AVIONICS magazine
TECH REPORT

AFDX: The Next-Generation Interconnect for Avionics Subsystems

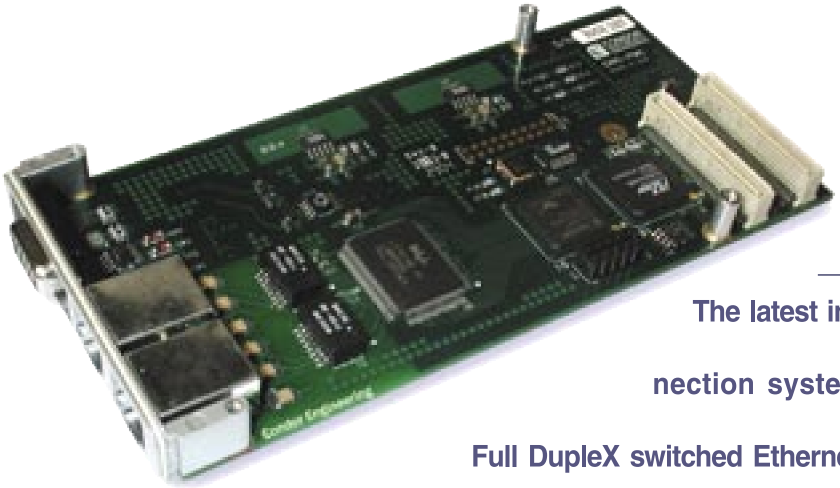

The latest in subsystem interconnection systems, AFDX (Avionics Full DupleX switched Ethernet network), already is on board the Airbus 380. When this huge aircraft enters service in 2006, it will carry aloft electronics that exhibit high-speed data transfer, enhanced reliability, and a reduction in wiring—vital attributes that are inherent in the design topology of AFDX. On the other side of the Atlantic, Boeing plans to install this state-of-the-art communications interconnect system in its B787 aircraft, which is expected to take its first flight in 2007 and begin service in 2008. As for earlier aircraft, AFDX or its offshoot, ARINC 664, Part 7, probably will replace older interconnection systems when retrofits are implemented.

Access Intelligence

AFDX is a standard that defines the electrical and protocol specifications (IEEE 802.3 and ARINC 664, Part 7) for the exchange of data between avionics subsystems. One thousand times faster than its predecessor ARINC 429, AFDX builds upon the original concepts introduced by Airbus. The European aircraft manufacturer devised AFDX and named it, as part of the evolution of its A380 aircraft. As a result AFDX, and its offshoot, ARINC 664, Part 7, have brought a number of highly significant improvements, both electrical and mechanical, to the interconnection of electronic subsystems aboard aircraft.

Many electronic subsystems are on board large aircraft, such as inertial platforms, control systems, sensors systems, and communication systems. They all demand high-reliability, high-speed information transfer. Control systems and avionics, in particular, rely on complete and up-to-date data delivery from source to receiver in a timely fashion. For safety-critical systems, reliable real-time communications links are essential—and that is where AFDX has brought about major improvements. AFDX builds upon a number of earlier bus structures (see “The History Behind AFDX”, page 4).

What is AFDX?

As shown in Figure 1, an AFDX system comprises the following components:

- *Avionics Subsystems*—the many traditional avionics subsystems on board an aircraft, such as the flight control computer, GPS and tire pressure-monitoring system. Together with an AFDX end system, an avionics computer provides a computing environment for hosting multiple avionics subsystems. Each avionics computer contains an embedded end system that connects the avionics subsystems to an AFDX interconnect, as shown in the illustration.
- *AFDX End System*–the system that provides the interface between avionics subsystems and an AFDX interconnect. The end systems guarantee a secure and reliable data interchange with other avionics subsystems. They each export an application program interface (API) to the various subsystems, enabling them to communicate with each other via a simple message transfer interface.
- *AFDX Interconnect*–which is a full duplex, switched Ethernet interface. It comprises a network of switches that forward Ethernet frames to their appropriate destinations. Because it is based on Ethernet technology, AFDX is a departure from the traditional ARINC 429 point-to-point technology and Mil-Std-1553 bus technology. This interface is discussed below, in detail.

