

AEROSPACE INFORMATION REPORT

SAE AIR4289

REV. A

Issued 1995-02
Reaffirmed 2001-10
Stabilized 2012-05
Superseding AIR4289

Handbook for the SAE AS4075
High Speed Ring Bus Standard

RATIONALE

This document has been determined to contain basic and stable technology which is not dynamic in nature.

STABILIZED NOTICE

This document has been declared "Stabilized" by the SAE AS-1A Avionic Networks 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 air4289a

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 revised, reaffirmed, stabilized, or cancelled. SAE invites your written comments and suggestions.

Copyright © 2012 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
SAE WEB ADDRESS: <http://www.sae.org>

**SAE values your input. To provide feedback
on this Technical Report, please visit
<http://www.sae.org/technical/standards/AIR4289A>**

FOREWORD

This High Speed Ring Bus (HSRB) Handbook has been prepared by the SAE AS-2B Ring Implementation Task (RIT) Group. The main purpose of the handbook is

- a. To provide an explanation of the standard in terms less formal and more readily understood than those which appear in the standard.
- b. To document the body of knowledge which directed the formulation of the standard.
- c. To provide notes on the implementation and use of the standard.

Initial investigation of high speed data transmission by the SAE began in 1979 with a study of requirements and the technology available at that time. It was recognized that progressively higher data rates were being required by the aerospace community and that novel architectures were being considered, capable of achieving not only high processing performance but also reduced life cycle cost. Having established requirements in the form of the HART (High Speed Data Bus Applications and Requirements Task Group) document, two Task Groups were set up in October 1984 to develop ring and linear high speed data bus standards AS4075 and AS4074 respectively. The Ring Implementation Task Group was established as part of the SAE AE-9B Subcommittee, subsequently the SAE AS-2 Committee, with a mandate to develop a ring data communication standard based on a token passing access method. The objective was, at a minimum, to meet all of the requirements of the HART document. The resulting standard, SAE AS4075, has been approved by the SAE Aerospace Council and meets all of the original objectives in full.

The conclusions of the HART committee were summarized in the SAE AE-9B High Speed Bus Requirements document. Among the fundamental findings, the document dictated that a high-speed network utilize serial data, with a range of from 1 to 4096 16-bit words per message. A serial bus has obvious size, cost and weight advantages over a parallel bus structure. The wide range of data word counts allows for very high throughput efficiency, minimizing the effects of message overhead.

FOREWORD (Continued)

A second general requirement of the HART document calls for distributed control. A token passing protocol meets this requirement and was selected because it readily lends itself to the processor-to-processor data communication applications foreseen for future military platform applications. In contrast to the command/response protocol structure of MIL-STD-1553, token passing reduces the overhead time required to control media access and enhances system reliability. The reduced overhead time arises from the fact that a round-robin token pass on a physical ring is less time-consuming than is the corresponding polling sequence for a command/response system, particularly at higher data rates. This is principally due to the absence of the round trip propagation delay that is inherent in a command/response protocol, but which is eliminated in a round-robin token passing protocol. The improvement in system reliability results from the elimination of the single point failure mode associated with a centralized bus controller for command/response networks.

Two other general requirements of the HART document involve addressing modes and classes of latency. The two addressing modes specified are physical addressing and logical addressing. Physical addressing assigns a unique address to each station on a network and is used for normal station-to-station communication. Logical addressing entails a message destination address not associated with any unique station. It may be used to identify a particular message (e.g., navigation information) to all stations on a network. Logical addressing supports multicasting (transmitting to a preassigned subset of all stations on a network) as well as broadcast messages (transmitting to all other stations on a network).

The requirement for different latency classes demands a mechanism for prioritizing of messages. By embedding such a mechanism in the token-passing protocol, it is possible to allow messages of high urgency, such as those involving imminent safety, to be transmitted prior to other traffic on the network. Priorities also allow the scheduling of real-time traffic.

The HART document also lists a number of requirements relating to the topology of the high speed networks. These specifications include a minimum of 128 stations, 300 m station separation (1000 m desired), electrical isolation as well as guidelines for station coupling, insertion and removal. The document demands that the standard be independent of both data rate and media type, allowing for both wire and fiber implementations. Additional HART requirements related to reliability and "robustness" issues, such as the necessity for rapid "healing" time for recovery from station and/or media faults. Other HART requirements deal with issues such as predictability, self-test, fault tolerance, fault isolation, initialization time and the capability to bridge to other networks.

In addition to the general requirements, the HART document also outlines a number of performance requirements and design goals. Since the emerging standards were to be targeted primarily for military applications, there is good deal of emphasis on real-time performance criteria. The most significant of these benchmarks is information rate. The average rate of data bus transmission is evaluated as a function of several parameters including the number of active stations, variable from 1 to 128, the use of broadcast and/or multicast messages, the effect of message traffic with varying latency levels and the effects of message overhead, short and long messages and the number of active and nonactive stations.

FOREWORD (Continued)

A second performance criterion is word error rate. This was specified as a 90% probability of no more than one detected word error every 100 ms and a 90% probability of no more than one undetected word error every 4 h.

Another important parameter considered was the effect on worst case latency in a system with varying message priorities. Latency is defined as the time delay from when a transmitting station is ready to send data until the time that a receiving station has received the full message. Specifically, the criterion is to minimize the latency of the lowest priority message on a network of 128 stations with 4 stations transmitting.

As mentioned, the SAE AE-9B Committee and HART Subcommittee gathered inputs from many sources. These included a number of questionnaires and presentations as well as computer simulation results to evaluate the relative performance of the various proposed networking schemes. After a prolonged effort to reach a consensus resulted in a stalemate, it was finally resolved to formulate not one, but two high-speed data communication standards. The two standards are AS4074 and AS4075.

The High Speed Linear Bus standard calls for a "broadcast" medium, whereby all stations are connected to a common physical medium. For a fiber optic implementation, this probably entails the use of a "star" coupler rather than a truly "linear" structure. Similar to HSRB, the linear bus utilizes a token-passing protocol. This is achieved by means of a "logical ring" superimposed on a physical bus topology.

A token ring communication system is characterized by a number of stations connected in a complete circuit by unidirectional point-to-point links. A message is transmitted on the ring by an originating station, passes through all intervening stations to the destination station, where it is copied and continues around the ring until it returns to the originating station, where it is removed. The right to originate a message, as opposed to the action of repeating data appearing at the input of a station, is transferred by means of a token. A token is created following power-up and initialization, and is claimed each time a message is transmitted. Following message transmission (MT) a new token is created in the station originating the message.

The performance advantages of the ring stem primarily from the compatibility of the token passing access method with the way in which stations are connected. In any token passing scheme, a free token is circulated around the system by passing it from one station to another. The ring requires no action from a station in passing a token other than to function as a passive repeater, allowing the token to pass through. This feature has implications not only for the time required to circulate the free token, but also in allowing for a truly equitable scheme of assigning message priorities to messages in a deterministic manner. A further important advantage of a ring system is its simple media interface requirements of unidirectional point-to-point links between stations.

FOREWORD (Continued)

Ring systems can be designed to be highly reliable. The individual point-to-point links of the ring offer the considerable advantage that they can be used to support dynamic redundancies. The basic topology of the HSRB is the counterrotating ring; while the HSRB can be used in simplex systems, the requirement for dual redundancy is most common, and dual redundant operation for the HSRB has been defined completely in the standard document. The dual ring provides static redundancy in the sense that there is a replicated secondary path for data should the primary path fail. Two further mechanisms provide dynamic redundancy, with the ability to reconfigure around a fault. The first mechanism is a bypass contained within the Ring Interface Module (RIM), activated in the event of station failure, and applicable to both simplex and dual redundant systems. The bypass not only overcomes the system failure mode for station failure but also provides a means for coupling signals into the ring and achieving spatial diversity in dual systems.

The second dynamic redundancy mechanism is loopback, a capability unique to counterrotating rings. Loopback operates by providing a path through a station to connect the two rings, forming a single ring and excluding the fault.

Development of the Standard:

This document has been developed in the form of a main document and slash sheets. The main document contains all of the generic information on semantics and protocols, and the slash sheets provide information specific to an implementation. This overall approach has been taken to maximize the useful life of the HSRB standard. Rings have the major advantage that, in principle, the clock rate can be increased as the implementation technology advances with an almost prorata increase in throughput without the need for modification of the protocols. Consequently there is no reference in the main HSRB document to data rates or other implementation specific parameters and the protocols and semantics have been designed specifically to allow for growth.

