 <i>An SAE International Group</i>	AEROSPACE STANDARD	SAE AS5370	REV. A
		Issued 2002-04 Stabilized 2011-03	
		Superseding AS5370	
Multi-Transmitter Bidirectional Fiber-Optic Data Bus for Distributed Aircraft Control Systems			

RATIONALE

A review of AS5370 by AS-3A has determined that there are no current applications for this document. The subcommittee has recommended that the document status be changed to stabilized.

STABILIZED NOTICE

This document has been declared "Stabilized" by the (Technical Committee) and will no longer be subjected to periodic reviews for currency. Users are responsible for verifying references and continued suitability of technical requirements. Newer technology may exist.

SAENORM.COM : Click to view the full PDF of as5370a

SAE Technical Standards Board Rules provide that: "This report is published by SAE to advance the state of technical and engineering sciences. The use of this report is entirely voluntary, and its applicability and suitability for any particular use, including any patent infringement arising therefrom, is the sole responsibility of the user."

SAE reviews each technical report at least every five years at which time it may be reaffirmed, revised, or cancelled. SAE invites your written comments and suggestions.

Copyright © 2011 SAE International

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of SAE.

TO PLACE A DOCUMENT ORDER: Tel: 877-606-7323 (inside USA and Canada)
Tel: +1 724-776-4970 (outside USA)
Fax: 724-776-0790
Email: CustomerService@sae.org
http://www.sae.org

SAE WEB ADDRESS:

**SAE values your input. To provide feedback
on this Technical Report, please visit**
<http://www.sae.org/technical/standards/AS5370A>

FOREWORD

This document, in conjunction with EIA-709.1-A Control Network Protocol Specification [1], defines a complete 7-layer protocol stack for communications on a single-fiber (half-duplex) fiber-optic channel. This channel supports communication rates of 1250 kbps, 2500 kbps, or 5000 kbps between multiple nodes, each of which consists of a fiber-optic transceiver, a protocol processor, an application processor, a power supply, and application electronics.

The EIA-709.1-A Control Network Protocol was developed as a communications protocol for networked control systems. EIA-709.1-A provides many degrees of flexibility. It supports both peer-to-peer and master-slave control strategies. It provides a variety of message delivery service options and multiple addressing options. In addition, it is capable of supporting a variety of communications media, e.g., twisted-pair, radio, AC power lines, and fiber. Hence, the protocol is suitable for a broad scope of applications, e.g., from small home automation applications to large building automation systems or distributed industrial control systems.

EIA-709 provides a rich repertoire of communications services and capabilities using a variety of media. The protocol provides three basic logical addressing modes (broadcast, group addressing, and subnet-node addressing) and a physical addressing mode (using node-unique IDs) that is normally only used for downloading code and performing other network maintenance functions. The communications packet layout illustrated below shows how message data is encapsulated as a serial bit stream on the bus. The packet header is depicted by the format, `src_addr`, `dest_addr`, and `domain` fields. The format field is defined to handle variable lengths of address fields (`src_addr`, `dest_addr`, and `domain`).

SAENORM.COM : Click to view the full PDF of as5370a




FIGURE 1 - Message Packet Fields

EIA-709 uses a modified carrier-sense-multiple-access technique (predictive CSMA) to gain access to the medium for transmission of messages. Hence the time it takes to send a message from one node to another varies substantially based on network loading. Message latency is not predictable and not bounded.

When an EIA-709 transceiver wants to transmit a message the following steps are taken:

- Check to see if the communications channel is idle.
- If the channel is idle, wait for a random time interval in the range of 1 channel propagation delay time to 32 channel propagation delay times x estimated message backlog of 1 to 255 messages.
- If the channel is still active, the message is sent. If not idle, go back to step a.

AS5370 defines a modification of the lower levels of EIA-709 that permits bounded message latency times for those applications that require more timing predictability than provided by EIA-709.

Message delivery with bounded latency is made possible by the creation of fixed repetitive beacon periods as shown below for a typical application where each of the time slots (T0 to T7) must be wide enough to accommodate the maximum length message transmission. One (and only one) EIA-709 priority message packet can be transmitted in one of the time slots T0 through T7 -- provided that a node "owns" that specific timeslot. Any number of EIA-709 non-priority message packets may be initiated in the time interval labeled Topen. No messages will be initiated in the interval Tguard band. There will be no collisions (lost messages) for priority messages on an AS5370 network provided that:

- All nodes on the network use the same timeline (definition of intervals T0..T7, Topen, and beacon period),
- No more than one node has "ownership" of a priority time slot, and
- One node (beacon master) has been commanded to generate the beacon.

