Avionics Networking Technology

By Michael Hegarty
Principal Engineer
Data Device Corporation

©Data Device Corporation

WHITE PAPER

November 2005

Avionics Networking Technology

Abstract

Avionics system architects have a wide variety of commercial off the shelf (COTS) choices for networking technologies. The challenge facing many system architects is choosing a cost effective technology that satisfies the functional requirements of the system. This paper provides insight into some of the key architectural features of two of the most popular emerging networking technologies used in avionics today, namely, Fibre Channel and Ethernet. An overview of each architectural feature is provided along with a discussion of the system level benefits associated with the feature. A new approach called high speed 1553 is also explored. The goal of this paper is not to evaluate these technologies against an arbitrary set of criteria to determine which is better for avionics systems, but rather to provide insight into the benefits of each technology. In reality, no one networking architecture will be "best" for all applications. Each application will contain its own unique requirements that will shape the selection decision. In many cases, the optimal solution may be a combination of several different networking technologies.

1 Introduction

MIL-STD-1553 is a robust serial data bus that has served as the primary command and control data network on board military aircraft for the last three decades. MIL-STD-1553's characteristics of high reliability, high availability, fault tolerance, and high interoperability have made it the data bus of choice for avionics systems. MIL-STD-1553 is still well suited for a large number of avionics applications; however, there are emerging requirements for high speed communication beyond MIL-STD-1553's 1 Mbps rate.

Traditional avionics systems have been implemented based on what is referred to as a federated architecture which consists of a series of independent subsystems which are interconnected with a fairly low speed command and control network (i.e. MIL-STD-1553). Information is processed within each subsystem and the results of the processed data are shared with other subsystems on an as-needed basis. Historically, the subsystems contained relatively slow processors and the amount of data passed between subsystems had been relatively low.

Today's generations of processors are considerably faster than those previously used in avionics systems. Modern avionics systems are now moving towards a more integrated architecture in which multiple subsystems are combined into single subsystems made up of multiple processors. These loosely coupled parallel processors need to share larger amounts of data. In addition, the advent of network enabled warfare and the increased desire to fuse sensor data from multiple sources (including off board data from other platforms or UAVs) are increasing the demand for high speed communication between subsystems on aircraft.

The demand for a high speed avionics data network is fairly obvious. There is a strong trend in the military to adopt commercial off the shelf (COTS) technology rather than designing a custom solution. The problem with COTS is reconciling the differences in use cases (i.e. requirements) between commercial and military applications. This paper will discuss several of the key architectural features of various networking technologies and their applicable benefits in avionics systems.

2 Background

Fibre Channel is the dominant market standard for Storage Area Networks (SANs) in the commercial market. Fibre Channel is used as the I/O interconnect between high end corporate enterprise servers and high speed, high capacity disk drives. In addition to its dominant commercial position, Fibre Channel is also being used in several military aircraft including F-18, B-1B, C-130 AMP, AH-64 Apache Longbow, and the new F-35 Joint Strike Fighter. Fibre Channel has also been selected an SAE working group for inclusion in a future release of MIL-STD-1760 as the standardized high speed data interface between aircraft and stores (smart weapons).

Ethernet is the dominant market standard for local area networks. A big contributor to Ethernet's success stems from its seamless upgrade path from 10 to 100 to 1000 Mbps and its cost effectiveness. Ethernet is now extending its reach to metropolitan area networks (MANs) and possibly wide area networks (WANs). Ethernet is continuing its evolutionary path with the introduction of 10G Ethernet. 100 Mbit Ethernet has been used in several military aircraft including KC-135, C-130, and several helicopter programs[i]. In addition a variation of 100 Mbit Ethernet, as defined in

ARINC 664, is being used by Airbus on the A-380 and by Boeing on the 787 Dreamliner.

3 Analysis of Various Networking Technologies

3.1 Fibre Channel

3.1.1 Implicit Login

3.1.1.1 Description of Feature

Fibre Channel includes features which facilitate a "plug and play" discovery process. The mechanism used by Fibre Channel to advertise capabilities and negotiate methods of communication is through a process called Login. A node may log into the switch fabric and other nodes to establish various communication parameters. Fibre Channel provides three different levels of login: fabric login, port login, and process login. These logins each fall into one of two categories; *explicit login* or *implicit login*.

The explicit login process requires nodes to physically exchange login with each other and the switch through the use of extended link services. Implicit login is a feature which allows selected parts of the system to "assume" configuration options rather than explicitly negotiate them. This allows a system to power-up and begin operating without executing a discovery process. A common use of implicit login in commercial applications is to enable nodes and switches from different vendors with different capabilities to interoperate.

3.1.1.2 Discussion of System Benefit

The system powers up faster. A key benefit of any plug and play is that it allows a system to autonomously react to changes in system configuration. A key disadvantage of plug and play schemes is that they take time to execute because it typically involves a discovery process to identify the devices connected to the network and then a negotiation process in which devices exchange configuration information. These processes take time to execute. One of the characteristics of avionics systems is that there is generally a requirement for systems to power up and be operational within a fairly short period of time. Avionics systems also need to be able to recover from temporary loss of power in a timely manner.

Implicit login allows system designers to selectively bypass the login process and utilize a known set of communication parameters. The ability to bypass the explicit login step is based on the fact that an avionics system will generally be a closed system which is not dynamically reconfigured. For example, if Fibre Channel were used to cross couple two redundant mission computers (i.e. identical computers) which are installed in an aircraft, the login communication parameters could be configured as part of the system integration effort, thus eliminating the need for an explicit login.

The system powers up the same way every time. One of the side effects of a plug and play system is that it may not be configured the same way every time it powers up based on variables such as the order in which devices power up. Repeatable, consistent performance of an avionics system requires that it power up with the exact same configuration options every time. Implicit login allows a system designer to define the communications parameters such that the system will power up with the exact same configuration every time.

3.1.2 Multiple Line Rates

3.1.2.1 Description of Feature

Fibre Channel's base transmission rate is 1.0625 gigabaud per second (1 Gig). The standard has extension of this base rate for 2, 4, 8, and 10 Gbps. Commercially 1, 2, and 4 Gbps are available and the higher rates are expected to be available soon. Fibre Channel also includes a facility called auto-speed negotiation, which allows for plug and play discovery of link speed capabilities.