Particular acknowledgement is given to the following task group members for their effort in prototype development:

Dr. Manvel Geyer
Rob DeCoco
Kevin Endlich

TABLE OF CONTENTS

1. SCOPE	8
2. REFERENCES	8
2.1 Applicable Documents	8
2.1.1 SAE Publications	8
2.1.2 Other Publications	8
2.2 Definitions	9
2.3 Acronyms	9
3 EXPLANATION OF THE STANDARD	15
3.1 General Description	15
3.1.1 Ring Bus Topology	15
3.1.2 Media and Signaling Rate Independence	18
3.1.3 Frames, Symbols and Code Bits	18
3.1.4 Token Passing	21
3.1.5 MT	22
3.1.6 Validity Checking	22
3.1.7 Reconfiguration	23
3.1.8 Error Detection and Recovery	24
3.1.9 Synchronization	25
3.1.10 Tradeoffs, Rationales and Advantages of the HSRB Standard	27
3.1.11 Token Passing Rings	31
3.2 Operational Requirements	33
3.2.1 Serial Transmission	33
3.2.2 Word Length	34
3.2.3 ML	34
3.2.4 Electrical Isolation	34
3.2.5 Self-Test	34
3.3 Performance Requirements	34
3.3.1 Station Delay	34
3.4 Semantics	34
3.4.1 Entities	36
3.4.2 Symbols	36
3.4.3 PDUs	36
3.4.4 Fields	43
3.5 Protocols	43
3.5.1 Protocol State Definitions	43
4. APPLICATION ISSUES	51
4.1 Topology	51
4.1.1 Ring Redundancy	51
4.1.2 Topology Design	52

TABLE OF CONTENTS (Continued)

4.2	Protocol	53
4.2.1	General System Considerations	53
4.2.2	Priority	53
4.2.3	Latency Control	55
4.2.4	Short Message Option	56
4.2.5	Message Retry	56
4.2.6	Addressing Modes	57
4.2.7	Loopback	58
4.3	Bridges	58
4.3.1	Bridging on the HSRB	59
4.3.2	Global Addressing	60
4.3.3	Four Rings with Serial Bridging	61
4.3.4	Four Rings with Parallel Bridging	64
4.3.5	Message RETRY - Through a Bridge	65
4.4	Global Time Reference	65
5.	IMPLEMENTATION ISSUES	68
5.1	RIU Architecture	68
5.1.1	RIU Example Implementation	70
5.2	RIU Storage Requirements	71
5.3	Synchronization and Elastic Buffer Operation	72
5.3.1	Background of Token Ring Synchronization	72
5.3.2	Conventional Independent Link Synchronization	73
5.3.3	Improved Independent Link Synchronization	74
5.3.4	Clock Tolerance	75
5.4	Media	76
5.4.1	Optical Media	76
5.4.2	Electrical Media	77
5.5	Bypasses	78
5.5.1	Electrical Bypassing	79
5.5.2	Optical Bypasses	83
6.	INTEROPERABILITY ISSUES	90
6.1	Physical Layer	90
6.2	Options	90
6.2.1	Retry Option	90
6.2.2	Multiple Short Message Options	90
6.3	Message Acknowledge	90
6.4	Message Priority	90
6.5	Logical Addressing	91
6.6	Bridges	92

TABLE OF CONTENTS (Continued)

6.7	Host/RIU Interface	92
6.7.1	Control	92
6.7.2	Status	92
6.7.3	Information	94
7.	PERFORMANCE CALCULATION	94
7.1	Introduction	94
7.2	Ring Equations	95
7.2.1	RRT	95
7.2.2	ML	99
7.2.3	ML Equal to RRT	102
7.3	Throughput Efficiency	105
7.3.1	One Station Transmitting	105
7.3.2	All Stations Transmitting	108
7.3.3	Other Cases	109
7.3.4	Effect of Priority Levels	109
7.4	Message Latency	110
8.	TEST CONCEPTS	110
8.1	Self-Test	110
8.2	Background Test	112

SAENORM.COM : Click to view the full PDF of air4289a

1. SCOPE:

This Handbook has been prepared by the Ring Implementation Task Group of the SAE AS-2 Committee, and is intended to support AS4075 by providing explanation of the standard itself and guidance on its use.

The principal objective in the preparation of a standard is to provide a statement of operational and performance requirements, and an unambiguous definition of the functions to be realized in any implementation, primarily from the view point of interoperability. While efforts have been made within the AS4075 standard to provide a readable general description of the HSRB, detailed explanations, rationale and guidance to the use are incompatible with the purpose and, indeed, the format of a standard. Accordingly, this Handbook contains a paragraph-by-paragraph explanation of the main sections of the standard, and a discussion of application and implementation issues.

2. REFERENCES:

2.1 Applicable Documents:

The following publications form a part of this specification to the extent specified herein. The latest issue of all SAE Technical Reports shall apply.

2.1.1 SAE Publications: Available from SAE, 400 Commonwealth Drive, Warrendale, PA 15096-0001.

AS4075	High Speed Ring Bus Standard
AS4075/1	Optical Implementation Relating to the High Speed Ring Bus (HSRB) Standard

2.1.2 Other Publications:

Bicknell, J: 'SAE High Speed Data Bus Standards', ERA Seminar on Military Avionics; London, November 1988.

Berggren, C: 'Communications for Distributed Real-Time Systems'; April 1987.

DeCoco, R J; Kroeger, B W; Kurtz, J J: 'An Overview of SAE AE-9B High Speed Ring Bus (HSRB) Performance'; SAE Aerospace Avionics Equipment Conference, April 1986.

Glass, M: 'A Look at the SAE High Speed Ring Bus'. Defense Electronics, Sept 1988.

Walser, R A: 'Concepts and Operational Experience for 9B Ring Bus Standards Development'; SAE AE-98 Meeting, April 24, 1986.

SAE 'Technical Rationale for the HART Document'; Issue 4, SAE AE-9B High Speed Data Bus Applications and Requirements Task Group, April 1986.

2.2 Definitions:

4B5B Encoding	An encoding method whereby four information data bits are encoded into five-bit symbols.
Delimiter	A symbol or sequence of symbols used to mark the boundary of a field. The HSRB standard specifies six types of delimiters: Token Starting, Token Ending, Message Frame Starting, Message Frame Ending, Beacon Frame Starting and Beacon Frame Ending.
Deterministic	Property of an item to which the future behavior can be predicted precisely.
Fault Tolerance	Capability of the system to endure component errors and/or failures without causing total system failure. Actions range from ignoring it, to retrying action, to complex actions of fault isolation, and then taking positive action to continue operation without the failed component (i.e., invoking backup, reconfiguring, invoking degraded operation, etc.).
Protocol	A set of related rules describing specific processes or activities.
Reconfiguration	Process by which the operation of the HSRB is initialized using Beacon frames following the application of power, the presence of failures, or the detection of certain errors using Beacons.
Semantics	The relationships between symbols and their meanings.
Topology	The arrangement or layout of elements comprising a system.

2.3 Acronyms:

AS	Aerospace Standard
AS-2	Interconnect Networks Committee of ASD
ASD	Avionic Systems Division
a.c.	Alternating current (refers to ac coupled)
AC	Address Control Field
ACK	Acknowledgement
ACK1	Message Acknowledged 1 subfield
ACK2	Message Acknowledged 2 subfield
ADJ	Adjustment subfield
AGC	Automatic Gain Control
ANSI	American National Standards Institute
AW0	Address Word Count Bit 0 (least significant)

2.3 (Continued):

AW1	Address Word Count Bit 1 (most significant)
AWC	Address Word Count subfield
BA	Bridge Access
BCON	Beacon Control
BFCS	Beacon Frame Check Sequence Field
BFCS0	Beacon Frame Check Sequence Bit 0 (least significant)
BFCS15	Beacon Frame Check Sequence Bit 15 (most significant)
BFED	Beacon Frame Ending Delimiter Field
BFSD	Beacon Frame Starting Delimiter Field
BLT	Beacon Loop Timer
BPI	Beacon Path Indicator subfield
BT	Beacon Type subfield
CON	Control field
CSMA	Carrier Sense Multiple Access
DA	Destination Address Field
DAL	Destination Address Logical subfield
DAP	Destination Address Physical subfield
dB	Decibel
DL	Destination Address Logical
DP8	Destination Address Physical bit 8 (most significant subaddress bit)
DP9	Destination Address Physical bit 9 (least significant physical address bit)
DP15	Destination Address Physical bit 15 (most significant physical address bit)
4B5B	Four bits five bits (an encoding scheme)
FDDI	Fiber Distributed Data Interface
FS	Frame Status Field
FCS	Frame Check Sequence
FET	Field Effect Transistor
g	Gravity
GA	Global Address Field
GAL	Global Address Logical subfield
GAP	Global Address Physical subfield
GL	Global Address Logical

2.3 (Continued):

