

## **Protocols for Aerospace Control Systems A Comparison of AFDX, ARINC 429, CAN, and TTP**

A number of new and existing data buses are being proposed for use by aircraft manufacturers and system suppliers. A data bus provides numerous physical and logical configurations for avionics architecture, data units/packets, protocols, message traffic, and so forth. This allows considerable design flexibility for system or subsystem suppliers. Flexibility can make it extremely difficult to establish a type design, to determinate compliance, and to maintain continued airworthiness. This paper gives a comparative overview of data bus protocols targeted for aerospace controls that should be considered by implementers when developing, selecting, or integrating aerospace control systems in the context of aerospace projects.

Current developments and future trends are having enormous effects on the aerospace industry. While the life cycle of aircraft is increasing, the life cycle of their electronic components is decreasing. As more and more electric applications are finding their way into airplanes, the complexity of the overall system is growing. But all-electric aircraft require higher levels of safety, reliability and availability. Additionally, the cost pressure for development and operation is constantly rising. This calls for reusability of software and hardware components and its certification basis across several aircraft projects – both commercial and military.

TTTech provides a well-proven solution to these challenges. The Time-Triggered Protocol (TTP) represents a communication platform for true modular avionics/controls (distributed or backplane). This fault-tolerant deterministic data bus is supported by an integrated tool chain and dedicated communication controllers. A network designed according to the principles of the time-triggered technology guarantees new levels of safety at lower overall cost. The features of TTP translate directly into benefits such as reduced wiring and weight, lower total life cycle costs, shorter time-to-market, ease of system integration and system upgradeability, as well as increased reliability of the overall system.

Aircraft manufacturers desire to take advantage of the advances in computing technology to reduce aircraft weight and the time for design, development, integration, assembly, and certification, thus reducing the total life cycle costs, while also increasing airborne system performance. This motivates them to consider replacing point-to-point wiring and unidirectional data buses with faster and weight reducing bidirectional data buses. Although ARINC 429 is the most widely used aviation data bus today, it seems no longer adequate for the majority of future airborne system applications. Therefore, several new data buses are being considered for use in next-generation aircraft. Currently, AFDX is primarily being used as a high-level system bus on commercial transport aircraft, i.e. Airbus A380, and announced for others, while TTP and CAN are usually employed on the subsystem level in commercial transport or e.g. TTP as system bus on general aviation aircraft.

The multitude of embedded controllers and communications equipment found on aircraft need to communicate with each other to collect data, pass commands, and health monitoring. Currently, a move from largely independent, federated subsystems to more integrated system designs can be observed throughout the industry. The increased amount of information shared on the data bus within and between systems and subsystems, with many of the shared parameters being vital for

---

## Protocols for Aerospace Controls

the operation of the aircraft, asks for airborne networks with high throughput, availability, and reliability. The throughput, however, will always come second to the need for reliability.

### Data Bus Systems

On aircraft in general, there are unidirectional and bidirectional data bus architectures. The accompanying communications protocols will then support one or the other system. The simpler and to this day the most prevalent architecture on-board commercial airplanes is the unidirectional data bus represented by ARINC 429. The system consists of a single transmitter and a number of receivers, each monitoring the line and listening for messages relevant for them. Communication back to the transmitter, if needed, is performed by a separate transmitter, receiver(s), and wires. Due to its simplicity, the system is robust, fault-tolerant, and extremely simple to design and implement. The disadvantage is a lot of labor-intensive wiring, which is also expensive and heavy, requires higher power consumption and leads to high operational costs of the aircraft.

More advanced data bus architectures such as ARINC 629 are bidirectional. ARINC 629 is applied on Boeing 777 and some Boeing 737 models. Here, any member or user of the system can transmit, receive, or both on a common, shared data bus. As the implementation of the physical layer is comparatively expensive, market penetration did not happen. More than 30 years ago, especially military avionic systems, which previously used discrete wiring, were upgraded to interface with the MIL-STD-1553B bidirectional data bus. This led to a significant weight saving, and the labor to run wiring harnesses could be reduced by thousands of hours. However, the bus control became more complex, and the engineering effort to integrate the system was no small task.