3.1.2.2 Discussion of System Benefit

Fibre Channel's range of transmission line rates allows a system designer to scale the performance of the network to meet the system requirements. The large number of options between 1 and 10 Gbps provides a cost effective solution for systems that require slightly more than 1 Gbps but don't need the full performance of an expensive 10 Gbps interface.

Fibre Channel's auto-speed negotiation feature allows two or more data rates to be mixed on the same network. This feature also allows for incremental upgrades in data rate on a node by node basis.

3.1.3 127 Level Priority

3.1.3.1 Description of Feature

Prioritization is a common method used to deliver improved quality of service to network communications. Fibre Channel allows communication schemes to be prioritized into one of 127 levels.

3.1.3.2 Discussion of System Benefit

Priority levels supported within Fibre Channel allows a system designer to control the latency of high priority data streams relative to low priority data streams. The large number of levels (127) provides a great deal of flexibility in ranking the relative time criticality of individual data entities.

3.1.4 Flow Control

3.1.4.1 Description of Feature

Fibre Channel provides two methods of flow control. *Buffer to buffer* flow control is used to regulate the rate of communication between a node and a switch. *End to end* flow control is used to regulate the rate of communication between two node ports through a switched fabric (or in a loop or point to point connection).

Buffer to buffer flow control is implemented using very short, low overhead primitive signals (known as ordered sets). End to end flow control is implemented using acknowledge frames.

3.1.4.2 Discussion of System Benefit

Avoid dropped frames due to overruns. In a switch-based system it is possible to transmit frames into the switch faster than they can be delivered to the destination port. Most switches implement ingress buffers which are used to temporarily store frames transmitted into the switch until the frames can be routed through the switch and delivered to the destination port. Congestion in the switch can lead to a case in which the ingress buffers are full, and the destination end port is not accessible (either because all the crossbar routing resources are being utilized or another frame is currently being transmitted to the end port). This buffer overrun condition typically occurs during random peak bursts of network activity and can lead to dropped frames.

It is also possible for a switch to deliver frames to a destination end port faster than the end port can process them. This end port overrun also leads to dropped frames.

The buffer to buffer flow control mechanism in Fibre Channel prevents buffer overrun conditions between end ports and the switch, between switches of a multi-switch element fabric, and between switch and end ports. Each switch port has a defined number of ingress buffers (for the purpose of buffering frames transmitted into the switch by an end port). The total number of buffers is communicated as part of the login process. The receiving port will notify the sending port each time an ingress buffer is freed up, thus allowing the node port to know if there is a physical buffer available. The node port will not transmit a frame into the switch if the switch does not have a buffer available to store it. The tracking of storage buffers allows the system to delay the transmission of frames rather than lose them due to congestion conditions.

Buffer to buffer flow control allows bursts of data to be transmitted at a high rate without the risk of dropping frames.

End to end flow control avoids deadlock conditions. Deadlock conditions can occur in complex systems when there are flow control dependencies between multiple nodes. Situations can exist in which a frame is queued up to be delivered to an end node but the end node has no receive buffers available. If a large number of frames become queued up for delivery to end ports with no buffers available it can create a deadlock condition in which the supply of buffers and routing resources become exhausted. To prevent a deadlock condition, Fibre Channel provides the option to implement end to end flow control in which a sending node will not transmit a frame into the switched fabric unless the node knows that the destination node port has an available receive buffer (there is no sense in sending a frame into the fabric if it cannot be delivered).

3.1.5 Reliable Transport

3.1.5.1 Description of Feature

Reliable transport is a common feature required in many systems. The most common implementation of reliable transport is for the destination to "acknowledge" the valid reception of a data item. The sending node will expect the acknowledgement to be sent within a finite amount of time. If an acknowledgement is not received within a defined period of time then the data item is assumed to be lost and error recovery schemes are invoked.

Fibre Channel provides a dual use for the end to end flow control discussed previously. In addition to controlling the rate at which data is sent between end nodes, the end to end acknowledgement can also be used to facilitate reliable transport. A sending node defines a timeout value for the maximum time it will wait for an acknowledgement after sending a frame or sequence. The timeout value must take into account the latency through the network as well as the receive processing time of the destination port. If the timeout value is exceeded then the sequence or exchange is deemed to have failed and the data can be resent.

3.1.5.2 Discussion of System Benefit

Faster Fault Detection and Recovery. Reliable transport can be implemented in a higher level protocol (such as TCP) or at the application level, but the time required to detect the timeout increases. Fibre Channel, on the other hand, implements the acknowledgement and timeout features at a relatively low level allowing the timeout values to relatively short. Fibre Channel also allows timeout values to be tuned for individual data exchanges.

Avionics systems typically require high reliability, high availability, and low latency. The implementation of reliable transport at a relatively low level in the protocol (normally implemented in hardware in the network controller) provides robust, low latency reliable transport. In addition the fast fault detection also facilitates faster failover to redundant nodes, thus improving the systems availability.

3.1.6 Choice of Copper or Optics

3.1.6.1 Description of Feature

Many people associate Fibre Channel with fiber optics and do not even realize that Fibre Channel is also implemented over copper cables. Fibre Channel includes multiple optical physical layers (including long and short wavelength) and multiple electrical physical layers (including differential pair and coax). The most popular electrical physical layer is 150 ohm twisted shielded pair. The most popular optical physical layer is short wavelength (770 to 860 nm for 1 Gb/s, 830 to 860 for 2 Gb/s) over 62.5 micron (core) multimode fiber.

3.1.6.2 Discussion of System Benefit

Electrical isolation options. Avionics applications typically include either transformer or capacitive coupling to provide electrical isolation for copper interfaces. Fiber optic interfaces have the advantage of providing natural electrical isolation.

Ruggedization. It has been found that copper based electrical interfaces provide a more robust tolerance to harsh environmental conditions than fiber optics. The performance of fiber optics in avionics systems has improved over the years and many military and commercial aircraft are successfully using fiber optic cabling. Methods have been developed to install and maintain fiber optics with an adequately high level of reliability.

Link Distances. Fiber optic interfaces have a distinct advantage in terms of link distances due to the low attenuation rate of the optical signal as compared to an electrical signal at these speeds.