NOTE: Theoretically nodes can have their time slots assigned in a different way as long as the time periods each node has access to the media are mutually exclusive. This is much easier to check if all nodes use the same common time slot values.




FIGURE 2 - Beacon Timing Elements

SAENORM.COM : Click to view the full PDF of as5370a

TABLE OF CONTENTS

1. SCOPE	6
2. REFERENCES	6
2.1 Applicable Documents	6
2.2 Related Documents	6
2.3 Definitions, Symbols, and Terminology	6
2.4 Abbreviations	8
3. OVERVIEW OF PROTOCOL LAYERS.....	8
4. PHYSICAL LAYER.....	11
4.1 Optical Interconnect	11
4.2 Optical Requirements	11
4.3 Transceiver Requirements	12
4.3.1 Photo-Optic Emitter/Detector	12
4.3.2 Transmit and Receive Control.....	14
4.3.3 Protocol Processor Interface.....	16
4.3.4 Communications Parameters.....	18
5. MAC SUBLAYER	18
5.1 Service Provided	18
5.2 Interface to the Physical Layer.....	19
5.3 Interface to the Link Layer	19
5.4 MPDU Format	20
5.5 AS5370 Media Access Algorithm.....	20
5.6 Predictive p-persistent CSMA - Overview Description.....	22
5.7 Idle Channel Detection.....	23
5.8 Randomizing	24
5.9 Backlog Estimation	24
5.10 Message Priority	25
5.11 Beta1, Beta2 and Preamble Timing Parameters.....	26
6. UPPER PROTOCOL LAYERS	27
7. NOTES.....	27
7.1 Key Words	27
7.2 Addressing Hierarchy.....	27

TABLE OF CONTENTS (Continued)

FIGURE 1	Message Packet Fields	2
FIGURE 2	Beacon Timing Elements	3
FIGURE 3	Network Topologies & Symbols.....	8
FIGURE 4	Transceiver Functions	12
FIGURE 5	Physical Layer Protocol Data Unit.....	13
FIGURE 6	Frame Format for AS5370 Compliant Transmitter	13
FIGURE 7	Sample Latency-Bound Time Line	15
FIGURE 8	Beacon Signal Pattern	15
FIGURE 9	Interface of MAC Sub-Layer with Physical and Link Layers	17
FIGURE 10	MPDU/LPDU Format.....	20
FIGURE 11	AS5370 Media Access Algorithm	21
FIGURE 12	AS5370 Beacon Period Timing Example	22
FIGURE 13	Predictive p-persistent CSMA Concepts and Parameters.....	23
FIGURE 14	Allocation of Priority Slots Within the Busy Channel Packet Cycle	25
TABLE 1	Protocol Layerings	9
TABLE 2	AS5370 Fiber-Optic Cable Types.....	11
TABLE 3	Fiber-Optic Cable Specifications.....	11
TABLE 4	Transceiver Sync States.....	17
TABLE 5	Communication Parameters for AS5370 Interoperable Transceivers	18

SAENORM.COM : Click to view the full PDF of as5370a

1. SCOPE:

This specification applies to a communication protocol for networked control systems. The protocol provides peer-to-peer communication for networked control and is suitable for implementing both peer-to-peer and master-slave control strategies. This specification describes services for all seven protocol layers. In the layer 7 specification, it includes a description of the types of messages used by applications to exchange application and network management data.

2. REFERENCES:

2.1 Applicable Documents:

Control Network Protocol Specification EIA-709.1-A – available from Global Engineering Documents, 16 Inverness Way East, Englewood, CO 80112-5704 (or call 1-800-854-7179).

2.2 Related Documents:

2.3 Definitions, Symbols, and Terminology:

Use of Terms, Definitions, Symbols and Graphical Representations Abbreviations.

The following section introduces the basic terminology used in this document and/or EIA-709.1-A. Terminology that differs from what is commonly used will be defined. For example, in general, bridges do selective forwarding based on the layer 2 destination address. There are no layer 2 addresses in this protocol, so bridges forward all packets, as long as the Domain address in a packet matches a domain of which the bridge is a member. Routers, for many other protocols, perform network address modification so that two protocols with the same transport layer but different network layers can be connected to form a single logical network. AS5370/EIA 709 routers may perform network address modification, but typically they only examine the network address fields and selectively forward packets based on the network layer address fields.

Channel: A physical unit of bandwidth linking one or more communication nodes. Refer to Annex E of EIA-709.1-A for further explanation of the relationship between a channel and a subnet.

Bridge: Device that connects two channels (x and y) that forwards all packets from x to y and vice versa, as long as the packets originate on one of the domain(s) to which the bridge belongs.