### Bus Traffic Control

Every non-100% time-triggered, Time Division Multiple Access (TDMA) based databus must be able to arbitrate data bus transmissions to ensure that only one transmitter is operational at a time and that the receivers are waiting for messages. And, unlike the unidirectional bus, a lot of thought has to go into the system design and integration. There are two approaches commonly used for traffic control of bidirectional buses: central control and distributed control.

The advantage of the central control or command/respond approach is that only one bus component has control of the bus traffic (master). All users are directed by the bus controller (master). If the data bus architecture has to change, only the bus controller (master) has to be modified to support the new configuration. But the most significant weakness of such architecture in an aircraft environment, where the bus controller or network failure could have catastrophic repercussions, is that the bus controller (master) represents a single point failure. When it fails, the entire system fails. Therefore, many systems based on MIL-STD-1553B are equipped with redundant bus controllers, only one having control of the network at a time.

In a fully distributed control system, all members of the network are in charge of their own access. Single fault tolerance can be achieved when appropriate fault handling mechanisms are implemented as well as true partitioning is in place (both within and between Line Replaceable Units, LRUs) to prevent fault propagation. So if any member fails or happens to operate erratically, the rest of the system will not be affected and will continue operating correctly. The weakness of the distributed control is exactly the opposite of the central control strength. It generates higher complexity in bus access sharing, and a change to the system configuration may require a change to every user of the system.

---

## Protocols for Aerospace Controls

### Distributed Control Principles

The bus controller in a centrally controlled system has a list of addresses of all bus users and sends a command to each one of them at a predetermined rate, giving each user the chance to access the bus and send data required by other users. The core of any distributed control is a media access control protocol to ensure that only one transmitter speaks at a time. In avionics networks, three fundamental approaches to controlling bus access are used: contention resolution (CAN, Ethernet, AFDX), time slot allocation (ARINC 629, ARINC 659/SAFEbus, TTP), and token passing (FDDI, AS 4074, LTPB). Some protocols use mixed approaches.

A contention protocol like Carrier Sense Multiple Access with Collision Detection (CSMA/CD) for Ethernet allows any bus user to start transmitting at any time, provided the bus is idle at the moment. But if two users start transmitting at the same time, a collision occurs. This is not an efficient use of the bandwidth and the throughput suffers during heavy traffic. The time slot allocation, on the other hand, assigns a specific time slot for every user. The user listens to the bus for inactivity and if the bus is inactive during the user's assigned slot, the user may begin to transmit. If the bus is busy, access is not attempted until the next time slot. Some protocols may decide to proceed quickly to the next time slot, if the owner of the current time slot did not start to transmit anything. In a token ring protocol, the user is allowed to transmit only after the node has received a unique bit pattern called token.

Some CSMA/CD protocols like Ethernet need collision detection to determine whether or not their attempt to transmit was successful. The users monitor their own transmissions, comparing them to the transmitted data internally stored. If they detect data corruption, they stop transmitting. Then, when the users detect that the bus is idle again, they wait for a random time and, provided the bus is still idle, attempt to access it again. The reason for the random delay is to avoid permanent collisions.

Another problem that must be addressed is propagation delays on the transmission line. One station/user waiting to transmit may see an idle bus while another user is already transmitting. Random time delays alleviate this problem, but the result is poor bus utilization and severe throughput problems under heavy traffic.

The strive for better throughput of avionic networks resulted in sophisticated time allocation methods and increase of data bandwidth. Generally, the bus control methods are so critical that they are made transparent to the designer of a user terminal by providing him with a dedicated chipset and couplings to build the interface. Token ring protocols are not commonly used in avionic systems but do exist.