Internal cabling. Electrical interfaces generally have an advantage for interconnects within a box because it allows for the use of traditional backplane connectors. Optical backplane connectors tend to be more expensive and more difficult to maintain. The short distances within a box lend itself to electrical cabling.

Line Rate. Fibre Channel's electrical physical layer is suitable for most 1 Gbps applications and for 2 Gbps applications with short distances. For speeds above 2 Gbps, fiber optics is a more effective physical layer. The ability to scale the line rates to higher link speeds in the future is a strong benefit for optical interfaces.

EMI. Fiber optic interfaces have an inherent advantage in terms of EMI immunity.

3.1.7 Multiple Upper Layer Protocols

3.1.7.1 Description of Feature

One of the key features of Fibre Channel is the support for multiple upper layer protocols. Fibre Channel is fairly unique compare to other data networking technologies in that it has no native command set for transferring data. Fibre Channel provides common communication services such as segmentation and reassembly, flow control, and reliable transport, but it does not provide a direct mechanism for initiating data transfers. Instead, Fibre Channel allows other Upper Layer Protocols (ULPs) to be mapped on top of a lightweight set of communication services. The use of ULPs allows applications to optimize the use of Fibre Channel for a specific function.

Common ULPs include FCP (SCSI), FC-IP (TCP/IP), FC-ASM, FC-AV, and FC-AE-1553. FCP is a mapping of Fibre Channel which is optimized for communicating to disk drives based on a SCSI protocol. FC-IP is an internet protocol encapsulation which allows TCP/IP traffic to be sent seamlessly across a Fibre Channel network. FC-ASM, or Anonymous Subscriber Messaging is a content label based messaging system which is optimized for the distribution of data entities in a real time avionics system. FC-AV is a lightweight encapsulation protocol which allows streaming audio and/or video to be efficiently sent across a Fibre Channel network. FC-AE-1553 provides a mapping for MIL-STD-1553's command/response protocol onto Fibre Channel.

3.1.7.2 Discussion of System Benefit

The ability to *tune* the network protocol through the use of different upper layer protocols allows a system designer to create a highly efficient network. The protocol stacks are relatively small, thus reducing the amount of hardware or software required to implement them and making them higher performance.

The ability to seamlessly mix frames carrying different ULPs in a common switched fabric with a common set of communication facilities (flow control, etc) allows a single network to be optimized to a variety of disparate communication functions.

FC-AV ULP is an upper layer protocol that is ideally suited for transferring streaming video from a display processor to a display head. The current trend in avionics systems is to decouple data processing from display processing and to decouple display processors from displays. This implies that systems will require the ability to distribute streaming video to displays over some kind of high speed network. The ability to route the video streams through a switched fabric is attractive in terms of flexibility and redundancy (ability to route the output of any display processor to any display head). FC-AV supports containerization of multiple video and audio streams, and provides standardized methods for identifying pixel characteristics, lines, frames, frame rates, and color information.

FCP is an upper layer protocol that is ideally suited for Storage Area Network (SAN) applications. Modern avionics systems include large mass storage devices for mission data, map data, terrain elevation data, etc. Fibre Channel's FCP protocol provides a high performance standardized block interface for mass storage devices. STANAG-4575 is an emerging NATO standard that stipulates SCSI/FCP for storing airborne surveillance and reconnaissance data on solid state storage devices.

FC-ASM is an upper layer protocol that is ideally suited for distributing real-time avionics data. Anonymous Subscriber Messaging (ASM) is a publisher / subscriber based protocol that distributes data entities to end systems based on Fibre Channel's multi-cast feature. Multicast is a feature which allows frames transmitted into the switch to be transparently delivered to multiple destinations. The switch implements a multicast server that is responsible for maintaining the distribution lists for multicast groups. ASM allows a system designer to add a new device to the multicast group for a selected data entity without having to change anything at the source of the data. For example an avionics system will include a GPS receiver that periodically sends out position information to subsystems that require it. If a new subsystem is added to the system that requires position information the system integrator can update the multicast group for the GPS data. The new equipment will be added to the appropriate multicast group with no impact on the source of the data (GPS receiver). This publisher/subscriber paradigm provides isolation between the source of data and the consumer of data entities, thus reducing the complexity of system integration. In addition FC-ASM includes security constructs such that the application in the sending node has no knowledge of what nodes are receiving the data and the application in the receiving nodes does not know which node was the source.

3.1.8 Multiple Topologies

3.1.8.1 Description of Feature

Fibre Channel supports three basic topologies: point to point, arbitrated loop, and switched fabric. Point to point allows two Fibre Channel nodes to be connected directly together with no additional external hardware. Arbitrated loop allows nodes to be connected together in a physical ring. Switched Fabric allows multiple nodes to be connected through a physical switch device. An arbitrated loop can be implemented using a hub (basically a static switch) that forms the physical loop. The advantage of a hub is that it provides ability to bypass a faulted or missing node in the loop.

3.1.8.2 Discussion of System Benefit

The support for multiple topologies allows a system to be configured based on a balance of cost and performance. A switch based system is more expensive than a loop based system because of the requirement for a switch, however, a switch based system will provide a much higher level of performance in terms of aggregate bandwidth and lower latency. The availability of these topologies and the ability to create a single “public loop” hybrid network with mixtures of high performance and low-cost subnetworks provides a system designer with an extra degree of freedom in terms of optimizing the cost of the system.

3.2 Gigabit Ethernet

3.2.1 1000Base-T Physical Layer

3.2.1.1 Description of Feature

1000Base-T is a physical layer option within the Gigabit Ethernet standard that allows communication over CAT 5 unshielded twisted pair. The gigabit data stream is divided across four independent twisted pairs within the CAT 5 cable. Each twisted pair has a symbol rate of 125 Mbps. PAM-5 coding is used to increase the number of bits per symbol (2 data bits plus a forward error correction bit). The data can be calculated as: 125 Mbps x 2 data bits per symbol x 4 pairs = 1 Gigabit per second.