GL3	Global Address Logical Bit 3
GL2	Global Address Logical Bit 2
GL1	Global Address Logical Bit 1
GL0	Global Address Logical Bit 0
GP	Global Address Physical Bit
GP0	Global Address Physical Bit 0 (least significant)
GP1	Global Address Physical Bit 1 (most significant)
HART	High Speed Data Bus Applications and Requirements Taskgroup
HKA	Highest Known Address Field
HKA0	Highest Known Address bit 0 (least significant)
HKA6	Highest Known Address bit 6 (most significant)
HSRB	High Speed Ring Bus
Hz	Hertz
IB	Information Bit
IB0	Information Bit 0 (least significant)
IB15	Information Bit 15 (most significant)
IED	Information Error Detected
IED1	Information Error Detected 1 subfield
IED2	Information Error Detected 2 subfield
IEEE	Institute of Electrical and Electronics Engineers
IFA	Interframe Adjustment field
IFCS	Information Frame Check Sequence field
IFCS0	Information Frame Check Sequence bit 0 (least significant)
IFCS15	Information Frame Check Sequence bit 15 (most significant)
INFO	Information field
INT	Integer function
IW	Information Word
km	Kilometer
LC	Liquid Crystal
LiNbO ₃	Lithium Niobate
LiTaO	Lithium Tantalate
LP	Logical Physical subfield

2.3 (Continued):

LTC	Loop Timer Counter
LUT	Look-Up Table
Mb	Millions of signalling bits per second
Mbaud	Megabaud (10^6 baud) (also see Mbd)
Mbd	Megabaud (also see Mbaud)
MCED	Message Control Error Detected
MCED1	Message Control Error Detected 1 subfield
MCED2	Message Control Error Detected 2 subfield
MCFCS	Message Control Frame Check Sequence field
MCFCS0	Message Control Frame Check Sequence bit 0 (least significant)
MCFCS15	Message Control Frame Check Sequence bit 15 (most significant)
MFED	Message Frame Ending Delimiter Field
MFSD	Message Frame Starting Delimiter
MHz	Megahertz (10^6 hertz)
ms	Millisecond (10^{-3} second)
NRZ	Non-Return to Zero
NRZI	Non-Return to Zero Invert
OSI	Open Systems Interconnect (an International Standards Organization reference model)
pf	Pico Farad (10^{-9} Farad)
P0	Token Priority Bit 0 (least significant)
P1	Token Priority Bit 1
P2	Token Priority Bit 2 (most significant)
PDUS	Protocol Data Units (also PDUs)
PM0	Message Priority Bit 0 (least significant)
PM1	Message Priority Bit 1
PM2	Message Priority Bit 2 (most significant)
PR	Priority subfield
PRM	Message Priority subfield
PRS	Priority and Retry Status
R0	Reservation Bit 0 (least significant)
R1	Reservation Bit 1
R2	Reservation Bit 2 (most significant)

2.3 (Continued):

RAM	Random Access Memory
RCVD	Received
RCVD1	Message Received 1 subfield
RCVD2	Message Received 2 subfield
RCVST	Receiver Status (Table 3.5.15-3)
RES	Reservation (contents 3.4.4.2.5)
RIM	Ring Interface Module
RIU	Ring Interface Unit (3.5.15.1)
RIUC	Ring Interface Unit Controller
ROM	Read-Only Memory
RSI	Retry Status Indicator subfield
S0	SMC Bit 0 (least significant)
S1	SMC Bit 1
S2	SMC Bit 2
S3	SMC Bit 3 (most significant)
S1	Unconnected State
S2	Reconfiguration State
S3	Inactive Ring State
S4	Repeater Ring State
S5	Active Ring State
S1.1	Unpowered State
S1.2	Powered State
S1.3	Quiescent State
S1.4	Wait State
S1.5	Synchronization State
S1.6	Listen State
S2.1	Reconfiguration Initialization State
S2.2	Vie State
S2.3	Master/Configuration State
S2.4	Slave State
S3.1	Inactive State
S4.1	Repeater State

2.3 (Continued):

S5.1	Await Token Starting Delimiter State
S5.2	Check Token Priority State
S5.3	Reserve Token State
S5.4	Repeat Message State
S5.5	Token Claim State
S5.6	Originate Message State
S5.7	Issue Token State
S5.8	Strip Message State
S5.9	Check Beacon State
S5.10	Warm Start State
SA6	Sending Address Bit 6 (most significant)
SA0	Sending Address Bit 0 (least significant)
SA	Sending Address field
SAE	Society of Automotive Engineers
SC	Station Count field
SC0	Station Count bit 0 (least significant)
SC6	Station Count bit 6 (most significant)
SMC	Short Message Count
T1	Token Status 1 bit (3.4.4.2.2)
T2	Token Status 2 bit (3.4.4.2.4)
TED	Token Ending Delimiter
TNLC	Twisted Neumatic Liquid Crystal
TSD	Token Starting Delimiter field
TS1	Token Status 1 subfield
TS2	Token Status 2 subfield
TSDC	Token Starting Delimiter Counter
UTIT	Uncontrolled Transmit Inhibit Time
VCXO	Voltage Controlled Crystal Oscillator
WC11	Word Count Bit 11 (most significant)
WC0	Word Count Bit 0 (least significant)
WC	Word Count field
WP	Wollaston Prism

3. EXPLANATION OF THE STANDARD:

An explanation of the general requirements, corresponding to Section 3, is given in the following pages. Sections 3.2 to 3.5 are a paragraph-by-paragraph explanation of the standard. Paragraph numbering has been preserved as far as possible for clarity. Some figures referenced by this document are found in AS4075 and are so referenced.

3.1 General Description:

- 3.1.1 Ring Bus Topology: The physical topology of the SAE ring bus is illustrated in Figure 3.1-1. While it is allowable under AS4075 to implement a single (nonredundant) ring network, most military applications will demand dual (or greater) redundancy. The data communications system shown in the figure consists of 2 to 128 stations connected by a series of point-to-point links over two parallel, closed paths. The two rings are counterrotating.




FIGURE 3.1-1 - Ring Bus Topology

The architecture of a typical station for the HSRB is illustrated in Figure 3.1-2. Note that each station consists of two Ring Interface Modules (RIMs) for media interface to the two rings, a single Ring Interface Unit (RIU) for protocol control and a host interface. The host is generally a microprocessor. The bypasses for the two RIMs allow ring operation to continue despite the presence of a failed or unpowered station.




FIGURE 3.1-2 - Ring Bus Station Architecture

3.1.1 (Continued):

The dual ring topology has three inherent mechanisms to enhance system reliability: first, dual redundancy allows for operation on the alternate ring in the case of a failed RIM or media segment; secondly, the RIM bypasses enable an unpowered or failed station to appear "transparent" so that the rest of the system can continue to operate; thirdly, in the case of media faults on both rings, it is possible to configure any given station in a loopback configuration (Figure 3.1-3). A loopback station receives its signal from one ring and retransmits the signal on the alternate ring. When configured for the loopback mode, one ring is used for the active reception and transmission of messages, while the RIUs on the alternate ring act as passive repeaters.




FIGURE 3.1-3 - Loopback Configuration

3.1.1 (Continued):

The HART document specifies that there must be a bridging mechanism, to support the interconnection of two or more buses (Figure 3.1-4). AS4075 provides the "hooks" for bridging to as many as three other rings, thereby increasing the number of stations that may be (indirectly) accessed from 128 to 512. Bridging is supported in HSRB protocol by the inclusion of two "global" address bits for physical addressing and four "global" address bits for logical addressing.




FIGURE 3.1-4 - Bridging Between Two Rings

- 3.1.2 **Media and Signaling Rate Independence:** One of the principal tenets of the HART document states that the high speed standards are to be independent of both data rate and media. AS4075 was partitioned into a main body and slash sheet(s) to provide this independence. The main body contains the media independent portion of the standard. Slash sheets provide the mechanism for specifying the media and signaling rate parameters for various implementations (electrical, fiber optic, etc.). This allows for maximum flexibility by allowing data rate, media and the parametric specifications regarding signal transmission and reception to be varied for each implementation.
- 3.1.3 **Frames, Symbols and Code Bits:** AS4075 specifies three types of frame: Token, Message and Beacon frames (Figure 3.1-5). The Token frame provides the mechanism for the distribution of access control for MT. The Message frame is used for the actual transfer of information. The Beacon frame is used during periods of power turn-on initialization and in fault recovery sequences for the purpose of ring reconfiguration control.

Figure 3.1-6 illustrates the encoding scheme specified by the HSRB standard. The three frame types are formed by combinations of symbols and code bits. As shown in Figure 3.1-6(a), the polarity of all signaling bits transmitted on the physical medium is determined by means of Non-Return-to-Zero Inverted (NRZI) encoding. That is, if the next symbol bit or code bit is a zero, the polarity of the next signaling bit will remain unchanged from the previous polarity; if the next symbol bit or code bit is a one, the polarity of the next signaling bit toggles from its present value. The formulation of symbol and code bits (discussed below) precludes the transmission of more than three successive signaling bits of the same polarity. Limiting the unbalance of signal patterns serves to reduce dynamic offset effects for electrical and fiber optic networks.