### ARINC 429

ARINC 429 was developed and is maintained by the Airlines Electronic Engineering Committee (AEEC) comprising members that represent airlines, government, and Aeronautical Radio, Incorporated (ARINC). The latter is a major company that develops and operates systems and services to ensure the efficiency, operation, and performance of the aviation and travel industries. It was organized in 1929 by four major airlines to provide a single licensee and coordinator of radio communications outside the government. Only airlines and aviation-related companies can be shareholders, although all airlines and aircraft can use ARINC's services.

## Protocols for Aerospace Controls

Page 4

ARINC 429 is a specification that defines, among many other things, how avionics equipment and systems should communicate with each other. They are interconnected by twisted pair wiring. The specification defines the electrical and data characteristics and the protocols that are used. ARINC 429 employs a unidirectional data bus standard known as Mark 33 Digital Information Transfer System (DITS) and can be used in an unidirectional multidrop system topology (see Fig. 1). Messages are transmitted at a bit rate of either 12.5 or 100 kbit/s to other system elements, which are monitoring the bus messages. Transmission and reception are on separate ports so that many wires are needed on aircraft equipped with a large number of avionics systems.



**Fig. 1:** ARINC 429 in unidirectional multidrop system topology

ARINC 429 has been installed on most commercial transport aircraft of leading manufacturers including Boeing and Airbus. A newer system specified as ARINC 629 was installed on the Boeing 777 for flight controls, and some aircraft use alternate systems in an attempt to reduce the weight of wire needed and to exchange data at a higher rate than is possible with ARINC 429. The unidirectional ARINC 429 system provides high reliability at the cost of wire weight and limited data rates.

### Key Attributes

- Unidirectional data bus standard MARK 33 Digital Information Transfer System (DITS)
- Bit rates: high speed 100 kbit/s, low speed 12.5 – 14.5 kbit/s
- Encoding: return to zero bipolar tri-state modulation
- Message length: 32-bit word, 255 word data block in block transfer mode
- Classes of service: periodic, sporadic and file transfer
- Media access: simplex single source multiple sink plus full duplex RTS/CTS handshake
- Topology: single source multiple sink
- Media: 78 Ohm unbalanced shielded twisted pair copper cable
- Number of nodes: 20 sinks, 1 source


---

## Protocols for Aerospace Controls

### CAN

Controller Area Network (CAN) is a serial data communications bus developed in the mid-eighties by Robert Bosch GmbH for the German car industry. It has since been adopted worldwide for automobile data communications, and in a wide range of other applications. The principal needs which drove the development of CAN were similar to those which spawned avionic data buses, i.e. to provide real-time data communication, with reduced cabling size and weight, and standardized I/O specification.

CAN belongs to the class of event-triggered protocols where the temporal control signals are derived primarily from events occurring outside or inside the computer system. Messages are sent if the host computer requests the transmission of a message, the channel is idle and the message priority wins over the messages that other nodes intend to send at the same time. CAN is deployed in a bus topology (see Fig. 2).



**Fig. 2:** CAN topology

There are two versions of CAN: CAN 1.0 and CAN 2.0. Most CAN controllers now in production are CAN 2.0, which is divided into two parts, known as A and B. Part A has a standard 11-bit identifier field, no specification for message filtering, and layered architecture based on Bosch's internal model. Part B extends the identifier field to 29 bits, has some message filtering, and supports the layer description of the Open Systems International (OSI) model. CAN 2.0 has been formalized in two ISO standards, ISO-11898 for high-speed applications (125 kbit/s to 1 Mbit/s) and ISO-11519 Part 2 for low-speed applications (125 kbit/s).

CAN is subdivided into different layers: the object layer, the transfer layer, and the physical layer. The CAN specification focuses on the first and second layers of the OSI model (data link layer and physical layer). A variety of protocols are used between the CAN data link layer and the various application layers. CAN implements five error detection mechanisms: Cyclic redundancy checks (CRC), frame checks, acknowledgment error checks, bit monitoring, and bit stuffing.

In the aerospace industry many different flavors of CAN are being used in order to compensate for the lack of determinism and fault tolerance.