1000Base-T also uses 4D 8-state Trellis Forward Error Correction (FEC) coding to reduce the impact of noise and cross talk. 1000BaseT provides full duplex communication through a technique referred to as dual-duplex transmission by transmitting and receiving data simultaneously in both directions at the same time on each of the four wire pairs. A 1000BaseT receiver makes use of echo cancellation techniques to separate the received signal from the transmitted signal on the same wire pair.

3.2.1.2 Discussion of System Benefit

The relatively low symbol rate on each wire pair (125 Mbaud) requires a lower bandwidth transmitted waveform than a Gigabaud signal. The lower bandwidth characteristics of a 1000Base-T waveform lead to several system level advantages.

Radiated Emissions. The radiated emission transfer function for a cable tends to increase at higher frequencies, meaning that higher bandwidth signals tend to radiate a stronger field strength for an equivalent spectral power density of a transmitted waveform on a cable. The advantages of a lower emission level may be offset by the specification limits in MIL-STD-461.

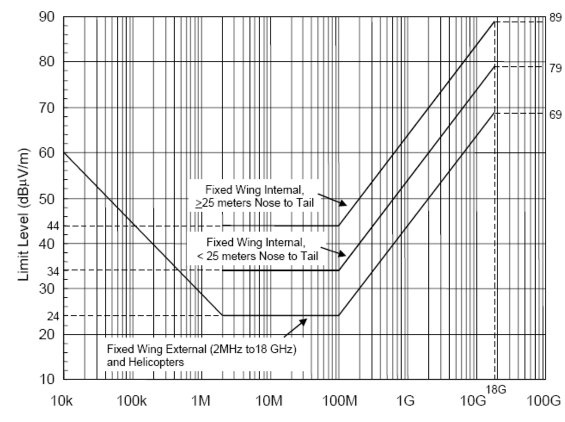

Figure 1 RE102 limit for aircraft and space system applications from MIL-STD-461E

The radiated emission specification limits defined in MIL-STD-461 are lower below 100 MHz than they are above 100 MHz (refer to Figure 1). Therefore a 1000Base-T system may not gain any advantage in terms of radiate emissions because the spec limit is lower below 100 MHz (where the bulk of the 1000Base-T signal power will reside).

Loss Budget. The loss budget of a cable plant can be a critical parameter in an avionics design. The potential for a large number of connectors, in a given signal path, leads to increased attenuation and distortion. 1000Base-T, due to its lower bandwidth signal, equalization, filtering, and forward error correction features, tends to have a rather robust loss budget. 1000Base-T also provides good performance in the presence of impedance mismatches (such as non-controlled impedance connectors).

3.2.2 Choice of Copper or Optics

3.2.2.1 Description of Feature

Ethernet provides options for an electrical physical layer (1000Base-T), and two fiber optic physical layers (1000Base-LX, and 1000Base-SX). 1000Base-T supports link distances up to 100m. 1000Base-SX specifies a short wavelength laser (770 to 860 nm) and supports link distances of up to 275m with 62.5/125 multimode fiber or up to 550m with 50/125 multimode fiber. Ethernet also includes 1000Base-LX which specifies a long wavelength laser (1270 to 1355 nm). 1000Base-T and 1000Base-SX are the two most popular physical layers.

3.2.2.2 Discussion of System Benefit

Electrical isolation options. Avionics applications typically include either transformer or capacitive coupling to provide electrical isolation for copper interfaces. 1000Base-T provides galvanic isolation through the use of isolation transformers. Fiber optic interfaces have the advantage of providing natural electrical isolation.

Ruggedization. It has been found that copper based electrical interfaces provide a more robust tolerance to harsh environmental conditions than fiber optics. The performance of fiber optics in avionics system has improved over the years and many military and commercial aircraft are successfully using fiber optic cabling. Methods have been developed to install and maintain fiber optics with an adequately high level of reliability.

Link Distances. Fiber optic interfaces have a distinct advantage in terms of link distances due to the low attenuation rate of the optical signal as compared to an electrical signal at these speeds.

Internal cabling. Electrical interfaces generally have an advantage for interconnects within a box because it allows for the use of traditional backplane connectors. Optical backplane connectors tend to be more expensive and more difficult to maintain. The short distances within a box lend itself to electrical cabling.

3.2.3 Commercial Software

3.2.3.1 Description of Feature

The prolific use of Ethernet in commercial networking applications has built up a large industry infrastructure of standard commercial middleware and software applications based on the TCP/IP protocol stack. The standardized interface provided by a TCP/IP socket allows third party software applications to transparently make use of the underlying networking interface without having to write custom device drivers.

3.2.3.2 Discussion of System Benefit

The use of Commercial Off-The-Shelf (COTS) software can greatly reduce the non-recurring cost of developing a system by leveraging the economy of scale associated with the commercial software package. The degree to which the use of COTS software will reduce the overall cost of development depends on several factors.

How well does COTS software satisfy the functional requirements of the system? COTS software is developed specifically for commercial applications which may have different needs than an avionics system. For example, Network File System (NFS) provides the ability to authenticate users and machines in a network through the use of Data Encryption Standard (DES) encryption and public key cryptography. The question becomes do the security features within NFS satisfy the security requirements of an avionics system? If it does not meet the security requirements what is the cost of augmenting the COTS solution to fully satisfy the system needs? Is the combined cost of the COTS software and custom extensions less than the cost of developing software specifically for the system? These are some of the challenging questions facing avionics system architects.

Does the COTS solution meet the system's real-time performance needs? Most commercial software applications are not specifically designed with real-time performance in mind. There may be a performance penalty associated with flexibility and open standards. The industry standards, on which many COTS software applications are written, tend to provide support for a broad class of applications. TCP/IP is a good example. TCP/IP can be used as a Local Area Network (LAN) to connect file servers, printers and workstations in a large corporate network. TCP/IP can also be used as a Wide Area Network (WAN) to provide connectivity between corporate networks in multiple cities. TCP/IP can also be used in a home computer to browse the internet through a dial-up or broadband connection.

TCP/IP clearly satisfies a very broad class of applications but what is the real-time implication? The flexibility of TCP/IP that makes is appropriate for such a wide range of applications also adds overhead that may impact the performance of a real-time embedded system.