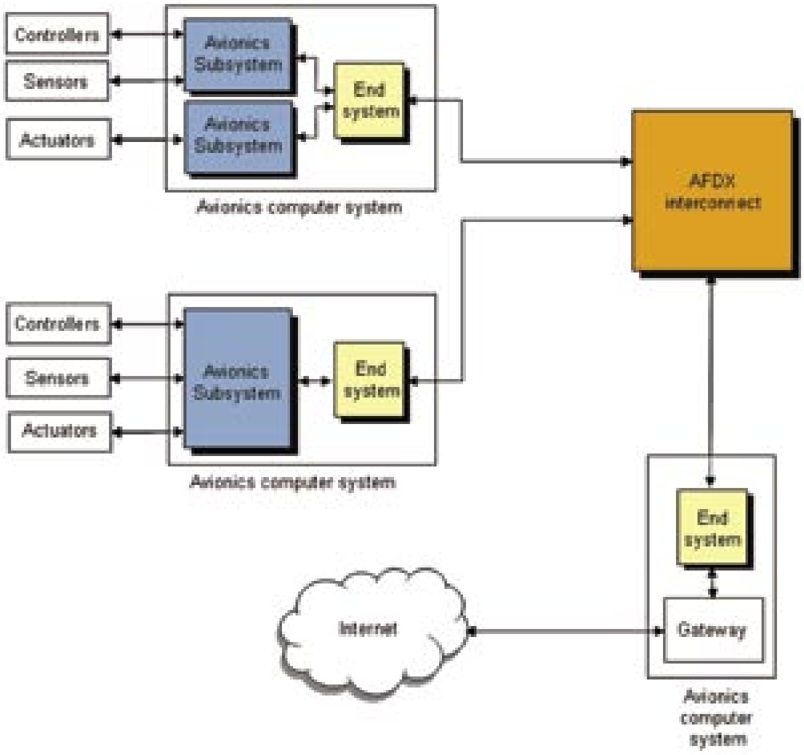

Figure 1— How subsystems interconnect with each other and the outside world.

Also shown in Figure 1, two of the end systems provide communications interfaces for three avionics subsystems, and the third end system supplies an interface for a gateway. The latter provides a communications path between the AFDX and external IP (Internet protocol) networks and typically is used for data loading and logging.

The Problem with Ethernet

Half-duplex switched mode Ethernet (see sidebar, page 4) is another name for the original Ethernet local area network (LAN). There is an issue when multiple hosts are connected to the same communication medium, as is the case with the coaxial cable depicted in Figure 5 (page 4), where there is no central coordination. In this case two hosts could transmit simultaneously, causing their transmissions to collide. A need exists, therefore, for the hosts to be able to detect transmission collisions.

A collision occurs when two or more hosts attempt to transmit at the same time. Each host, then, has to retransmit their data—which opens the possibility of another collision and the need to retransmit again. To avoid this phenomenon, each host must select a random transmission time from an interval for retransmitting the data. If a collision is again detected, the host selects another random time for transmission from an interval that is twice the size of the previous one, and so on. This is often referred to as the “binary exponential backoff strategy.”