The encoding scheme used to convert symbols to signaling bits is 4-bit/5-bit (4B5B). As shown in Figure 3.1-6(b), 4B5B encoding converts 4 bits of data to 5 symbol bits. The symbol bits are then converted to signaling bits by means of the NRZI encoding. Each of the 16 possible 4-bit (hex) combinations is converted to a unique 5-bit symbol. Therefore, each 16-bit word in the Information field (INFO) of a Message frame consists of four symbols.

In addition to the 16 symbols representing hex bytes, Figure 3.3-6(b) lists a number of "nonhex" symbols, used for special purposes. The principal use of these symbols is to form the Start and End Delimiter fields for the three frame types. The Token Starting Delimiter (TSD) field is made up of 10 signaling bits: a "J" symbol followed by a "K" symbol. Message frames are appended to Claimed Tokens by means of a Message Frame Starting Delimiter (MFSD): a "J" symbol followed by an "A" symbol. During periods of reconfiguration, all Beacon frames start with a Beacon Frame Starting Delimiter (BFSD) made up of a "K" symbol followed by a "J" symbol. For all three frame types, the End Delimiter consists of a single "T" symbol.

IDLE ("I") symbols consist of five consecutive ones, resulting in a pattern of alternating ones and zeros transmitted on the physical medium. IDLEs are used in Token and Message frames as part of the HSRB Clock recovery (synchronization) scheme. IDLE symbols are also used following some MT sequences and during reconfiguration as a means of keeping the entire ring "full" of signaling bits in order to maintain synchronization.





FIGURE 3.1-5 - Ring Bus Frame Types (a) Token Frame; (b) Message Frame; (c) Beacon Frame



Symbol	Code Group	Assignment
0	11110	Data Symbol '0000' Binary
1	01001	Data Symbol '0001' Binary
2	10100	Data Symbol '0010' Binary
3	10101	Data Symbol '0011' Binary
4	01010	Data Symbol '0100' Binary
5	01011	Data Symbol '0101' Binary
6	01110	Data Symbol '0110' Binary
7	01111	Data Symbol '0111' Binary
8	10010	Data Symbol '1000' Binary
9	10011	Data Symbol '1001' Binary
A	10110	Data Symbol '1010' Binary. Also used with a preceding 'J' symbol as the message frame starting delimiter.
B	10111	Data Symbol '1011' Binary
C	11010	Data Symbol '1100' Binary
D	11011	Data Symbol '1101' Binary
E	11100	Data Symbol '1110' Binary
F	11101	Data Symbol '1111' Binary
I	11111	Idle. No information currently being transmitted.
J	11000	Used for token, message frame, and beacon frame starting delimiters.
K	10001	Used for token and beacon frame starting delimiters.
Q	00000	Quiet. No signal being transmitted.
T	01101	Terminate. Ending delimiter for token, message frame and beacon frame.
S	11001	Never transmitted as a symbol, but may occur in a token or frame status.
V	00001	Invalid. These symbols shall not be used as they violate consecutive code-bit zeros, duty cycle, or prevent unique starting delimiter patterns.
V	00010	
V	00100	
V	01000	
V	10000	
V	00011	
V	00101	
V	00110	
V	00111	
V	01100	

FIGURE 3.1-6 - 4-bit/5-bit NRZI Encoding (a) Simplified Block Diagram;
(b) 4B5B Symbol Coding Table

3.1.3 (Continued):

Code bits, rather than symbols, are used in the various "Control" subfields of Token frames as well as in the "Frame Status" field of Message frames. In these fields, it is necessary for stations to flip various bit values in "real-time." By eliminating the 4B5B encoding/decoding for such functions as the Token Status, Priority and Reservation subfields in Token frames as well as the Error and Acknowledge bits in Message Check Sequences, it is possible to specify a maximum "per station" delay of six signaling bit times.

3.1.4 Token Passing: During normal operation, access to the HSRB is controlled by continually passing a Token frame around the ring. When in the token passing mode, each station behaves as a passive repeater.

There are two types of Token frames: free tokens and claimed tokens. The two Token Status bits are used to distinguish between the two types. When there are no stations vying to transmit, the free token will continuously circulate around the ring. When a station has a message pending, it flips the Token Status bits creating a claimed token and appends its message to the claimed Token frame, forming a Message frame. Following the end of the Message frame, the transmitting station issues a new free token.

One of the salient features of the HSRB is its protocol for message prioritizing. This is accomplished by means of the Priority and Reservation fields of the Token frame. When a message is circulating around the ring, any station may use the Reservation field to indicate that it has a message pending. This is done by updating the value of the Reservation field of the claimed token preceding the current message to the priority of the vying station's message, providing that the priority of the pending message is higher than the priority currently indicated by the Reservation field. The Reservation field may assume a value between 0 and 7, with 0 indicating the highest priority level. When the transmitting station issues the new free token, the value of its Priority field is assigned the value from the Reservation field of the returned claimed token. The free token may then be claimed, and another Message frame transmitted by any downstream station with a message priority equal to or greater than the priority of the free token.

An exception to the normal token passing sequence involves Short Message protocol. A short message is defined as a message that is shorter than the time to propagate around the ring plus the time to process the reservation field of the returning token. In some scenarios using the Short Message protocol, it is possible to have two or more messages propagating on the ring at the same time. By means of the Short Message Count (SMC) field, the HSRB standard allows for as many as 15 consecutive short messages to be transmitted without regard to the normal priority protocol. This is implemented by loading the value of 7 (lowest priority) into the token Priority and Reservation fields following each short message (up to 15), instead of waiting for the claimed token portion of the Message frame to return to the transmitting station, as is otherwise done. After 15 short messages have been transmitted, it becomes necessary to observe the normal priority protocol by transmitting IDLE symbols until the claimed token returns and then loading the returned Reservation field into the Priority field of the new free token. In systems where message prioritization ordering is considered more important than bus bandwidth, the HSRB standard allows for disabling the Short Message option by always loading a value of 15 into the SMC field of free tokens when they are issued.

- 3.1.5 MT: Information transfer on the HSRB is implemented by means of the Message frame. In Message frames, data are always transferred from the station claiming the free token to one or more destination stations. As shown in Figure 3.1-5(b), the Message frame is formed by appending additional fields to a claimed Token frame. The Message Priority field indicates the actual priority level of the message, which must be equal to or greater than the Priority field of the claimed token. Provision for a single message retry is included in the HSRB protocol. Such messages are identified by means of the Retry Status bit. The number of 16-bit words to be transmitted may vary from 1 to 4096 and is indicated by the Word Count field. The Sending Address field specifies the 7-bit physical address of the transmitting station.

As mentioned, AS4075 supports two types of destination addressing: logical and physical. For a particular message, the addressing mode is specified by the Logical/Physical subfield. In addition, the standard supports Global Addressing, providing the "hooks" to communicate with stations on the other rings, via bridges.

For physical addressing, the two-bit Global Address Physical (GAP) subfield identifies the number of the destination ring. By definition, zero (0) is always designated as the Global Address of the ring containing the originating station. The 16-bit Destination Address Physical (DAP) field is subdivided into a 7-bit Physical Address supporting 128 physical stations and a 9-bit subaddress field, allowing for 512 physical subaddresses per station.

In the logical addressing mode, it is possible to send data to the stations on the originating ring and/or to any combination of up to three other rings, via bridging. The destination ring(s) is (are) identified by the four bits of the GAL subfield. For any particular message, the length of the Logical Address field is identified to be 16, 32, 48, or 64 bits by the Address Word Count (AWC) subfield. This specifies the width of the Logical Address subfield to be from one to four words, as determined by system requirements.

The Information/Adjustment field contains between 1 and 4096 16-bit data words, as indicated by the Word Count Field. In addition to the data words, this field also contains adjustment subfields (ADJs), providing that the Word Count is 256 or greater. ADJs, which consist of 6 IDLE symbols following by a J and A symbols, are embedded following every 256 words of the INFO to allow compensation for clock differences as part of the standard's synchronization (clock recovery) scheme.

- 3.1.6 Validity Checking: As a mechanism for verifying the valid transfer of a transmitted Message frame, the HSRB standard calls for two check sequences. Both are based on the same residue generator polynomial. The Message Control Frame Check Sequence (MCFCS) field encompasses the Priority/Retry Status, Word Count, Source Address, Address Control, Global Address (GA) and Destination Address (DA) fields. A second check field, the Information Frame Check Sequence (IFCS) field covers the INFO.

3.1.6 (Continued):

The final field of the Message frame is the Frame Status (FS) field. The FS field enables the station receiving a message to inform the transmitting station that the message has been received, acknowledged and verified as correct. The active bits of the FS field are all initialized by the transmitting station and may later be changed by the receiving station. The two Message Acknowledge bits indicate that the full message has been received and made available to the receiving station's host. The two Message Received bits indicate that the station has received a message containing its address. The two Message Control Error Detected and the two Message Error Detected bits indicate disparities between the received and calculated values of the MCFCS and IFCS fields, respectively.