What is the cost of qualifying COTS software? Software Quality Assurance (SQA) is an aspect of a COTS solution that needs to be considered. The methodology used to write software for a flight control computer on a commercial aircraft that carries hundreds of people is different from the development methods used to write purely commercial software. Commercial software companies do not employ the same levels of SQA that developers of commercial avionics software because the cost is too high. Software cost of quality is often described as being exponential. In general, it is not cost effective for a commercial software developer to eliminate every bug in the application. That is not to say that commercial software is poor quality. The economical reality of commercial software is that there is a point of diminishing returns with regard to debugging and verification.

The question facing the avionics system architect is does the COTS software solution meeting the reliability requirements of the system? The differences in software development and verification methodologies between commercial software and avionics software may be offset by the level of maturity of the product based on the potentially large number of users. Pure COTS applications, through the reporting of bugs and corresponding bug fixes, may supply a high level of reliability but it may be hard to quantify.

3.2.4 Network File System

3.2.4.1 Description of Feature

Network File System (NFS) is a client/server application that provides a client on the network with file access to storage on a server. NFS abstracts the network implications from higher level applications. These higher level applications access the NFS file system over the network the same way they access local storage.

3.2.4.2 Discussion of System Benefit

The use of NFS supports a programming paradigm that is more consistent with commercial software applications. For example, the ability to implement a standard file system on a remotely located mass storage device in an avionics system enables could enable the use of a COTS database engine for the storage and retrieval of mission data, digital terrain elevation data.

NFS is standardized communication protocol between clients and servers. There are numerous COTS implementations of both NFS client and NFS server applications. These applications are typically included as part of an operating system's I/O support package. NFS may be implemented on a system without the need to develop custom software or device drivers.

3.2.5 UDP with Real-Time Middleware

3.2.5.1 Description of Feature

User Datagram Protocol (UDP) is part of the TCP/IP protocol stack. Internet Protocol (IP) implements the network layer which consists mainly of the task of routing packets through the network. The next layer above IP is either the Transport Control Protocol (TCP) or the UDP layer. TCP is a fairly complex layer which implements the transport layer which consists of segmentation and reassembly, flow control, and reliable transmission (acknowledgement and retry). UDP is a simple datagram layer that allows the next higher layer software to send packets of information across the network without the overhead (or benefits) of a transport layer. This shifts the responsibility for transport layer functions to a higher software layer.

3.2.5.2 Discussion of System Benefit

The low overhead of UDP/IP can be leveraged to implement a system that combines the best of two worlds. UDP/IP can be used with real-time middleware to provide a system that leverages the COTS UDP/IP protocol stack and extends its capabilities by performing real-time communication services. The implementation is based on the assumption that UDP/IP is a fairly efficient datagram protocol. The real-time middleware package would add deterministic services on top of the inherently non-deterministic combination of Ethernet and UDP/IP. The result is a system which provides reasonable performance on a commercial networking technology.

3.2.6 IP Network Layer

3.2.6.1 Description of Feature

Internet Protocol (IP) implements the network layer portion of the TCP/IP protocol stack. The main functions of IP include address resolution, determining routes for sending packets, and segmentation (segmentation only when routing to networks with inconsistent packet sizes). IP also includes extensions to aid in Quality of Service (QoS).

3.2.6.2 Discussion of System Benefit

The power of IP lies in its ability to seamlessly route data packets across disparate networks. This allows applications on separate computers to transparently communicate across a heterogeneous combination of networks. The implication for avionics system is that it provides the underlying framework for building the envisioned network of networks. The overwhelming commercial success of the internet demonstrates the raw power of connectivity. The emerging trend within the military avionics community is to provide increased situational awareness.

The first step in achieving increased situation awareness is establishing the required connectivity. IP provides a mechanism for addressing the issues associated with the large variety of networking technologies including onboard networks as well as off-board data links. The question will be whether or not the ubiquitous IP (and its successor IPv6) will provide the necessary performance in terms of quality of service, determinism, security and a host of other issues.

3.2.7 802.3ad Link Aggregation

3.2.7.1 Description of Feature

Link aggregation, as defined in 802.3ad, provides the ability to increase the effective bandwidth between an end system and a switch by combining (or trunking) multiple physical links into one logical link. Other devices on the network send frames to the logical link destination address and the switch will deliver those frames across the multiple physical links. It is completely transparent to other nodes that the designation has a logical link. The inherent features of link aggregation include load balancing and redundancy.

3.2.7.2 Discussion of System Benefit

Scalable bandwidth. Link aggregation provides a method of transparently increasing the effective bandwidth to an end port by adding additional links. Ethernet technology supports 1 Gbps and 10 Gbps link speeds. Link aggregation may be an attractive feature for applications that require more than 1 Gbps but less and than 10 Gbps and don't want to incur the higher cost of 10 Gigabit.

Load balancing. The load balancing attribute of link aggregation allows the switch to spread frames across the multiple physical links that define the aggregated logic link. Load balancing reduces congestion in the switch and increases the overall data throughput to the end node.

Redundancy is a natural extension of link aggregation. Since there are multiple paths (multiple physical links) between the switch and the end node there is inherent redundancy. The link aggregation control protocol provides mechanisms for recovering from and working around failed physical links within a logical link. This allows the system to operate in the presence of a fault, although with degraded performance because of the lower bandwidth due to the absence of the failed link.

3.2.8 1G to 10G

3.2.8.1 Description of Feature

Gigabit Ethernet's various physical layers provide a 1 gigabit per second data transmission rate. The next generation of Ethernet beyond Gigabit Ethernet is 10 Gigabit Ethernet (10G Ethernet). The physical layers defined in 10G Ethernet provide a 10 gigabit per second data transmission rate. 10G Ethernet's physical layer is primarily fiber optic. 10G Ethernet (802.3ae) defines seven optical physical layers (10GBase-SR, 10GBase-SW, 10GBase-LX4, 10GBase-LR, 10GBase-LW, 10GBase-ER, 10GBase-EW). The optical physical layer which is most likely to be applicable to avionics system is 10GBase-SR (850 nm short wavelength).

There is a 10GBase-CX4 (802.3ak) which utilizes 8 twinax cables and only operates over very short distances. The use of 10GBase-CX4 is rather limited. There is work being done on developing 10GBase-T (802.3an) which will run over 4 twisted pair (cat 6 or cat 7 cable). 10GBase-T is expected to be available in mid-2006.

