commercial aerospace applications since 1964 and MIL-STD-1553 products for more than 25 years. DDC's design and manufacturing facility is located in Bohemia, N.Y.

References

1. Airlines Electronic Engineering Committee; ARINC Specification 664P7; Aircraft Data Network Part 7 Avionics Full Duplex Switched Ethernet (AFDX) Network; Aeronautical Radio, Inc.; Annapolis, MD; 2005.

2. Airlines Electronic Engineering Committee; Draft 4 of Supplement 1 to ARINC Specification 664: Aircraft Data Network Part 2 – Ethernet Physical and Data Link Specification; Aeronautical Radio, Inc.; Annapolis, MD; 2006.

3. ANSI INCITS 332-1999, Fibre Channel-Arbitrated Loop-2 (FC-AL-2); American National Standards Institute, New York, New York; 1999.

4. ANSI INCITS 373-2003, Information Technology – Fibre Channel – Framing and Signaling (FC-DS); American National Standards Institute, New York, New York; 2003

5. ANSI INCITS TR-42-2007, Fibre Channel – Avionics Environment – Upper Layer Protocol MIL-STD-1553B Notice 2 (FC-AE-1553); American National Standards Institute, New York, New York; 2007

6. MIL-STD-1553B Military Standard, Digital Time Division Command/Response Multiplex Data Bus Notice 2; Defense Automated Printing Service; Philadelphia, PA; 1986.

7. Glass, Mike; Data Device Corporation; "How fibre channel can be enabled to meet military-specific requirements"; *Boards and Solutions Magazine*; pages 15-18; ICC Media; Pliening, Germany; September 2005.

8. Goldman, Steve Rood; "Adapting Gigabit Ethernet to Military Systems"; Data Device Corporation; Bohemia, NY; 2006.

9. Goldman, Steve Rood; Data Device Corporation; "Ethernet: A TOE-to-TOE Comparison"; Data Device Corporation; Bohemia, NY; 2005.

10. Hegarty, Michael; "Avionics Networking Technology"; Data Device Corporation; Bohemia, NY; 2005.

11. IEEE Std 802.3-2002; IEEE 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; Institute of Electrical and Electronics Engineers; Piscataway, NJ; 2002.

12. AS5653, Draft v19; High Speed for MIL-STD-1760; Society of Automotive Engineers, Warrendale, PA; 2006.

13. Trend Communications website; "Introduction: Gigabit Ethernet"; Berkshire, UK.