The HSRB information transfer sequence is summarized as follows: the station claiming the token appends its transmitted Message frame to the claimed Token frame. The transmitted message is repeated by each successive station around the active ring, including the receiving station(s). Receiving station(s) copy and acknowledge the Message frame, as the message is repeated by all stations. When the message returns to the transmitting station, it is stripped off the ring. In the process of stripping the message, the transmitting station alerts its host regarding any and all fault conditions indicated by the returned FS field. Immediately after completing transmission of the Message frame, the transmitting station issues a new free Token frame, followed by IDLE symbols. An exception to this involves the completion of the last (fifteenth) consecutive short message to be transmitted; in this instance, IDLE symbols are transmitted until the Message frame returns, and then the free token is issued.

- 3.1.7 Reconfiguration: The Beacon frame is used only during periods following station power turn-on or during fault recovery sequences. The purpose of the reconfiguration sequence, which entails the transmission of Beacon frames among the currently operational stations in a system, is to establish one of three possible signaling paths to utilize for normal operation. The three reconfiguration possibilities are: (1) Ring 0 Active, (2) Ring 1 Active, and (3) Loopback.

The Beacon Control (BCON) field consists of two subfields: the Beacon Type (BT) and Beacon Path Indicator (BPI) subfields. The three bits of the BT subfield identify which of the seven Beacon types is presently being passed around the ring. The seven types of Beacon frames are: Warm Start, Warm Recover, Restart, Vie, Configure Ring 0, Configure Ring 1, and Configure Loopback. The Warm Start and Warm Recover Beacons are used for the Warm Start sequences, involving recovery to "soft" errors. All of the other Beacon types are utilized in the Restart, or Reconfiguration sequence.

Following certain "soft" fault conditions, a Warm Start sequence may be initiated by any station detecting such an error. After the "master" station repeats the Warm Start Beacon, it issues a Warm Recovery Beacon followed by a free token to resume normal operation.

3.1.7 (Continued):

Scenarios requiring the full Restart configuration sequence include following station power-up, as a result of one of several "catastrophic" fault conditions, or following a failed Warm Start sequence. The reconfiguration sequence is initiated by the transmission of a Restart Beacon. Next, Vie Beacons are transmitted and retransmitted in a point-to-point manner by all active stations. The Vie portion of the reconfiguration sequence involves making a determination of the functional communication path. The BPI bit is used in the Vie process to determine if there is continuity in the Ring 0 and/or Ring 1. The Station Count (SC) field is used to determine which of the two ring paths has the greater number of functional stations. The Highest Known Address (HKA) field is used to identify the station with the highest physical address, which is then designated as the "master" station on the ring. The "master" stations then determines the active transmission ring path, and issues the appropriate type of Configuration Beacon. Following the Configuration Beacon, the "master" station issues a free token to commence normal operation. Similar to the MCFCS and IFCS fields for Message frames, the Beacon Frame Check Status (BFCS) field provides a mechanism for checking the validity of Beacon frames.

- 3.1.8 Error Detection and Recovery: As required by the HART document, the HSRB standard calls for a number of mechanisms to support robust operation in a real-time environment, with minimal recovery times in response to a number of fault conditions. Most of these self-healing mechanisms are directly embedded in the lower level protocol and are, therefore, implemented with minimal or no involvement of the station host.

AS4075 calls for a variety of fault recovery schemes ranging from "no action required" to full reconfiguration. Faults including Token Priority too low, claimed Token Reservation too high or too low, and SMC errors have only a minor effect on ring operation and require no corrective action. Single (or multiple) bit errors in Token or Message frames will generally be evidenced by the receipt of invalid starting delimiters or symbol(s) and/or incorrect MCFCS or IFCS.

These conditions require no immediate corrective action to be taken; however, a station detecting one of these errors is required to notify its host of the fault. Similarly, a transmitting station will notify its host if it detects a nonreception or nonacknowledged condition in the FS field of a Message frame. The host has the option of selecting a mode incorporating a message retry. The HSRB standard supports this by the inclusion of the Retry bit in Message frames, which indicates the retry status of the Message frame.

Another class of detected errors indicates either catastrophic fault conditions and/or conditions that may "hang" the bus. These conditions result in a "Warm Start" beaconing sequence in order to reinitialize the ring without having to go through a full reconfiguration sequence. The Warm Start error conditions include Lost Token Start Delimiter, Token Priority too high, Token Status bit disagreement, Token frame format errors, and fault conditions that are continuously monitored by a number of timer and counter mechanisms. The Lost Free Token Counter detects two consecutive messages containing the same Sending Address without an intervening free Token. Another timer mechanism, the Uncontrolled Transmit Inhibit Timer (UTIT), inhibits MT after 1.1 times the maximum message length (ML) of 4096 words.

3.1.8 (Continued):

For the most catastrophic failure modes, the HSRB standard calls for a full "cold start" reconfiguration sequence. The conditions requiring full reconfiguration are, loss of input on active ring, acquisition of input from either receiver, Warm Start failure and station insertion.

3.1.9 Synchronization: One of the major topics of discussion during the development of the HSRB standard was the specification of the synchronization, or clock recovery scheme. A synchronization mechanism is required in order for each station on the ring to synthesize a decoder clock, sample its received bit stream and repeat the signal to the next downstream station on the ring.


During the HSRB standard development process, two synchronization schemes were considered. The two techniques, as illustrated in Figure 3.1-7, are the "single clock master" scheme and the "point-to-point" scheme.

The "single clock master" scheme is utilized in IEEE standard 802.5 (the IBM Token Ring) and is shown in Figure 3.1-7. In this scheme, one station emerges from the reconfiguration process as the clock master station. The timing of the signaling bit stream transmitted by the master station is derived from a crystal oscillator in the station. All stations on the ring include a clock recovery device (possibly a phase-locked loop). The function of the clock recovery device is to synthesize a decoder sampling clock that is synchronized to the transition times of the incoming bit stream and ultimately to the master station.

For the "single clock master" scheme, all stations other than the master station synchronize their transmission as well as reception to the synthesized clock. In so doing, the transmission of signaling bits for all stations on the ring is essentially synchronized so that there is no "bit slippage" possible between adjacent stations, since they all transmit and receive at the same controlled rate. The clock master station includes a 128-bit elastic buffer, or bit FIFO between its receiver and transmitter sections. In addition to enabling the circulation of a Token frame on a very short ring, the elastic buffer provides a mechanism to compensate for effects, such as delay drifts in transmitters, receivers and media due to warm-up or temperature, that may either "shorten" or "lengthen" the effective length of the ring (total bit delay).

The "single clock master" algorithm was seriously considered for the HSRB standard, but was ultimately rejected in favor of the "point-to-point" scheme. While the "single clock master" technique did offer a slight performance advantage and has been successfully implemented at 10 MHz in IEEE Std. 802.5, it was felt that the accumulated phase jitter resulting from intersymbol interference and zero-crossing noise would result in difficult and unreliable operation for a 128-station 100 MHz ring.

The "point-to-point" synchronization scheme is utilized in the ANSI Fiber Distributed Data Interface (FDDI) standard and is shown in Figure 3.1-7. Similar to the "single clock master" scheme, each station contains a clock recovery device (one for each ring). However, instead of deriving the timing for the retransmitted signaling bit stream from the synthesized clock, each station synchronizes its transmitted signal from its internal crystal oscillator, independent of the decoder clock.



- (a) Point-to-point synchronization scheme
(b) Single clock master scheme

FIGURE 3.1-7 - Ring Synchronization Schemes

3.1.9 (Continued):

One problem that must be overcome in implementing the "point-to-point", or independent, synchronization scheme is the problem of bit slippage between independent stations, due to clock frequency differences. Bit slippage can cause a station to start losing bits, thus causing errors in messages. This bit slippage is avoided in the HSRB by specifying a maximum clock frequency tolerance of $\pm 0.009\%$, adding an Adjustment Field consisting of IDLE symbols at strategic times in the Message frame and requiring a small elastic buffer for all stations on the ring. The clock frequency tolerance guarantees a maximum bit slippage from one station to the next. The IDLE symbols provide areas of the Message frame to add to or delete bits in the elastic buffer in order to compensate for the bit slippage. The Message frame includes 6 IDLE symbols (at origination) between every 256 words in the INFO fields of Message frames to allow for the elastic buffer operation.

For the "point-to-point" synchronization scheme, the small elastic buffers distribute the adjustment function among all the stations on the ring, eliminating the need for the large elastic buffer required for the "single clock master" scheme. However, the "point-to-point" technique does demand that a small buffer (40 bits) be implemented in a single station, designated as the "master" (as a result of reconfiguration) for the purpose of allowing token passing operation on a very small ring.

As specified in the HSRB standard, the independent synchronization scheme permits the use of simple, economical clock recovery devices while providing reliable system operation. The performance penalty, in comparison to the "single clock master" scheme is small: an additional 2-bit station delay time and a maximum 0.8% throughput reduction for the INFO fields of Message frames.