3.2.8.2 Discussion of System Benefit

The existence of a growth path for Gigabit Ethernet is important. The ability to scale a system as bandwidth needs increase in the future is mandatory. 10G Ethernet is relatively expensive today compared to Gigabit Ethernet. The relative cost of 10G Ethernet is expected to decrease over time as more commercial applications start adopting 10G. Today it is more economical to utilize multiple Gigabit links (link aggregation) rather than a 10G link. Economies of scale will reverse the trend and 10G Ethernet will become a more economical high bandwidth interconnect.

Today the only effective option for implementing 10G Ethernet is with fiber optics. Fiber optics has been shown to work in military avionics systems. Over time there may be electrical alternatives (such as 10GBase-T) that are also viable for avionics systems.

3.2.9 Quality of Service (Priority Tags, DiffServ, and Token Bucket)

3.2.9.1 Description of Feature

The convergence of voice, video, and data over a common network is driving the need for at least a rudimentary level of Quality of Service (QoS) in commercial networks. The parameters that constitute QoS include bandwidth, latency, jitter, and loss. Several features exist within commercial Ethernet which help deliver a basic level of QoS. The most common of these features are priority fields and differentiated services (Diffserv). Priority fields (also know as class of service) is a layer 2 feature (i.e. 802 layer) while diffserv is a layer 3 feature (i.e. IP layer).

In 1998 the 802.3ac standard was released defining an extension to Ethernet called Virtual Local Area Network VLAN Tagging. The format of the actual VLAN Tag control information is defined in 802.1Q. The VLAN Tag protocol, defined in 802.1Q, includes a three bit User Priority Field. The use of the eight level Priority Field is defined in 802.1p.

Internet Protocol includes an 8-bit Type Of Service (TOS) field that includes a 3-bit precedence field and 5 other bits. RFC 2474 (Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers) redefines the TOS field to be a 6-bit Differentiated Services Codepoint

(DSCP). Routers use the DSCP to apply scheduling behavior, also know as Per Hop Behaviour (PHB).

A Common use of Priority Fields and Diffserv is to dedicate physical queues within switches and routers for various priority levels. The queue servicing mechanisms employed are typically strict priority, weighted round robin or both. The combination of separating packets packet classes into separate physical queues and servicing policies allows for real-time performance tuning of the network. Most managed switches provide the ability to control physical allocation and servicing policies.

Other QoS facilities available in a TCP/IP system is mapping TCP and UDP port numbers to physical queues (switch or router management option), and traffic shaping. Traffic shaping can be done through several methods. A more common method is through the use of token bucket algorithms.

Token bucket algorithms define a method of bandwidth management that is regulated by end nodes and policed by the switch. The token bucket algorithms provide provisions for regulating not only the average transmission rate from a node but also the maximum burst size. It is important to regulate maximum burst size in addition to average transmission rate. Sporadic bursts can lead to congestion and delay through switches, and, in extreme cases, can lead to overrunning the physical queues in the switch (i.e. loss of data). The switch can monitor or police the traffic from each node and may drop or reclassify packets that don't follow the traffic shaping profile defined by the token bucket algorithms.

The "token bucket" analogy uses tokens as an indication of permission to transmit data into the network (i.e. into the switch). The transmitting agent within an end node accumulates tokens (or permissions to transmit bytes) at a rate referred to as the Committed Information Rate (CIR). The transmitting agent within the end node may not accumulate an indefinite number of tokens. The maximum number of tokens that a transmitting agent may accumulate is referred to as the Committed Burst Size (CBS). If a transmitting agent accumulates the maximum number of tokens (i.e. fills the bucket) then new accumulated tokens are discarded (i.e. once the bucket is full any additional tokens will flow over the top and be lost). As tokens are consumed (i.e. data is transmitted) the level of the bucket will fall below the maximum and more tokens may be accumulated.

3.2.9.2 Discussion of System Benefit

Quality of Service is an important attribute of an avionics system. In order for the avionics system to deliver predictable performance the data transferred through the networks needs to be predictable. The level of predictability (or determinism) in a system is a function of the system architecture and functional requirements.

Priority Tags, Diffserv, and port mapping (Class of Service) are common features found in most higher-end managed switches. Token bucket algorithms are a slightly more exotic feature and not generally implemented in standard commercial hardware and software. Class of Service features, although very common, provide a lower level of determinism than token bucket algorithms. System architects have several options to work with in defining a system that delivers the required level of determinism.

3.2.10 iSCSI

3.2.10.1 Description of Feature

Internet SCSI (Small Computer System Interface) defines a method of implementing a storage area network (SAN) over IP. Today Fibre Channel is by far the dominant market standard for storage area networks (i.e. the connection between enterprise servers and storage systems). The combination of Ethernet and TCP/IP is the dominant market standard for Local Area Networking (LAN) in corporate enterprise applications (i.e. the connection between servers and work stations).

iSCSI is an attempt to force a convergence of the LAN and SAN networking domains into a single common technology – Ethernet. There are advantages of maintaining a single network in terms of commonality of hardware and in-house technical expertise. However, iSCSI has shown to provide lower performance than Fibre Channel. The performance limitations of iSCSI are being addressed through the use of Remote DMA (RDMA) or more specifically IP extension for RDMA (iSER). However, the implementation of iSER eliminates the cost advantage of Ethernet and still does not provide as high a level of performance as Fibre Channel. It appears, at least for the near term, that Fibre Channel will remain the dominant market standard for storage area networks.

3.2.10.2 Discussion of System Benefit

iSCSI is applicable to SAN applications which include a file server as opposed to Network Attached Storage (NAS) applications that utilize protocols such as Network File System (NFS) or Common Internet File System (CIFS). The pros and cons of a SAN (or server) based system versus a NAS based system is a raging debate in the commercial market. SAN protocols, such as iSCSI and Fibre Channel, are block oriented while NAS protocols are file oriented.

The tradeoff between NAS and SAN centers around the type of communication between the client and the server. If the server is simply acting as a volume manager and all I/O between the client and server is file oriented then NAS could be used. If, however, the server is performing other functions, such as database, mail server, or authentication then a server with a SAN is more appropriate.