Domain: A virtual network that is the network unit of management and administration. Group and subnet (see below) addresses are assigned by the administrator responsible for the domain, and they have meaning only in the context of that domain.

Subnet: A set of nodes accessible through the same link layer protocol; a routing abstraction for a channel; EIA-709.1-A subnets are limited to a maximum of 127 nodes.

2.3 (Continued):

Node: An abstraction for a physical node that represents the highest degree of address resolvability on a network. A node is identified (addressed) within a subnet by its (logical) node identifier. A physical node may belong to more than one subnet; when it does, it is assigned one (logical) node number for each subnet to which it belongs. A physical node may belong to at most two subnets; these subnets must be in different domains. A node may also be identified (absolutely) within a network by its Unique_Node_ID.

Group: A uniquely identifiable set of nodes within a domain. Within this set, individual members are identified by their member number. Groups facilitate one-to-many communication and are intended to support functional addressing.

Router: Device that routes data packets to their respective destinations by selectively forwarding from subnet to subnet; a router always connects two (sets of) subnets; routers may modify network layer address fields.

Network Variable: A variable in an application program whose value is automatically propagated over the network whenever a new value is assigned to it.

Service Pin Message: A network management message containing a node's Unique_Node_ID. Used by a network management device that receives this message to install and configure the node. May be generated by application or system code. May be triggered by external hardware event, e.g., driving a "service pin" input low.

Beta1: Period immediately following the end of a packet cycle. A node attempting to transmit monitors the state of the channel, and if it detects no transmission during the Beta1 period, it determines the channel to be idle.

Beta2: Randomizing slot. A node wishing to transmit generates a random delay T . This delay is an integer number of randomizing slots of duration Beta2.

Figure 3 shows the basic topology of networks based on this protocol and the symbolic representations used in this document. Fiber-optic connections are achieved with single fiber point-to-point links as shown in Figure 3a. Multiple nodes may be interconnected in a string provided that the transceiver in each node supports two links. Connecting the last node on the string to the first node results in a ring topology. Note that this is an active ring in that the transceivers at each node in the ring must be powered in order to maintain data propagation from one link to the next link. With an active ring topology each node in the ring receives and regenerates all incoming message packets. Regeneration reshapes the received signal and restores the power level of outgoing signal.




FIGURE 3 - Network Topologies & Symbols

2.4 Abbreviations:

The EIA-709.1-A standard defines data elements flowing between protocol layers as Protocol Data Units. The abbreviations used for these PDUs are:

PPDU	Physical Protocol Data Unit, or frame
MPDU	MAC Protocol Data Unit, or frame
LPDU	Link Protocol Data Unit, or frame
NPDU	Network Protocol Data Unit, or packet
TPDU	Transport Protocol Data Unit, or a message/ack
SPDU	Session Protocol Data Unit, or request/response
NMPDU	Network Management Protocol Data Unit
DPDU	Diagnostic Protocol Data Unit
APDU	Application Protocol Data Unit

3. OVERVIEW OF PROTOCOL LAYERS:

The protocol specified by AS5370 consists of the layers shown in Table 1 with the governing standard defined for each layer or sublayer. This section provides a summary description of each layer starting with layer1 and working up to layer 7.

Whereas EIA-709.1-A supports multiple Physical Layer protocols and data encoding methods, AS5370 addresses the fiber-optic physical media.

TABLE 1 - Protocol Layerings

LAYERS		Governing standard
6,7	Application & Presentation Layers Application: network variable exchange, application-specific RPC*, etc. Network Management: Network management RPC*, diagnostics	EIA-709.1-A
5	Session Layer Request-response	EIA-709.1-A
4	Transport Layer Acknowledged and unacknowledged unicast and multicast Authentication Server Transaction Control Sublayer Common ordering and duplicate detection	EIA-709.1-A
3	Network Layer Connection-less, domain-wide broadcast, no segmentation, loop-free topology, learning routers	EIA-709.1-A
2	Link Layer framing, data encoding, CRC error checking MAC Sublayer Optional predictive p-persistent CSMA: collision avoidance; optional priority and collision detection; Optional latency-bound media access: no collisions with priority, predictive p-persistent CSMA with no priority	EIA-709.1-A AS 5370
1	Physical Layer Fiber-optic media	AS 5370

* RPC means remote procedure call

3. (Continued):

The MAC (Medium Access Control) sublayer employs a collision avoidance algorithm called Predictive p-persistent CSMA (Carrier Sense, Multiple Access). AS5370 extends the media access options of EIA-709.1-A to include a latency-bound media access algorithm.

The Link Layer supports a simple connection-less service for simplicity and compatibility with the multicast protocol. Its functions are limited to framing, frame encoding, and error detection, with no error recovery by re-transmission.