3.1.10 Tradeoffs, Rationales and Advantages of the HSRB Standard: In terms of topology, protocol and survivability, this document offers a number of advantages over other networking standards.

The ring topology provides a number of inherent advantages over linear and star buses. The ring configuration, consisting of a series of point-to-point repeater nodes, greatly reduces the amount of transmitter power required and simplifies the design of transmitters and receivers. This is in contrast to networks employing a broadcast medium, which encounter inevitable tradeoffs between total media length, the maximum number of taps, transmitter power, receiver sensitivity and bit error rate. Since each station on a ring bus is itself a repeater, there is essentially no limit to the number of stations or total ring length.

In addition to the transmitter/receiver considerations, the ring topology is particularly advantageous for a fiber optic implementation. Unlike a system employing a "broadcast" medium, there is no need for tree or star couplers. In addition to limiting the number of nodes on a network (at this time, 32 is considered the most for a passive star configuration), the use of a star coupler reduces system reliability by introducing a single point failure mode. Another advantage of the ring configuration is that it allows stations to be positioned arbitrary, independent of physical address with no effect on performance.

3.1.10 (Continued):

In some respects, the counterrotating ring topology provides more inherent reliability than does a linear bus structure. For example, media faults on the two redundant buses between adjacent nodes renders a linear bus inoperative; for the same condition on a ring bus, continued operation is made possible by reinitializing to the loopback configuration. In addition, the point-to-point topology provides protection against a station with an "uninhibited transmitter" fault. Networks employing a broadcast medium cannot continue operation in such a situation. In general, the inclusion of bypasses and the loopback mode for a point-to-point ring network allows for robust system operation over a wider range of fault scenarios than is possible for a network using a broadcast medium.

The use of token-passing protocol, in contrast to Carrier Sense Multiple Access (CSMA), provides excellent determination of latency. This is particularly important for real-time military systems, where it is imperative that individual nodes be serviced within a known, predetermined maximum time period. While CSMA provides good performance at low levels of loading (less than 20%), its throughput efficiency and worst case latency degrades severely at high utilization levels, due to access contention.

In addition to the performance and system reliability advantages provided by the use of token passing, the protocol scheme provided by the HSRB standard includes a number of other desirable features. Perhaps the most salient of these features is the access control scheme that is specified in the HSRB standard. By taking full advantage of the physical ring topology, the time required to circulate a Token frame around to all stations on a ring is held to an absolute minimum. This is due to the fact that for a repeater ring bus, the passing of Token frames (as well as Message frames) is a pipelined operation, rather than a purely serial operation (reference Figure 3.1-8). That is, the station-to-station delay time for token passing on a ring bus consists entirely of a station's total throughput delay plus the media propagation time to the next station. The maximum allowable station throughput delay is specified as six signaling bit times plus transmitter and receiver delay times, or a total of about 8 signaling bit times (80 ns) at 100 Mb. By contrast, the station-to-station token passing time for a linear or star bus is comprised of the time needed to transmit the entire Token frame plus the media delay time between the two stations (that are not necessarily in close proximity) plus the Automatic Gain Control (AGC) adjustment, resynchronization and response times for the next downstream station on the logical ring. This situation is particularly pronounced for the fiber optic star topology, in which each station-to-station token pass requires propagation from one station to the star coupler and then from the coupler to the next station on the logical ring (Figure 3.1-8(c)).

The net effect of employing a physical ring rather than a logical ring is that the time required to pass the token is minimized, allowing relatively more time for the transmission of Message frames, resulting in a very high throughput efficiency. This is of paramount importance for a ring supporting a large number of stations, particularly where long separation distances are involved and where arbitrary assigning of physical addresses is desired.




FIGURE 3.1-8 - Token Passing (a) Ring Bus; (b) Linear Bus; (c) Star Coupled Bus

3.1.10 (Continued):

Aside from throughput efficiency, the implementation of a token-passing protocol on a physical ring provides superior performance in the areas of predictability, short message operation and message prioritization. Average message latency is minimized by the efficiency of the HSRB token passing algorithm. The latency for very high priority messages is held to an absolute minimum by the use of the Reservation and Priority fields of Token frames.

Use of the optional Short Message protocol provides a means to further improve throughput performance by suspending the usual priority rules following short messages. The use of the Short Message protocol may be disabled for applications where system predictability is favored over average throughput.

To a large extent, the use of the point-to-point synchronization scheme for the HSRB simplifies the design of receivers and clock recovery devices. Another aspect of the HSRB protocol that reduces the demands on receiver AGC, filter, phase-locked loop and decoder design is the fact that under normal operation, there is always a stream of signaling bits present at each active receiver's input. This is in contrast to command/response, CSMA or token-passing on a broadcast medium where there is always "dead time" between bursts of bus activity. A receiving station is susceptible to noise-induced errors during this "dead time", particularly the time period just prior to and just following the reception of Beacon, Token or Message frames. As a result, the requirements for AGC bandwidth and rapid phase lock for a receiver on a ring bus are much less demanding than on a broadcast media. This is because under normal conditions, a receiver's AGC and clock recovery circuits are able to remain locked on to the same input signal, rather than having to continually adapt to different signals of varying amplitude, frequency and phase. In addition to eliminating the added overhead of preamble fields associated with networks using broadcast media, this contributes to achieving a very low system bit error rate. The bit error rate for the ring topology is, in any case, very low since the full signal power from one station's transmitter is received by only one station, rather than shared between all other stations in the system, as is the case for a broadcast medium.

The encoding scheme employed in the HSRB, 4B5B NRZI, provides a fundamental advantage over the Manchester encoding employed in other standards. 4B5B NRZI provides a ratio of 0.8 data bits per signaling bit, as opposed to 0.5 data bits per signaling bit for Manchester encoding. In other words, to attain a data bit rate of 100 MHz requires 200 million signaling bits per second for Manchester, but only 125 million signaling bits per second for 4B5B. This difference in encoding schemes yields a significant advantage, particularly for operation at very high data rates.

Some of the features supported by the command/response protocol of MIL-STD-1553 include RT-to-BC transfers, RT-to-RT transfers and mode code messages. In the interest of simplifying the design, fabrication, test, software development and maintenance of the HSRB hardware, these functions are not directly implemented by the HSRB standard. Nevertheless, it should not be noted that command/response or any of the other "1553-like" features, including any or all of the mode code functions, may be effectively emulated in a token passing network by means of the high level system design.

3.1.10 (Continued):

The large number of fault detection and recovery mechanisms called for in the HSRB standard ensures robust operation and fast healing to a variety of error and fault conditions. This provides the network survivability and reliable, real-time operation demanded for military platform applications, with recovery times on the order of a millisecond for station and media faults.

3.1.11 Token Passing Rings: Both ANSI FDDI and SAE AS4075 (HSRB) are high speed Token Passing Ring buses.

Both buses use a 4B5B encoding scheme. FDDI transmits at 125 Mb (100 Mbit/s data rate) while the HSRB standard does not limit the bit rate.

(The slower transmission speed of the initial 100 Mb fiber optic implementation was chosen in order that the necessary components, both fiber optic and crystals with sufficiently tight tolerance to meet the synchronization requirements, could be obtained to meet the wide temperature range experienced in some military systems).

To understand one of the major differences between the HSRB and FDDI one must first consider the other major Token Passing Ring standard, namely IEEE 802.5. The 802.5 standard is considerably slower (signaling rate of 1 or 4 Mb) than the HSRB or FDDI. However, many of the underlying principles of the Token Passing technique are exhibited by the standard.

A Token Passing Ring contains an attribute of "Ring Length". Ring Length is made up of the transmission delay on the media together with the delay through the registers within the ring stations that are in the direct path of the ring. The Ring Length of a system will appear to be longer if the transmission rate is higher as more bits can be accommodated within the ring if the bits are transmitted faster, therefore appearing to take up less room on the ring.

In the case of 802.5 the transmission rate is comparatively slow and a single message therefore tends to fill the ring. This fact allows an 802.5 station to use the mechanism of not releasing a Free Token until it receives its previously transmitted message header back. This allows the use of the priority protocol that is implemented in 802.5. A circulating message collects the priority of the other messages awaiting transmission on the ring (in the priority reservation field of the message header) and a Free Token can then be released at the highest priority of message that is awaiting transmission. This mechanism gives the 802.5 standard a true priority (highest priority goes first) scheme.

With increasing transmission speed the length of a ring becomes effectively longer. One message may well not fill the ring. If the 802.5 protocol were used on such a ring its effective bandwidth would be markedly reduced as it would spend a considerable amount of time transmitting padding idle symbols in order to fill up the ring while awaiting the return of message headers.

The simple answer to this problem is for a station not to wait for the return of its own message before transmitting a Free Token. In this way it is possible to have more than one message on the ring at the same time and greatly increase the amount of data being transmitted, reduce the amount of idle being transmitted and increase the effective bandwidth of the ring.

3.1.11 (Continued):

The price to be paid for such a scheme is that if a station does not wait for its own message header to return before issuing a Free Token then the issued Token cannot be allocated a priority.