Mass storage applications within avionics systems are most likely to be file oriented. NFS is an attractive method of implementing a networked storage system because it eliminates the need for a separate file server. The NFS server acts as both the storage device and a simple file server. If the system requires the functionality of a full blown file server then Fibre Channel would provide a more efficient SAN interface in terms of performance and cost-effectiveness.

3.2.11 Simple Hardware, Complex Software

3.2.11.1 Description of Feature

One of the key architectural features of Ethernet implementations is that the hardware portion of the interface (i.e. Ethernet MAC and PHY) is relatively simple while the software portion is more complex than other networking technologies. The rationale behind this method of implementation is economics and the attributes of the target commercial market.

Ethernet has become a pervasive technology based on its cost effectiveness. The hardware portion of an Ethernet interface consists of a Media Access Controller (MAC) or link layer function and a Physical layer function (PHY). The amount of logic required to implement these functions is relatively small and the production volumes are very high, thus the cost is low. The penalty for having inexpensive hardware is that there is an increased burden on the host processor to implement the vast majority of the communication protocol (i.e. TCP/IP).

The commercial market for Ethernet is very cost sensitive. The commercial applications that utilize Ethernet, such as Windows based workstations, have sufficient excess processing capacity that the software-centric nature of Ethernet and TCP/IP is very cost effective. For example, a Windows based workstation running a word processing program to edit a file on a network file server is not using a significant portion of the bandwidth available on its multi-gigahertz processor. The computer has more than enough spare MIPs that it is very efficient to implement as much of the protocol as possible in software.

3.2.11.2 Discussion of System Benefit

The characteristics of embedded processors used in avionics systems are different from those used in typical commercial applications. Avionics system tend to real-time and “compute intense.” The compute intense nature of embedded real-time avionics systems is at odds with the software-centric nature of TCP/IP protocol processing. For avionics applications that require relatively low data transfer rates over the Ethernet network, the software-centric model of TCP/IP is very cost effective.

There is a growing trend in avionics systems to pass larger amounts of data between subsystems. As the data rates increase, the burden on the processor increases dramatically. In extreme cases the processor becomes so consumed implementing the TCP/IP protocol that there is insufficient processor bandwidth to implement the intended mission function.

For low data rate systems (such as systems transferring less than 10 megabytes per second) a conventional Ethernet interface consisting of a hardware link layer (MAC) and PHY with a software implementation of TCP/IP is very cost effective. For systems requiring higher data rates (e.g. above 10 MBps) some kind of hardware acceleration or TCP/IP Offload Engine (TOE) is generally required to free up processing bandwidth on the host CPU.

3.2.12 Need for a TOE

3.2.12.1 Description of Feature

The commercial industry has recognized that there applications, such as Windows based workstations, that have more than enough excess processing bandwidth to handle the additional burden of TCP/IP processing. However, there are other commercial applications, such as network file servers, that do not have as much excess processing bandwidth. Enterprise file servers tend to be heavily utilized. These server applications would benefit from a shift of the communication protocol processing from the computer to the network card.

The emerging trend in commercial high performance enterprise servers is to utilize an Ethernet interface with a TCP/IP Offload Engine (TOE). There are several levels of offload. The most common (and lowest level of offload) is the implementation of TCP checksums in hardware. Checksum calculation is a CPU intense function that can be easily implemented in hardware. The implementation of hardware TCP checksum offloading requires that the TCP/IP protocol stack running on the host processor be capable of shifting that function to hardware. Normally TCP checksum calculations are an integral part of the TCP/IP software stack and are not easily moved to hardware accelerators (it generally requires a modification of a third party TCP/IP protocol stack to support the specific hardware implementation). Some commercial operating systems are beginning to provide hooks for adding offload capability (such as Microsoft's chimney architecture) but this is a more difficult task with embedded real-time operating systems (RTOS).

There are other TOE implementations that offload the entire TCP/IP protocol stack to the network interface adapter. These "full offload" devices have the advantage of being able to bypass the host processors TCP/IP stack as opposed to having to modify the stack. Full offload adapters tend to be easier to integrate with an embedded RTOS than partial offload adapters.

3.2.12.2 Discussion of System Benefit

The key benefit of a TOE lies in its ability to reduce the host processor bandwidth required for networking I/O. A TOE will allow more the processor's bandwidth to be applied to the mission function. The reduction in CPU bandwidth spent on networking I/O also improves the real-time responsiveness of the computer and will increase the overall determinism of the processing function. The benefit associated with a TOE has an inherent price associated with it because the cost of the hardware for the network interface controller will increase, however, for many applications the increased cost of the network interface controller will be offset by the value of increased available host processor bandwidth.

3.3 High Speed 1553

3.3.1 Description of High Speed 1553

MIL-STD-1553 is a serial, time division multiplex data bus that has been used as the primary command and control data interconnect in military aircraft for the past three decades. MIL-STD-1553's robust performance, high level of interoperability, large installed based, and well established infrastructure of vendors has made MIL-STD-1553 the network of choice for military avionics systems. While MIL-STD-1553's one megabit per second data rate is adequate for the current avionics applications, there are emerging applications that require higher bandwidth. The challenge facing the military avionics industry is finding cost effective methods of supplementing MIL-STD-1553 with higher bandwidth data communication channels.

Traditional avionics system have been implemented based on what is referred to as a federated architecture which consists of a series of independent subsystems which are interconnected with a fairly low speed command and control network (i.e. MIL-STD-1553). Information is processed with each subsystem and the results of the processed data are shared with other subsystems on an as needed basis. In general, the amount of data passed between subsystems had been relatively low.

The advent of network enabled warfare and the increased desire to fuse sensor data from multiple sources (including off board data from other platforms or UAVs) is increasing the demand for high speed communication between subsystems on aircraft. Satisfying this demand for higher bandwidth communication on existing aircraft requires that either new cabling is installed or higher data rates are run over existing cabling.

Research was conducted by Data Device Corporation aimed at exploring the option of supporting high bit rate transmissions over existing MIL-STD-1553 networks. A combination of empirical and theoretical methods was used to determine if a MIL-STD-1553B network contains sufficient capacity to support high data rate transmissions. The results of DDC's analysis is that, for some MIL-STD-1553 buses, there is sufficient bandwidth to implement a broadband system in which legacy 1 Mbps 1553B waveforms could coexist with new 200 Mbps waveforms, thus providing an incremental high speed communication channel to existing MIL-STD-1553 buses.