## Protocols for Aerospace Controls

Page 6

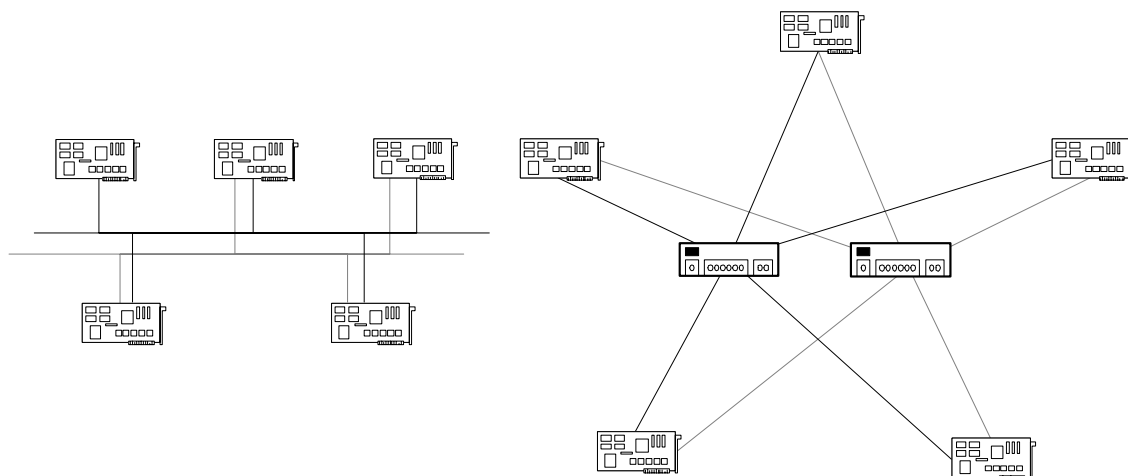
### Key Attributes

- Multimaster priority-based serial communications protocol supporting distributed real-time control and multiplexing using non-destructive contention-based arbitration
- Bit rates: 1 Mbit/s (with 40m bus), 100 kbit/s (with 500m bus)
- Encoding: non return to zero (NRZ)
- Message length: 0 to 8 bytes
- Classes of service: periodic and sporadic
- Media access: carrier sense multiple access with collision avoidance (CSMA/CA)
- Topology: terminated differential two wire bus
- Media: screened or unscreened twisted pair or flat pair telephone cable

### TTP

Time-Triggered Protocol (TTP) has initially been developed more than 20 years ago at the Vienna University of Technology and is now being maintained by the international cross-industry consortium TTA-Group since 2001. The protocol is at the heart of the Time-Triggered Architecture (TTA), which is a distributed computer platform for highly dependable real-time systems in the automotive, aerospace and other transport industry applications. A TTA system has effective consistency services and error detection mechanisms implemented to provide a maximum degree of fault tolerance, safety, and availability. Fault tolerance is dependent on the network topology used.

TTA can be implemented in a bus or star topology with controllers as bus interface units (see Fig. 3). The bus topology uses a broadcast bus where all nodes are electrically connected to each other. The star topology has all nodes connected to each of the replicated channels of the interconnection network via bidirectional links. In the bus topology each node is equipped with a local bus guardian; in the star topology two redundant central bus guardians (CBG) are implemented. A central bus guardian is a unit that acts as a failure mode converter to protect the communication channels from temporal transmission errors.



**Fig. 3:** TTP in star and bus topology

---

## Protocols for Aerospace Controls

TTP belongs to the class of the time-triggered protocols, where the temporal control signals are solely derived from the progression of time. TTP was specifically designed for high-dependability hard real-time applications, where timely error detection, fault isolation, and fault tolerance must be provided. Media access is controlled by a conflict-free Time Division Multiple Access (TDMA) strategy. Every node is assigned to a sending slot in a TDMA round. Additional slots can be defined for nodes to be added later. Every TTP controller contains the dispatching table (MEssage Descriptor List, MEDL) that contains the information which node is allowed to send which message at a particular point in time. The MEDLs of a cluster are constructed before run time and are common knowledge to all nodes. Every TTP controller contains two replicated communication channels in order that the loss of one channel can be tolerated, but can also operate only on one channel.