Both FDDI and HSRB utilize schemes that allow multiple messages on the ring, giving the advantage of increased effective bandwidth. However, the different methods which the two standards use to overcome the problem of priority handling have a profound effect on the operation of the two standards within a system.

FDDI has overcome the problem by the adoption of Token Rotation Timers. Very simply when Free Tokens are rotating round the ring quickly the ring must be lightly loaded and therefore messages of low priority can be transmitted. When Free Tokens are rotating around the ring slowly the ring must be loaded and low priority messages are therefore not allowed to be transmitted.

This scheme does not offer a true priority scheme as there is no guarantee that high priority messages will always be transmitted before lower priority messages. The scheme is, in fact, more akin to bandwidth sharing with an attempt being made to allocate more bandwidth, and therefore better service, to the more important messages. FDDI therefore offers good throughput with bandwidth sharing and is a scheme that is well suited to its intended application, i.e., that of general purpose data processing.

The HSRB standard takes a different approach to handling the problem of utilizing the available bandwidth while still maintaining a priority scheme.

When messages are "longer" than the ring then a priority reservation scheme is used in the same way as in the IEEE 802.5 standard. This ensures a true priority scheme where the higher priority messages go before the lower priority messages. When messages are "shorter" than the ring then the priority scheme is suspended and messages are dealt with as they are encountered. However, in order to ensure that higher priority messages are not locked out by the presence of a large number of short messages the priority scheme is only suspended for a maximum of 16 short messages. After 16 short messages in a row a long message is forced, by padding the message with idle, causing the re-establishment of the priority scheme.

The number of short messages that can be sent in a row (sixteen) was not chosen randomly but was chosen such that a highest priority message would not experience considerably more delay when waiting for the short messages to complete than it might have had to experience while waiting for a maximum length message to complete. The other factor to be borne in mind is that if the system contains considerable number of short messages then these messages, which would have caused great inefficiency if dealt with under a priority scheme will be dealt with very quickly, and efficiently, under the scheme used without detriment to the priority scheme as applied to the other messages.

The HSRB was designed specifically for use in Real-Time systems. As such special attention was paid to making both the semantics and the protocol of the standard as efficient as possible.

3.1.11 (Continued):

The measures taken to make the semantics as efficient as possible can be seen by comparing the semantics of the HSRB with that of the FDDI standard. For instance the use of a 16-bit error detection code in the HSRB (which is quite adequate for the purpose) compared to the 32-bit one used in FDDI.

The above semantic differences, while having an effect on efficiency, cause a small effect compared with the bit delay allowed through a station in the two standards. The FDDI standard allows up to 94 bits delay in a station on the ring while the HSRB standard allows only 6 bits delay in all the stations on the ring except one master station that is allocated up to 40 bits. This difference has a profound effect on the Ring Rotation Times (RRT) of the two standards (as RRT is made up of transmission delay and total station bit delay).

To summarize the differences in the two standards, the HSRB offers a true priority scheme whereby higher priority messages go first. The FDDI offers a bandwidth sharing scheme which attempts to allocate sufficient bandwidth to high priority information.

Although both standards offer considerable, and efficient, throughput the HSRB standard imposes far less delay in the path of any particular message.

A comprehensive guide to the calculation of the performance to be expected of the HSRB is given in Section 7.

HSRB provides levels of performance and reliability eminently suitable for airborne, shipboard, vehicle and ground based applications. In addition to exhibiting superior throughput and latency performance in comparison to other networking standards, the HSRB offers a great number of fault detection and recovery mechanisms to support truly robust operation in a real-time environment. Since the HSRB standard specifies point-to-point transmission and a protocol that is essentially independent of both media and clock rate, it is inherently upward compatible. For example, it is practical to plan and implement a 100 MHz system today, with an eye towards making future upgrades to the gigahertz range.

3.2 Operational Requirements

- 3.2.1 **Serial Transmission:** Data in a communications system can be transferred in either a parallel or serial fashion. A parallel transfer system is characterized by a data bus (8, 16, or 32 bits wide) and multiple control and timing signals while a serial transfer system combines control and data on a single transmission medium. Although the parallel system is capable of achieving higher bandwidths, the serial system has the cost, weight and power advantages of requiring only one driver and receiver per station (two for a redundant system) and only one transmission media per link between stations. Moreover, message latency is governed primarily by the line rate in parallel systems and there is little benefit in adopting a parallel structure for real-time, general purpose communications. The HSRB is a serial data transfer system offering high signaling rates with either wire or fiber optic media and sufficient bandwidth for modern avionic systems.

- 3.2.2 **Word Length:** The HSRB defines a message format which incorporates a word length of 16 information data bits. This word definition is used to count the amount of data in a message and to generate the check word at the end. The actual data quantities transferred can be bits, bytes, 16-bit words, 32-bit words, or any other size desired. Data are assembled into 16-bit words only to satisfy the message format.
- 3.2.3 **ML:** The Message frame format of the HSRB allows for a Message frame containing from one to 4096 data words. Messages greater than 4096 data words must be segmented into Message frames of an allowable size for transmission.
- 3.2.4 **Electrical Isolation:** The signal path between any physical stations on the HSRB should not provide for a direct current carrying connection. A wire media must therefore be AC coupled. An electrical shield, should be continuous and will therefore constitute a direct connection between RIMs.
- 3.2.5 **Self-Test:** The standard requires that a station guarantees its health without interfering with the ring operation, before joining the ring. Since the self-test is implementation dependent, no specific self-test requirements are defined in the standard. The self-test may employ any techniques necessary to meet the requirement consistent with interoperability. After the station joins the ring following self-test, it shall be capable of performing the HSRB protocols. Guidance on self-test is provided in Section 8.

3.3 Performance Requirements:

- 3.3.1 **Station Delay:** Station delay is the only performance parameter of the RIU which is specified in the standard. A maximum logic delay is specified in order to allow calculations of some timer settings, but more importantly to maintain a deterministic high throughput and low latency.

Station delay affects the choice of coding used for fields of the Token frame and the FS field of a Message frame. Due to the small delay time specified, it is not possible to adequately process 4B5B symbols in the Token frame. Code bits are used to allow immediate bit processing of such information as the Priority, Token Status, and Reservation. Likewise, code bits are used for the FS field of the Message frame so that the appropriate bits may be changed "on-the-fly".

The station delay will directly affect the latency and throughput of messages; as station delay increases, latency will increase and throughput will decrease. Section 7 addresses the effects of station delay on the network performance.

3.4 Semantics:

A definition of the semantics for the HSRB is given in 3.4. Entities, symbols and Protocol Data Units (PDUs) are first defined, followed by the elements which make up the PDU: frames, subframes, fields and subfields. The definition of each PDU continues until the appropriate entity level is reached according to the scheme shown in Figure 3.4-1. Signals are continuously present in the ring, in the form of either the PDUs or Idle. In this way bit synchronization is maintained and a constant signal is available for activation of AGC circuits in optical systems.




FIGURE 3.4-1 - Scheme for PDU Definition

- 3.4.1 Entities: In any serial data transmission system it is necessary to provide a means of determining in a receiving station the bit boundaries of an incoming data stream. In Manchester encoding this is achieved by ensuring that a transition occurs at the centerpoint in time of each bit, but at a significant cost in terms of channel bandwidth. Since, as indicated above, a signal is always present on the ring, such a frequent transmission of checking information is not needed. As a compromise, a self-clocking scheme has been chosen for the HSRB which, at the cost of one additional bit for every four data bits, ensures that the number of bit times between transitions is a maximum of three.

In the HSRB standard the bits transmitted on the physical medium are termed signaling bits; these are generated by a NRZI encoding of a bit stream consisting of code bits and symbols, giving a transition for a logic 1 and no transition for logic 0.

- 3.4.2 Symbols: In order to minimize station delay, a parameter which has a significant effect on performance, it is necessary to manipulate certain parts of the data stream without decoding symbols into data bits and re-encoding. The standard was designed so as to allow this process to be implemented in hardware and carried out on-the-fly. For this reason code bits are defined. Code bits are always transmitted in groups of five with specified bits assigned to logic 1 to guarantee no more than three bit times between transitions, regardless of the state of the bits not assigned to logic 1 in each group.

- 3.4.3 PDUs: Three PDUs are used: Token frames, Message frames, and Beacon frames.

- 3.4.3.1 Token Frame: The Token frame is the means by which the right to transmit is transferred from one station to another. It exists only in the form of a free token, i.e., available to be claimed. Once claimed it becomes a Claimed Token subframe and forms part of a Message frame.

The Token frame consists of a TSD, a Control field (CON) and a Token Ending Delimiter. Within the standard the contents of the Token frame can be traced from the Token frame definition in 3.4.3.1 to the field and subfield definitions in 3.4.4.1 to 3.4.4.3. The Delimiters are defined in terms of symbols but, since these are fixed, they can also be viewed as groups of code bits. The CON field is defined in terms of code bits to allow rapid bit manipulation without the need for group decoding/encoding.