Since Ethernet has no central control, the packets of data could, theoretically, collide repeatedly, despite a binary exponential backoff strategy. There is, in fact, a chance that an infinite chain of collisions

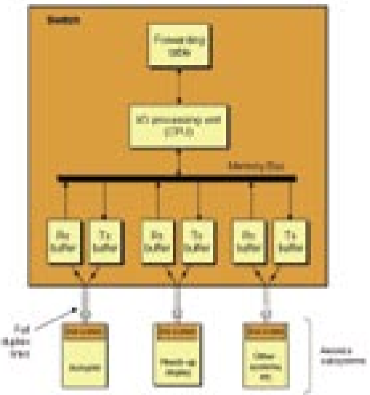

Figure 2—A full-duplex Ethernet switch that solves the collision problem.

could occur, and the packet would never be successfully transmitted. Therefore, in a half-duplex mode, very large transmission delays are possible due to collisions—a situation that would be unacceptable in an avionics data network.

So, what is required—and what was implemented in AFDX—is an architecture in which the maximum amount of time for one packet to reach its destination becomes known. Achieving this meant ridding the system of "contention."

Eliminating Contention

To do away with contention and hence the indeterminacy as to how long a packet takes to travel from sender to receiver, AFDX adopted full-duplex switch Ethernet. It overcomes the issue of collisions inherent in half-duplex based Ethernet. As shown in Figure 2, each avionics subsystem—autopilot, heads-up display, etc.—connects directly to a switch over a full-duplex link that comprises two twisted pairs, one for transmit (Tx) and one for receive (Rx). The switch, which comprises all the components contained in the large box, is able to buffer packets for both reception and transmission.

Figure 2 also shows that both the Rx and Tx buffers are capable of storing multiple incoming/outgoing packets in a FIFO (first in, first out) order. The role of the input/output (I/O) computer-processing unit (CPU)—moving packets from the incoming Rx buffers to the outgoing Tx buffers—is achieved by examining each arriving packet that is next in line in the Rx buffer to determine its destination address (virtual link identifier). The CPU then checks with the forwarding table to determine which Tx buffer(s) are to receive the packet. The packet is subsequently copied into the Tx buffer(s), via the memory bus and transmitted again in FIFO order to the selected avionics subsystem or to another switch. This type of switching architecture is referred to as "store and forward."

Consequently, the full-duplex switch architecture eliminates contention, which causes collisions. (In practice, AFDX mandates that two redundant switch architectures are employed.)

Theoretically, an Rx or Tx buffer could overflow. But this will not happen if the buffer requirements were planned correctly in the original design and made large enough. With AFDX you may not be able to get your message out immediately, but you are guaranteed that it will not be delayed for more than a given time interval, say no more than 400 microseconds.

Instead of collisions and retransmissions, switching architecture can result in jitter due to the random delay introduced by one packet waiting for another to be transmitted. The extent of jitter introduced by an end system and switch must be controlled if deterministic behavior of the overall avionics system is to be achieved.

Reducing Wire Runs

In addition to the enhancements already described, AFDX delivers other benefits when compared to ARINC 429. Some of these distinctions are illustrated in Figure 3. In ARINC 429, as depicted in Figure 3a, is