DDC is currently engaged in the development of technology to enable high speed data communication over existing MIL-STD-1553 data buses. The goal of this development is to deliver a network interface controller that provides independent access to both legacy MIL-STD-1553 and a new high speed 1553 devices simultaneously on a single 1553 data bus.

3.3.2 Discussion of System Benefits of High Speed 1553

The implementation of new high speed interfaces becomes an economic tradeoff between the costs of adding additional cabling and associated electronics versus the cost of updating the existing electronics to increase the data rates on the existing MIL-STD-1553 buses. The US Air Force has estimated that it would cost approximately one million dollars to rewire a fighter aircraft to supplement or replace MIL-STD-1553 cabling with high speed cabling[ii].

The implication of reusing existing cabling goes beyond the basic cost tradeoff. There is also the aspect of time to consider. Rewiring an aircraft implies a larger scope of work than adding or replacing electronics in existing systems. A major retrofit requires that an aircraft be taken out of service for a fixed period of time while the update is being implemented. Only a small portion of the active inventory of aircraft can be taken out of service at any given time. This implies that it could take a long period of time to field a retrofit. It is not uncommon for the retrofit of a fleet of aircraft to take 5 or 10 years. High speed 1553 is an enabling technology that contributes to the ability to perform incremental upgrades of aircraft.

4 Summary/Conclusion

The tables below provide a summary of the selected key features and benefits of Fibre Channel, Ethernet, and high speed 1553 that have been discussed in this paper.

Table 1. Summary of Fibre Channel Features and Benefits

Fibre Channel	
Feature	**Benefit**
Implicit Login	• Faster system power-up • Fast recovery from loss of power • System powers up same way every time
Multiple Line Rates	• Scale line rate to meet system throughput requirements • Auto-speed negotiation allows for a mixture of line rates on the same network
127 Level Priority	• Quality of Service
Flow Control	• Avoid dropped frames due to congestion • Avoid deadlock conditions
Reliable Transport	• Faster fault detection and recovery
Choice of Copper or Optics	• Flexibility to choose • EMI tolerant fiber optics • Mechanically rugged copper • Optics provides path to higher data rates • Optics support longer link distances
Multiple Upper Layer Protocols	• Tune the protocol to a specific application • FC-AV for audio and video streams • FCP for Storage Area Network • FC-ASM for real-time avionics data entities • FC-IP for transparent bridging to other networks (WAN)
Multiple Topologies	• Flexibility between high performance switched and low cost loop architectures • Allows a mixture of switched and loop nodes

Table 2. Summary of Ethernet Features and Benefits

Ethernet	
Feature	**Benefit**
1000Base-T Physical Layer	• Lower high frequency radiated emissions (EMI) • High loss budget / ability to pass through multiple connectors with minimal distortion
Choice of Copper or Optics	• Flexibility to choose • EMI tolerant fiber optics • Mechanically rugged copper • Optics provides path to higher data rates • Optics support longer link distances
Commercial Software	• Reduced development cost by using COTS • May or may not meet functional, real-time, or quality requirements
Network File System	• COTS solution for Network Attached Storage (NAS) • Eliminate need for a separate file server and storage area network
UDP with Real-Time Middleware	• Low overhead datagram protocol on IP network layer functionality • Real-time middleware provides reasonable performance on a COTS protocol stack
IP Network Layer	• Seamlessly route data packets across disparate networks • Integrate Global Information Grid (GIG) with platform network
Link Aggregation	• Scalable bandwidth • Load balancing • Redundancy
1G to 10G Ethernet	• Path to higher bandwidth
Quality of Service	• Predictable performance • Deterministic communication
iSCSI	• Storage Area Network over IP
Simple Hardware, Complex Software	• Low cost hardware (MAC + PHY) • Utilize excess CPU bandwidth to implement protocol
TCP/IP Offload Engine (TOE)	• Reduce CPU bandwidth required for communication • More CPU bandwidth available for mission function

Table 3. Summary of High Speed 1553 Features and Benefits

High Speed 1553	
Feature	**Benefit**
Operate over existing cable	• Reduced cost of installation for new electronics • Faster installation of new electronics • Possibility of field upgrades • Reduce time required to retrofit a fleet
Concurrent with legacy waveform	• Add new functions without interfering with existing systems

Example Architecture. A system could include a mixture Fibre Channel and Ethernet. The system may include FCP for storage, FC-AV for streaming video from display processors to displays panels, FC-ASM for distributing real-time mission data between loosely coupled processors, Ethernet for connections between embedded processors and man-machine interfaces (workstations), and TCP/IP for wide area network (i.e. interface to data links). In addition high speed 1553 may be used in places where running new cable is cost prohibitive, such as through the wings to pylon mounted weapons.

Avionics system architects are faced with a myriad of choices when it comes to data networking technologies. Countless trade studies have been done to determine what the "best" choice is for avionics networking. The reality is that there is no common best choice for every application. Each system has its own unique requirements. In the end, the optimal solution for a system is probably a combination of several different technologies.

Michael Hegarty
Principal Engineer
Data Device Corporation

For more information, contact Michael Hegarty at 631-567-5600 ext. 7257 or 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. The company's design and manufacturing facility is located in Bohemia, N.Y.

References

•

[i] Rockwell Collins, "IPC-7000 Integrated Processing Center 7000", http://www.rockwellcollins.com/ecat/gs/IPC-7000.html?smenu=103, 14-Nov-05.

[ii] William Wilson, "High-performance and cost-effective avionics network upgrade solution and approach."

Robert W. Kembel, *Comprehensive Introduction to Fibre Channel*, Northwest Learning Associates, Inc, 1998.

MIL-STD-1553 Designer's Guide, Data Device Corporation, New York, 1998.

IEEE 802.3-2002, The Institute of Electrical and Electronics Engineers, Inc, 2002.

Douglas E. Comer, *Internetworking with TCP/IP, 1. principles, protocols, and architecture*, Prentice Hall, 2000.

Ed Taylor, *Networking Handbook*, McGraw-Hill, 2000.