The Network Layer handles packet delivery within a single domain, with no provisions for inter-domain communication. The Network service is connection-less, unacknowledged, and supports neither segmentation nor re-assembly of messages. The routing algorithms employed by the network layer to learn the topology assumes a tree-like network topology; routers with configured tables may operate on topologies with physical loops, as long as the communication paths are logically tree-like. In this configuration, a packet may never appear more than once at the router on the side on which the packet originated. The unicast routing algorithm uses learning for minimal overhead and no additional routing traffic. Use of configured routing tables is supported for both unicast and group addresses, although in many applications a simple flooding of group addressed messages is sufficient.

The Transport and Session layers encompass functionality that is important for reliable message delivery. A common Transaction Control Sublayer handles transaction ordering and duplicate detection for both. The Transport layer is connection-less and provides reliable message delivery to both single and multiple destinations. Authentication of the message sender's identity is included as a transport layer service, for use when the security of sender authentication is required. The authentication server requires only the Transaction Control Sublayer to accomplish its function. Thus Transport and Session layer messages may be authenticated using all of the addressing modes other than broadcast.

The Session Layer provides a simple Request-Response mechanism for access to remote servers. This mechanism provides a platform upon which application specific remote procedure calls can be built. The EIA-709.1-A network management protocol, for example, depends upon the Request-Response mechanism in the Session layer.

A transport layer acknowledged message expects indication of message delivery from remote destination(s). A session layer Request message expects indication that application-specific remote task(s) have been completed. A given message uses only one or the other type of service, but not both.

The EIA-709.1-A standard includes the Presentation Layer and the lowest level of the Application Layer. These layers provide services for sending and receiving application messages including network variables, and other types of messages such as network management and diagnostic messages and foreign frames (see section 11 of EIA-709.1-A). For a network variable update, the APDU header provides information on how the APDU is to be interpreted. This application independent interpretation of the data allows data to be shared among nodes without prior arrangement.

4. PHYSICAL LAYER:

This section defines the fiber-optic physical layer of the AS5370 protocol.

4.1 Optical Interconnect:

This section defines options for fiber-optic cables, connectors, and optical wavelengths. The cables listed in Table 2 shall be used for the network wiring.

Standard readily available commercial connectors that may be used include FSMA (threaded) and bayonet (quarter-turn quick disconnect, e.g., Lucent ST™) connector body types.

The standard optical wavelength is 850 nanometers $\pm$ 35 nanometers. An alternative wavelength is 1320 nanometers $\pm$ 35 nanometers, but all cable on a given network shall use the same wavelength.

TABLE 2 - AS5370 Fiber-Optic Cable Types

Fiber Type	Fiber (Core) Diameter	Fiber (Cladding) Diameter	Comments (informative) See Table 3
multi-mode, graded index	62.5 μ m	125 μ m	Acceptable if lower emitted power does not reduce link margins below 6 db.
multi-mode, graded index	100 μ m	140 μ m	Preferred when maximum coupled launch power is needed or desired, and has been widely adopted for aircraft use.

TABLE 3 - Fiber-Optic Cable Specifications

Optical Parameters	62.5/125 Fiber Cable	100/140 Fiber Cable
Attenuation at 880 nm	< 3.5 dB/km	< 5 dB/km
Attenuation at 1320 nm	< 1.5 dB/km	< 3 dB/km

4.2 Optical Requirements:

Single multi-mode fibres operated in half-duplex send/receive mode shall be used to interconnect each node to two other nodes on the fiber-optic ring at the selected optical wavelength. Optical power level transmitted shall not exceed -10 dBm, where 0 dBm equals 1 milliwatt. Receiver sensitivity shall be at least -26 dBm. Thus, conforming AS5370 transceivers can provide an optical link power budget of 16 dBm. Link power margins (after all link losses are accounted for) should be at least 6 db for reliable operation.

4.3 Transceiver Requirements:

A transceiver supporting dual fiber links is shown in Figure 4. A dual link transceiver shall have a fiber-optic transceiver that includes the following functions:

- A processor interface,
- Send/receive control logic to control transceiver operation and data flow, and
- Two half-duplex photo-optic emitter/detectors,

Note that the repeater function and one emitter/detector may be deleted if a node is required to support only one link.




FIGURE 4 - Transceiver Functions

- 4.3.1 Photo-Optic Emitter/Detector: The photo-optic emitter/detector translates received light into electrical signals or electrical signals into light based on control inputs from transmit/receive control logic. The emitter/detector function requires a single high performance electro-optic device operating in half duplex mode. This device shall be capable of being switched between light emitting (LED) and light detecting (photodiode) mode at the specified wavelength. Using this device eliminates the need for physically separate emitters and detectors and the use of an optical splitter and combiner. This wavelength shall be either 850 or 1320 nanometers. Figure 5 shows the physical layer protocol data unit (PPDU) which is used on the fiber-optic network.