With TTP both event-triggered and time-triggered messages can be sent. The transmission of event-triggered messages is performed over an event channel. Bandwidth is reserved for event transmissions inside the TDMA slots, and the messages use identifiers. A typical event channel mechanism is a CAN emulation, in which a CAN-compatible interface is provided and the CAN messages are transmitted inside TTP data frames. Modularity of the TTP system is maintained when using event transmission in this way because bandwidth is not arbitrated among different nodes but only among different functions within a node. Timing and bandwidth analysis for event transmissions is therefore done on a per-node basis and does not need system-level design.

With its deployment in a variety of aerospace applications, TTP is a key contender for next-generation aircraft. Honeywell uses TTP for General Electric's F110 full authority digital engine control (FADEC) system on the Lockheed Martin F-16 fighter aircraft and for Honeywell's F124 full authority digital engine control system on the Aermacchi M-346 fighter trainer aircraft. Honeywell's APEX integrated cockpit with Nord-Micro has selected TTP as communication protocol for the Airbus A380 cabin pressure control system. Alcatel uses TTP as field bus protocol for their railway signaling system ELEKTRA 2, which will be used in Switzerland, Austria and Hungary.

### Key Attributes

- Real-time communication protocol for interconnection of LRUs in distributed real-time fault-tolerant systems but also backplane (see FADEC)
- Bit rates: currently available TTP controllers support up to 25 Mbit/s synchronous and up to 5 Mbit/s asynchronous transmission. Upper range is not limited by TTP itself and can be extended with new implementations of the TTP communication controller.
- Encoding: modified frequency modulation (MFM) and Manchester (both up to 5 Mbit/s), and MII (up to 25 Mbit/s) with currently available TTP controllers
- Message length: up to 240 byte with 4-8 bit header and 24-bit CRC
- Media access: distributed Time Division Multiple Access
- Topology: dual channel linear bus and star topology or mixed
- Media: copper and optical fiber
- Number of nodes: up to 64 nodes

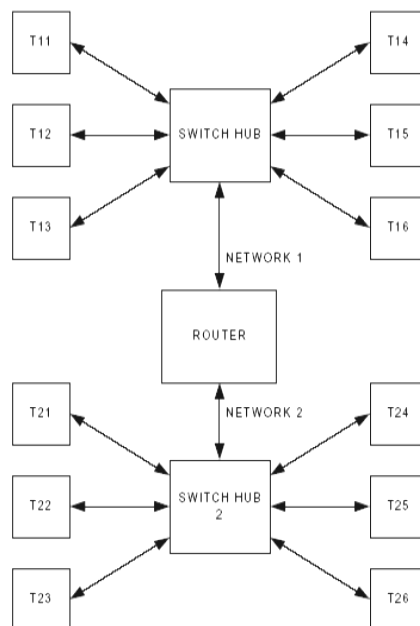
## Protocols for Aerospace Controls

Page 8

### AFDX (ARINC 664)

Avionics Full-Duplex Switched Ethernet (AFDX) is used as the main avionics data bus network on board the Airbus A380. Based on commercial 10/100 Mbit/s switched Ethernet, AFDX uses a special protocol for deterministic timing and redundancy management to provide secure and reliable communications of critical and non-critical data. The AFDX communication protocol has been derived from commercial data bus standards to achieve the deterministic behavior for avionics applications. AFDX protocol is specifically defined by ARINC 664 part 7 and is supported by other parts of ARINC 664. End systems (equipped with an AFDX controller) communicate based on virtual links by use of bandwidth allocation gaps. Switches (central AFDX switch) incorporate functions for filtering and policing, switching based on configuration tables, end-system and network monitoring.

Switched Ethernet is based on a star topology with full-duplex transmission, i.e. a node may send and receive at the same time (see Fig. 4). All nodes are connected to a switch that forms the center of the star. This switch distributes the messages in the system. If two or more nodes start to transmit at approximately the same time, their messages will be buffered in the switch and relayed sequentially. When the buffers are overrun and packages are dropped, timeliness can no longer be guaranteed.