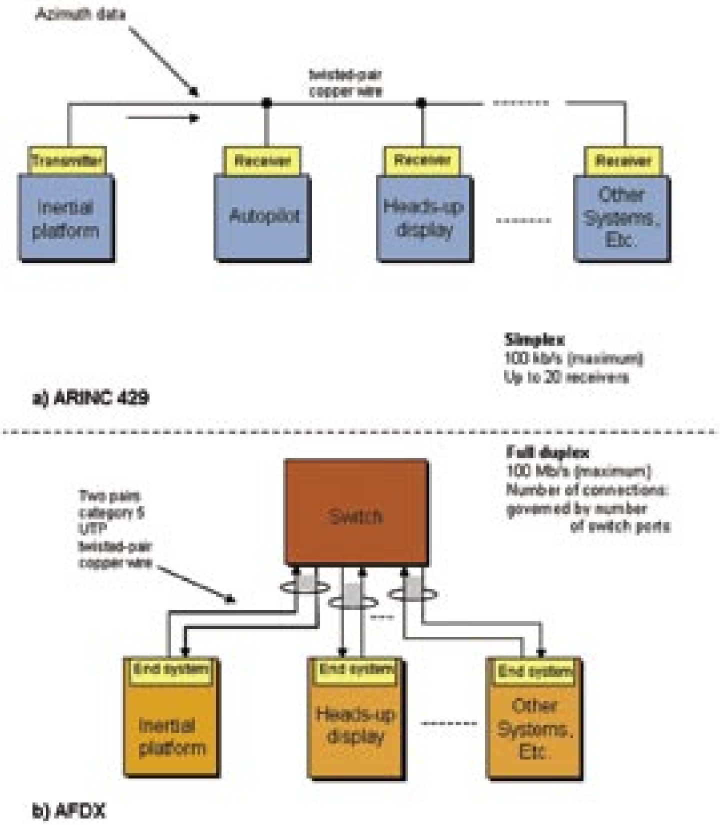

Figure 3—Unlike ARINC 429 (a) with AFDX (b) each subsystem requires only a single connection to the switch, which can have an arbitrarily large number of ports.

Condor Engineering and AFDX

Condor Engineering has been engaged in the support of avionics protocols for more than 16 years. In keeping with that tradition, Condor has recently introduced an innovative and technically advanced PMC (PCI mezzanine card) for AFDX protocol (figure 4). The Dual-Port ARINC 664 Card supports full throughput, simultaneously, on all channels. Its two independent, full-bandwidth ports can be employed for traffic monitoring, traffic generation, external triggering, and analysis. The card also offers high-resolution, time-tagging of incoming Ethernet packets, as well as highly accurate traffic generation. An IRIG-B (Inter-Range Instrumentation Group, Standard B) receiver/generator is included for synchronization to an external IRIG-B time source and for synchronization between multiple network interface cards.

Advanced AFDX traffic generation and end system libraries are included. The end system library implements the complete AFDX protocol stack and uses an XML (extensible markup language)-based configuration file approach for specifying an AFDX network.

For further details regarding the AFDX protocol, refer to Condor Engineering's AFDX Protocol Tutorial at www.condoreng.com/support/downloads/tutorials/index.shtml.

the 'bus drop' topology, and Figure 3b represents an Ethernet switched topology introduced in the discussion of ARINC 429 in the panel. In this example the twisted pair must link the transmitter of the inertial platform to the receiver of every device meant to receive the azimuth signal. With this point-to-multipoint property the avionics system must include an ARINC 429 bus for each communication path, and in a system with many end points, that can represent a major overhead. Also, because it requires some huge wiring harnesses, the ARINC 429 bus system adds undesirable weight.

With AFDX, as shown in Figure 3b, each subsystem is connected to the switch. So no matter how many subsystems require the azimuth signal from the inertial platform, none need to be connected individually to the inertial platform. Instead additional subsystems can be added by simply connecting just once to the switch.

Also in the case of ARINC 429, a transmitter can fan out to only 20 receivers. Whereas, with AFDX, the number of fan-outs from the inertial platform is limited only by the number of ports on the switch, which can be an arbitrarily large number. (This is connoted by the 'breaks' at the ends of the memory bus in figure 2, which denote that an arbitrarily large number of avionics subsystems can be added to the bus.

To learn more about AFDX, e-mail sales@condoreng.com or visit www.condoreng.com.

The History Behind AFDX

The evolution of AFDX and ARINC 664, Part 7, builds upon a noble heritage. In 1970 the University of Hawaii deployed a packet radio system, called the Aloha Network, to provide data communications between stations located on the islands comprising the Hawaiian chain. The communications system had no centralized control. Thus collisions (simultaneous transmission between two or more stations) could occur.

Ethernet

In 1972 Robert Metcalfe and David Boggs at Xerox Palo Alto Research Center built upon the Aloha Network idea. Using a coaxial cable as the communication medium, they invented Ethernet. As was the case with the ALOHA protocol, Ethernet also lacked centralized control, so again, transmissions from different stations or hosts could collide.