FIGURE 5 - Physical Layer Protocol Data Unit

- 4.3.1.1 Transmit Mode: Consistent with the PPDU format shown in Figure 5, the emitter/detector in transmit mode shall generate a bit stream on the medium in accordance with Figure 6. 'T' is the bit period (for example 800 nanoseconds at 1250 Kbps), equal to $1/(\text{Bit Rate})$.




FIGURE 6 - Frame Format for AS5370 Compliant Transmitter

The emitter shall employ Differential Manchester encoding of data and clock information. This scheme provides a transition (referred to as “clock transition”) at the beginning of every bit period for the purpose of synchronizing the receiver clock. Zero/one data are indicated by the presence or absence of a second transition (the “data transition”) halfway between clock transitions.

Clock transitions occur at the beginning of a bit period, and therefore, the last valid bit in the packet does not have a trailing clock edge. Thus, polarity is arbitrary at the start of transmission.

The emitter shall transmit a preamble at the beginning of a packet to allow the other nodes to synchronize their receiver clocks. The preamble consists of (i) a bit-sync field and (ii) a byte-sync field. The bit-sync field is a series of Differential Manchester “1”s. The byte-sync field is a single-bit Differential Manchester “0” that marks the end of the preamble and the beginning of the first byte of the LPDU/MPDU delivered to the physical layer by the data link layer/media access sub-layer.

4.3.1.1 (Continued):

The emitter shall terminate the packet by forcing a Differential Manchester line code violation, i.e., it shall hold the data output transitionless long enough for the receiver(s) to recognize an invalid bit code (line code violation). To the receiver(s), this shall signal the end of the packet. Specifically, at the end of packet transmission, the transmitted waveform shall remain transitionless for at least three 130 microseconds after the final clock transition (excepting the final data transition, if present). The emitter shall ensure that no light is emitted during the idle periods between packets.

4.3.1.2 Receive Mode: Consistent with the PPDU format shown in Figure 5, the emitter/detector in receive mode shall accept a bit stream on the medium in accordance with Figure 6.

The receiver shall detect edge transitions during two windows for each bit period, T . The first window is at $T/2$, and determines if a "0" is being received. The second window is at T , and defines a "1". The transition during the second window sets up the next two windows at $T/2$ and T . The receiver shall receive data as long as transitions continue. Timing instability, or jitter, may be caused by changes in the communication medium, or instability in the input clocks of the transmitting or receiving nodes. The value of $T/2$ shall be 400 nanoseconds $\pm$ 80 nanoseconds and T shall be 800 nanoseconds $\pm$ 80 nanoseconds.

For the receiving node to reliably terminate reception of a packet, the received line-code violation period must have no transitions until the receiving node detects the end of the packet. The receiving node shall terminate a packet if no clock transitions are detected after the last bit.

The receiver shall ensure reception of data by the node under all conditions specified in this document.

4.3.2 Transmit and Receive Control: The transmit and receive control logic shall control the flow of message data between each Emitter/Detector and the Processor Interface by controlling whether the Emitter/Detector is in receive mode or transmit mode. Note that the emitter/detector will be in the detect mode until there is a need to transmit a message or beacon signal.

The transmit and receive control logic controls the repeating of data received on one fiber port to the other fiber port when the network topology is a ring. The repeater regenerates optical signals from the optical port receiving data to the other optical port.

When the transceiver is operating in latency-bound mode it provides additional functions. Specifically, the transmit control shall be capable of generating the beacon signal when configured as a beacon master and the receive control must be capable of detecting receipt of a beacon signal when the transceiver is not a beacon master. A typical beacon period (the latency-bound time line) is illustrated in Figure 7. The beacon signal is a sequence of pulses that will be seen by all receivers as Manchester code violations (hence not interpreted as data bits). The timing intervals t_1 , t_2 , and t_3 in Figure 9 must each be 2.5, 0.5, and 3.0 data bit times respectively.