**Fig. 4:** Two networks connected by a router in AFDX

AFDX uses several concepts to apply the switched Ethernet standard for the usage in safety-critical real-time applications. The two main additional concepts are: a traffic shaping policy and the replication of the communication channel. By executing the traffic shaping policy the timeliness property for message transmission can be addressed.

---

## Protocols for Aerospace Controls

Page 9

### Key Attributes

- Full-Duplex Switched Ethernet
- Bit rates: 10/100 Mbit/s cross data rate
- Encoding: PAM5x5
- Message length: 64 – 1518 bytes
- Classes of service: station to station, multicast and broadcast
- Topology: dual redundant star network
- Media: copper and fiber
- Number of nodes: up to 1024 (without bridges)

### Summary

This paper gives a comparative overview of four communication protocols used in avionics and aerospace controls systems. From an assessment of their key features it can be concluded that ARINC 429 is a well-defined, reliable, easy-to-implement, inexpensive protocol whose throughput and reliability have been adequate for most applications in the past. More sophisticated protocols are needed when large amounts of data have to be moved around quickly and reliably. CAN seems to be most appropriate for avionic systems that have less stringent timeliness and reliability requirements. If avionic systems require more throughput and reliability, then the TTP protocol appears to be the more suitable choice at subsystem level. Future airborne system applications will be most likely based on TTP and AFDX. These two data buses complement each other because they concentrate on different application areas, with TTP focusing on aircraft subsystems and general aviation and AFDX being optimized for higher level data transfer functions in large transport aircraft. TTP is ideally suited for low-speed (up to 25 Mbit/s), low-cost, fault-tolerant, robust applications and AFDX is best used for high-speed (up to 100 Mbit/s), fault-tolerant, robust applications.

## Protocols for Aerospace Controls

	<b>AFDX</b>	<b>ARINC 429</b>	<b>CAN</b>	<b>TTP</b>
<b>Max. frame length</b>	1518 bytes	32 bits	8 bytes	240 bytes
<b>Frame types</b>	normal	broadcast, file transfer	data, remote, error, overload	initialization, normal, x-frames
<b>Max. bit rate of current implementations</b>	10/100 Mbit/s	100 kbit/s	1 Mbit/s	25 Mbit/s
<b>Media access</b>	direct	direct	CSMA/CA	TDMA
<b>Max. bus length</b>	< 100 m	~ 65 m	40 m recommended	typically < 100 m (not limited by protocol)
<b>Latency</b>	depends on network load	very small; only controller-dependent	depends on message priority and network load	TTA usage: typically < 20 us
<b>Jitter</b>	depends on network load	very small; only controller-dependent	depends on message priority and network load	< precision (configurable, typically 5 us)
<b>Error containment</b>	yes; by independent switch device	no; single error can cause total loss of communication	no; single error can cause total loss of communication	yes; per node (communication network interface, bus guardian)
<b>Error handling strategy</b>	shut off erroneous nodes	ignored by receivers	immediate retry, passive state via error counters, shut off erroneous nodes (bus-off)	replicated channels, shut off erroneous nodes (fail-silent in bus, fail-operational in star)
<b>Components</b>	available, incl. TCP/IP protocol implementations, expensive	available, relatively inexpensive;	available, inexpensive	available, low-cost
<b>Voltage mode</b>	shielded twisted pair or optical fibers	twisted pair	twisted pair	shielded twisted pair, optical fiber or 100Base T/F (Ethernet-like)
<b>Price per chip</b>	250 – 700 \$	30 – 200 \$	1 – 2 \$	15 – 40 \$
<b>Net data rate</b>	–	53 kbit/s	< 300 kbit/s for 500 kbit/s network	< 4000 kbit/s (dual) for twisted pair < 20000 kbit/s (dual) for 100BaseT/F
<b>Cost per switch</b>	~ 15,000 \$	–	–	~ 75 \$