The Ethernet communication protocol is referred to as CSMA/CD (Carrier Sense, Multiple Access, and Collision Detection—see illustration below). Carrier sense means the hosts can detect whether the medium (coaxial cable) is idle or busy. Multiple access means that multiple hosts can be connected to the common medium. Collision detection means that, when a host transmits, it can detect whether its transmission has collided with the transmission of another host (or hosts). The original Ethernet data rate was 2.94 Mbs/s. The Ethernet local area network (LAN) is a half-duplex system in which collisions do occur.

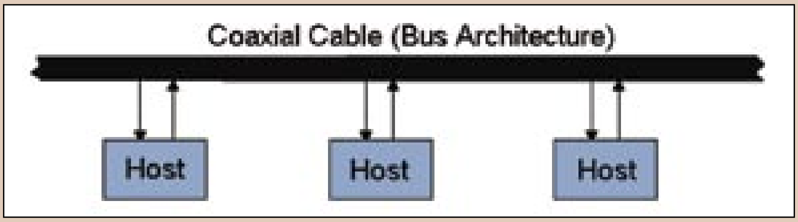

Figure 4—An Ethernet LAN is a half-duplex system in which collisions do occur.

ARINC 429

The ARINC 429 specification was introduced in 1978, and the current version, 429-15, was adopted in 1995 by the Airlines Electronic Engineering Committee (AEEC). The specification defines the hardware and the data formats required for bus transmission. The point-to-multipoint property of ARINC 429 requires that the avionic systems include an ARINC 429 bus for every pair-wise communication. The hardware comprises a single transmitter that can be connected to up to 20 receivers.

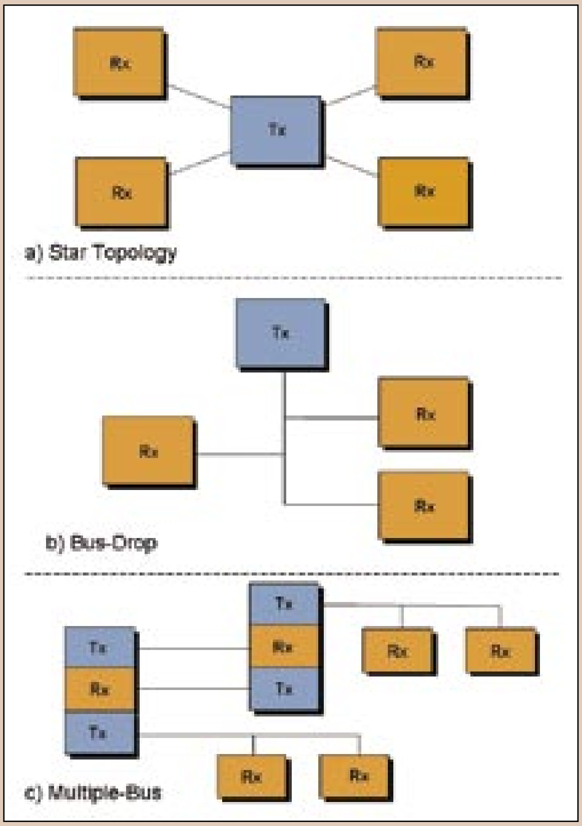

Figure 5—As shown here, all ARINC 429 topologies are single-direction (simplex) connections.

Data is transmitted in one direction only. This is also known as "simplex." The topology can take the form of a star, as shown in illustration A (above), or a bus-drop topology as shown in B. Bidirectional transmission requires two or more buses, as depicted in C. Transmission rates may be 12.5 or 100 kb/s. For interconnections the transmission buses employ 78-ohm shielded twisted pairs.

An ARINC 429 message consists of 32-bit words. A transmitter "talks" to any number of receivers, with each receiver monitoring continuously for data intended for it. But the receivers are unable to acknowledge receipt of data over the same interconnect.