FIGURE 7 - Sample Latency-Bound Time Line




FIGURE 8 - Beacon Signal Pattern

4.3.2 (Continued):

Latency-bound operation requires definition of the following parameters:

- a. Priority time slots must be wide enough to accommodate the transmission of a complete message including preamble, header, address fields, message data, 16-bit CRC, and the end of message code. Note that the width of priority slots for the CSMA mode are only wide enough to ensure that all other nodes on the network can sense that the node with priority access has exercised his option to commence sending a priority message.
- b. The non-priority time slot (Topen interval) must also be wide enough to accommodate the transmission of a complete non-priority message. The width of the non-priority slot can be minimized if a particular system transmits all data with priority slots. It is common practice to set the non-priority time interval to 1 or 2 milliseconds in order to support test equipment and development systems that use non-priority messaging.
- c. The guard band interval should be of sufficient duration to allow completion of a non-priority message without its being corrupted by the beacon signal.
- d. The beacon period is equal to the sum total of all the priority time slots, the non-priority time slot, and the guard band interval.

- 4.3.3 Processor Interface: The physical layer handles the actual transmission and reception of binary data. The physical layer shall interface with the MAC sub-layer and the Application layer as defined in this section. The bit error rate presented to the MAC layer is expected be equal to or better than 1 in 10^{-9} . For compatibility with the higher layers, the physical protocol must support the following defined service interfaces (see Figure 9):

P_Data_Indication (Frame) - Physical layer provides this indication to the MAC sublayer and the link layer once per incoming LPDU/MPDU.

P_Data_Request (Frame) - The MAC sublayer uses this primitive to pass the Frame, the encoded LPDU/MPDU, to the physical layer for immediate transmission. The bit transmission order is defined in Annex D.

P_Data_Confirm (Status) - The physical layer returns Status as to whether the frame was transmitted. Status has two possible values: success—indicating the frame was transmitted or request_denied—indicating that activity was detected on the line prior to transmission.

P_Channel_Active () - The physical layer uses this primitive to pass the status of the channel to the MAC sublayer. This is an indication of receiver activity, not necessarily of valid data.

P_Xcvr_Request () - The MAC sublayer uses this primitive to request transceiver status.

P_Xcvr_Status (xcvr_present, sync_state) - The physical layer uses this primitive to pass the status of the fibre_optic channel to the MAC sublayer. The definition of sync states is defined in Table 1.

The processor interface also provides the path by which a host application can configure the transceiver for proper operation in the latency-bound mode of operation. As a minimum, the following controls shall be provided by the host processor to the transceiver's transmit control:

- a. A Mode Control flag for a Boolean parameter, pCSMA or Latency-bound,
- b. A Beacon Master Enable register for a Boolean parameter, enabled or disabled,
- c. An Application Offset Timer Value register (16-bit value with LSB equal to 16 bit times)
- d. A Beacon Watchdog Timer Value register (8-bit value with LSB equal to 8192 bit times)
- e. A Beacon Watchdog Enable flag for a Boolean parameter, enabled or disabled,
- f. Time value registers (15-bit value with LSB equal to 16 bit times) for at least 29 priority time slots, and
- g. Time slot owned flag (Boolean assigned to this node or not assigned) for at least 29 priority time slots.

As a minimum, the following status information shall be provided by the transceiver's transmit control to the host processor:

- a. An Application Offset Timer expired flag and
- b. A Beacon Watchdog Timer expired flag




FIGURE 9 - Interface of MAC Sub-Layer with Physical and Link Layers

TABLE 4 - Transceiver Sync States

State	Priority Xmit	Xmit	Comments
SYNC_CSMA	Yes	Yes	Transceiver in predictive CSMA mode (Fall back state)
SYNC_LOCK	No	No	Transceiver in latency-bound mode & waiting for priority slot
SYNC_SLOT	Yes	No	Transceiver in latency-bound mode & current slot allocated to this node
SYNC_OPEN	No	Yes	Transceiver in latency-bound mode & this is the non-priority slot (all nodes may attempt to transmit)

- 4.3.4 Communications Parameters: The following communication parameters (Table 5) apply to an AS5370 channel. Note that these parameters are not relevant to the latency-bound method of media access.

TABLE 5 - Communication Parameters for AS5370 Interoperable Transceivers

Communication Rate	1250 kbps	2500 kbps	5000 kbps
Minimum No. of Bit Sync Bits	7	7	7
Preamble Length	125 μ s	6.4 μ s	3.2 μ s
Beta2 Width	24.0 μ s*	24.0 μ s*	24.0 μ s*
Beta1 Width	24.6 μ s*	24.6 μ s*	24.6 μ s*

NOTE: The value for Beta2 and Beta1 depend on fiber length and the number of AS5370 nodes on the network. Add two bit-times per node for networks larger than 10 nodes. (1.6 μ s for 1250 kbps, 0.8 μ s for 2500 kbps, 0.4 μ s for 5000 kbps)

5. MAC SUBLAYER:

This section defines the Service Provided, Interface to the Link Layer, Interface to the Physical Layer, MPDU Format, AS5370 Media Access Algorithm, Predictive p-persistent CSMA, etc.

5.1 Service Provided:

The Media Access Control (MAC) sublayer facilitates media access with optional priority. It uses a protocol called Predictive p-persistent CSMA (Carrier Sense, Multiple Access), which has some resemblance to the p-persistent CSMA protocol family.

Predictive p-persistent CSMA is a collision avoidance technique that randomizes channel access using knowledge of the expected channel load. A node wishing to transmit always accesses the channel with a random delay in the range (0..w). To avoid throughput degradation under high load, the size of the randomizing window, w, is proportional to the estimated channel backlog BL:

$$w = BL * W_{base}$$

where W_{base} is the base window size. W_{base} is measured in time. Its duration, derived from Beta2 (see 5.7), equals 16 Beta2 time intervals.

5.2 Interface to the Physical Layer:

The Physical layer handles the actual transmission and reception of binary data. Multiple physical layer protocols are supported by the EIA-709.1-A control network protocol. For compatibility with the higher layers, all physical protocols must support the defined service interface (see). These interface definitions are described in 4.3.3.

5.3 Interface to the Link Layer:

The MAC sublayer is closely coupled to the Link layer, described in Section 6 of EIA-709.1-A. With the MAC sublayer being responsible for media access, the Link layer deals with all the other layer 2 issues, including framing and error detection. For explanatory purposes, the interface between the two layers is described in the form shown in Figure 9.

Although the service interface primitives are defined using a syntax similar to programming language procedure calls, no implementation technique is implied. Frame reception is handled entirely by the Link layer, which notifies the MAC sublayer about the backlog increment via the Frame_OK() primitive.

The interface between the Link and the MAC layers has the following service interface primitives:

M_Data_Request (Priority, delta_BL, ALT_Path, LPDU)

This primitive is used by the Link layer to pass an outbound LPDU/MPDU to the MAC sublayer. Priority defines the priority with which the frame is to be transmitted; delta_BL is the backlog increment expected as a result of delivering this MPDU. ALT_Path is a binary flag indicating whether the LPDU is to be transmitted on the primary or alternate channel, baud rate, etc. See 5.4 for how ALT_Path is set.

Frame_OK (delta_BL)

On receiving a frame and verifying that its CRC is correct, the Link layer invokes this primitive to notify the MAC sublayer about the backlog increment associated with the frame just received.

M_Data_Indication()

The MAC sublayer provides this indication to the link layer once per incoming LPDU/MPDU.

5.4 MPDU Format:

The combined MPDU/LPDU format is shown in Figure 10. (Annex D of EIA-709.1-A contains the details of the NPDU frame).

The MAC sublayer uses the L2Hdr field, which has the following syntax and semantics:

- Pri
- Alt_Path
- Delta_BL
- 1-bit field specifying the priority of this MPDU: 0 = Normal, 1 = High
- a 1-bit field specifying the channel to use. This is a provision for transceivers that have the ability to transmit on two different channels and receive on either one without prior configuration. The transport layer sets this bit for the last two retries, unless requested to specify the alternate path for every transmission
- a 6-bit unsigned field that specifies channel backlog increment value to be added as a result of delivering this MPDU (value can be zero).




FIGURE 10 - MPDU/LPDU Format

5.5 AS5370 Media Access Algorithm:

This section defines the algorithm for either latency-bound media access or predictive p-persistent CSMA media access. Figure 12 illustrates the CSMA as well as the latency-bound media access algorithms for AS5370. Figure 14 provides an example of the beacon period timing for the AS5370 latency-bound mode of operation for a 4-node network.




FIGURE 11 - AS5370 Media Access Algorithm




FIGURE 12 - AS5370 Beacon Period Timing Example

5.6 Predictive p-persistent CSMA - Overview Description:

Like CSMA, Predictive p-persistent CSMA senses the medium before transmitting. A node attempting to transmit monitors the state of the channel (see Figure 10), and determines the channel to be idle if it detects no transmission during the Beta1 period. Nodes without a packet to transmit during this Beta1 period shall remain in synchronization for the duration of the priority slots (see 5.10), and at least W_{base} randomizing slots. This maintenance of synchronization allows a packet that arrives in the output queue of the MAC sublayer after the end of the Beta1 time to be transmitted in a valid slot according to the other nodes with a packet to transmit.

Next, the node generates a random delay T (transmit) from the interval $(0..BL \cdot W_{base})$, where W_{base} is the size of the basic randomizing window and BL is an estimate of the current channel backlog. T (transmit) is defined as an integer number of randomizing slots of duration Beta2 (see 5.7 and 5.8). If the channel is idle when the delay expires, the node transmits; otherwise, the node receives the incoming packet, and then repeats the MAC algorithm. In Figure 8, D_{mean} is the average delay between packets, and, since the random delay T is uniformly distributed, D_{mean} is given as $W_{base}/2$ for small values of BL .




FIGURE 13 - Predictive p-persistent CSMA Concepts and Parameters

5.6 (Continued):

By adjusting the size of the randomizing window as a function of the predicted load, the algorithm keeps the collision rate constant and independent of the load. Provided that the estimated backlog is greater than or equal to the real backlog, the following holds:

$$\text{Collision Rate} = \text{Error Pkt Cycles} / \text{Error Free Pkt Cycles} \leq 1 / 2W_{\text{base}}$$

A base window size of 16 is used for the EIA-709.1-A Control Network Protocol. This implies that there are an average of 8 randomizing slots of width Beta2 and one slot of width Beta1 between each packet. Also, the width of the Beta2 period is crucial to efficient utilization of the channel.

5.7 Idle Channel Detection:

The idle channel condition is asserted whenever the following two conditions are met:

1. The current channel state reported by the physical layer via the P_Data_Indication () primitive is low; and
2. No transition has been detected during the last period of Beta1.

The length of the Beta1 period is defined by the following constraint:

$$\text{Beta1} > 1 \text{ bit time} + (2 * \text{Tau}_p + \text{Tau}_m)$$

5.7 (Continued):

The first term assumes a data encoding method that guarantees a transition and/or carrier during every bit time. If encoding methods are used which do not meet this constraint, then the first term must be adjusted to be the longest time that the channel may appear idle without being idle, i.e. the longest run in legal data transmission without a transition and/or carrier asserted on the medium. The second term takes care of propagation and turnaround delays, which are:

τ_p is the physical propagation delay defined by the media length;
 τ_m is the detection and turn-around delay within the MAC sublayer; this is the period from the time the idle channel condition is detected, to the point when the first output transition appears on the output. On media where there is a carrier, this time must include the time between turning on the carrier, and it being asserted as a valid carrier on the medium.

5.8 Randomizing:

At the beginning of the randomizing period, a node wishing to transmit generates a random delay T (transmit) from the interval $(0, BL * w_{base})$. The node then waits for this period, while continuing to monitor channel status; if the channel is still idle when the delay expires, the node transmits.

The transmit delay T (transmit) is an integer number of randomizing slots of duration β_2 ; the length of the randomizing slot must meet the following constraint:

$$\beta_2 > 2 * \tau_p + \tau_m$$

Parameters τ_p and τ_m are defined in 5.7.

5.9 Backlog Estimation:

The predictive aspect of the MAC algorithm is based on backlog estimation. Each node maintains an estimate of the current channel backlog BL , which is incremented as a result of sending or receiving an MPDU and decrements periodically—once every packet cycle. The increment to the backlog is encoded into the link layer header, and represents the number of messages that the packet shall cause to be generated upon reception. After sending or receiving a packet with a non-zero backlog increment, the node's backlog estimation is incremented by the backlog increment.

The backlog is decremented under one of the following conditions:

On waiting to transmit: If $BL * w_{base}$ randomizing slots go by without channel activity

On receive: If a packet is received with a backlog increment of '0'

On transmit: If a packet is transmitted with a backlog increment of '0'

On idle: If a packet cycle time expires without channel activity.

5.9 (Continued):

The backlog always has a value ≥ 1 . The algorithm post-increments rather than pre-increments the backlog by the amount associated with the MPDU being transmitted, because the number of expected responses is of no importance until after transmitting the MPDU.

5.10 Message Priority:

The EIA-709.1-A protocol supports optional priority on a channel by channel basis.

Priority slots, if any, occur immediately after the Beta1 period following the transmission of a packet (Figure 10). The number of priority slots per channel ranges from 0 to 127. Priority slots are typically not contended for, but rather are uniquely assigned to nodes on the channel. Nodes that have been assigned a priority slot do not have to use it with every message; the node decides on a message by message basis whether or not to use the assigned priority slot. This determination is made by examining the priority bit within the LPDU header (Figure 7).

For AS5370 in latency-bound mode, priority messages are sent within fixed time intervals (time slots) as described in 5.5.




FIGURE 14 - Allocation of Priority Slots within the Busy Channel Packet Cycle

The predictive p-persistent CSMA protocol provides no synchronization among the nodes. Therefore, if the channel has been idle for longer than the randomizing period ($\text{Beta1} + \text{number of priority slots} + D_{mean}$ above), access to the link is random without regard to priority. Once the link returns to the busy state, access to the link shall be in priority order.