- 3.4.3.1.1 TSD Field: The TSD is defined in 3.4.4.1 and consists of a J symbol followed by a K symbol, which is a uniquely recognizable bit sequence even in the absence of a knowledge of symbol boundaries and can be used for the initialization of a state machine.
- 3.4.3.1.2 CON Field: The CON field is defined in 3.4.4.2 and its subfields in 3.4.4.2.1 through 3.4.4.2.5. When taken as a complete field of 20 code bits it can be seen that logic 1's have been inserted into the appropriate bit positions to generate the requisite number of transitions following NRZI encoding.

3.4.3.1.2 (Continued):

The function of the subfields within the CON fields are described in 3.5. However, two points are worthy of note here.

- a. The ordering of the subfields allows bit manipulation without access to all the subfields in parallel (the priority level of a token must be known before a station with a message to transmit can claim it, the token status, claimed or free, must be known before a station can make a reservation).
- b. The two token status bits, which should be identical, are separated by the SMC to provide some immunity to burst errors. In reality, however, the duplication of the Token Status bits provides little more than a warm feeling and a measure of error monitoring capability.

3.4.3.1.3 Token Frame Ending Delimiter Field: The Token Frame Ending Delimiter field is defined in 3.4.4.3 and consists of a T symbol. It can be used to provide confirmation that the token is free, as indicated by the Token Status bits.

3.4.3.2 Message Frame: The Message frame is used to transmit information. It consists of a Claimed Token subframe which was the free token claimed by the station transmitting the Message frame with the Token Ending Delimiter stripped off, followed by a Interframe Adjustment field (IFA), various address and message control fields, an INFO field and a FS field. Within the standard the contents of the Message frame can be traced from the Message frame definition in 3.4.3.2 to the field and subfield definitions in 3.4.4.4 through 3.4.4.17. All fields are defined in terms of symbols or data bits with the exception of the FS field which, since rapid bit manipulation is required, is defined in code bits.

The main content of the Message frame following the Claimed Token subframe, the IFA and the MFSD field is organized around 16-bit boundaries for ease of implementation. Thus, the first 16-bit word contains the Priority and Retry Status, and Word Count fields, the second 16-bit word contains the Sending Address, Address Control and GA fields, the third 16-bit word contains the DA field and so on. The word boundaries of the Message frame are shown in Figure 3.4.2.

3.4.3.2.1 IFA: The IFA is defined in 3.4.4.4. It consists of six Idle symbols and is provided for the purpose of synchronization.

3.4.3.2.2 MFSD Field: The MFSD field is defined in 3.4.4.5 and consists of a J symbol followed by an A symbol. It is therefore distinguishable from a TSD and is used to redefine the symbol boundaries following the variable length preamble between the Claimed Token subframe and the remainder of the Message frame.




FIGURE 3.4-2 - Main Part of Message Frame Shown in 16-bit Words

- 3.4.3.2.3 **Priority and Retry Status Field:** The Priority and Retry status field is defined in 3.4.4.6, and consists of the Priority subfield and the Retry Status subfield, defined in 3.4.4.6.1 and 3.4.4.6.2 respectively. The Priority subfield consists of three data bits indicating the priority of the message, and can be used to order incoming messages in a queue and to establish priority for messages passing through a bridge. It should be noted that the priority in the Claimed Token subframe may be lower than the priority of the message itself. The HSRB provides for a single retransmission of a message (using the next available token at the appropriate priority and deferring to general messages in the station at a higher priority) in the event of an information error to avoid the need to reinsert a retry message in the input queue. The Retry subfield contains a single code bit to indicate whether or not the current Message frame is a retry for station and network management purposes.
- 3.4.3.2.4 **Word Count Field:** The Word Count Field is defined in 3.4.4.7 and consists of three symbols to indicate the number of words to be expected in the INFO field. The Word Count could be used to prepare an input buffer for reception and to detect a transmission too short or too long. Since the minimum ML is one word, the value zero is used to indicate 4096 words.
- 3.4.3.2.5 **Sending Address Field:** The Sending Address field is defined in 3.4.4.8 and consists of two symbols. Since up to 128 physical addresses are supported, the most significant bit of this field is not required and is set to logic 0.
- 3.4.3.2.6 **Address Control Field (AC):** The AC field is defined in 3.4.4.9 and its subfields in 3.4.4.9.1 and 3.4.4.9.2. The Logical Physical (LP) subfield consists of two data bits, one of which is preassigned to logic 0, the other indicating whether the Global Address and Destination Address following is a logical address or a physical address.

The AWC subfield consists of two data bits and is used only in the logical addressing mode. In the physical addressing mode bits are assigned to logic 0.

In the logical addressing mode the AWC indicates the number of 16-bit address words following the first word. In this way 1, 2, 3, or 4 words of logical address can be transmitted.

It should be noted that the system designer must configure his system for compatibility between the length of address transmitted and the capability of a receiving station to decode it. A variable length address field has been provided to avoid the need for stations in smaller systems to decode an address of maximum length. The logical address zero has been reserved for broadcast. If the AWC is higher than the number of words capable of being received, the message will be repeated but not copied. The address transmitted always corresponds to the least significant bits; for example, a 16-bit address is the least significant 16 bits of a 64-bit address with the most significant 48 bits set to zero.

- 3.4.3.2.7 GA Field: The GA field is defined in 3.4.4.10 and its subfields in 3.4.4.10.1 and 3.4.4.10.2. In order to provide a nice format for the standard, but with the result that the reader may be totally confused, the GA field is defined in the logical and physical address modes as subfields. It is important to note that these subfields do not coexist; in the physical address mode, as determined by the logical physical bit in the AC field, the GA field becomes the GAP subfield and in the logical address mode the GA field becomes the GAL subfield.

The GA field supports the use of bridges for up to four rings as described in 4.2. The GAP subfield, defined in 3.4.4.10.1 consists of four data bits. The Bridge Access (BA) bit indicates that transfer across a bridge is required and the GAP bits indicate the destination ring number. A GAP of zero indicates that the sending station and the destination station are on the same ring.

The GAL subfield indicates the ring or rings for which the Message frame is intended. Since one GAL bit is provided for each ring a Message can be sent to any ring or combination of rings.

- 3.4.3.2.8 DA Field: The DA field in 3.4.4.11. Common with the GA field, the DA field is defined in terms of the DAP and the Destination Address Logical subfields (DAL) which do not coexist. The DA field becomes, in the physical mode, the DAP subfield, and in the logical mode, the DAL subfield.

The DAP subfield is defined in 3.4.4.11.1. It consists of four symbols, the lower nine bits of which are assigned to a subaddress and the upper seven bits which are assigned to the physical address. In this way up to 128 physical addresses with up to 512 subaddresses each are supported on each at 4 rings.

The DA Logical subfield is defined in 3.4.4.11.2 and consists of 4, 8, 12, or 16 symbols as prescribed by the value of the AWC. The most significant nonzero word is transmitted first.

- 3.4.3.2.9 MCFCS Field: The MCFCS field is defined in 3.4.4.12 and consists of four symbols. The purpose of this field is to allow error detection over all address and CON fields of the Message frame. A 16-bit sequence has been specified not only to provide considerable error detection capability for these important fields but also for commonality with the IFCS to allow the same circuitry to be used for all FCSs.

- 3.4.3.2.10 INFO Field: The INFO field is defined in 3.4.4.13 and contains between four and 16384 symbols (between one and 4096 words) in accordance with the value transmitted in the Word Count field. Information words (IW) are transmitted most significant bit first; since there is no store and forward operation in the protocol, however, words will be received in the order in which they are transmitted, wither most significant word first or vice versa. The INFO field is divided into 256 word groups separated by adjustment fields which, when originated, consists of 6 idle symbols followed by JA. The adjustment field is provided for the purpose of initialization of an elastic buffer.

- 3.4.3.2.11 IFCS Field: The IFCS field is defined in 3.4.4.14 and consists of four symbols. The purpose of this field is to allow error detection over the INFO field excluding the adjustment fields.

In any implementation where common hardware is used for the MCFCS and the IFCS the FCS register must be read and reinitialized at the start of the INFO field.

The IFCS does not cover the FS field; error detection is carried out before the FS field is received so that a change of Frame Status can be indicated as the message passes through a station.

- 3.4.3.2.12 Message Frame Ending Delimiter Field: The Message Frame Ending Delimiter field is defined in 3.4.4.15 and consists of a T symbol. It marks the end of the information and the beginning of the FS field.

- 3.4.3.2.13 FS Field: The FS field is defined in 3.4.4.16 and its subfields in 3.4.4.16.1 through 3.4.4.16.8. The FS field, as a whole, consists of 15 code bits with appropriate bits preassigned to logic 1 to provide the requisite number of transitions following NRZI encoding. Status bits are provided for indication of message control error, information error, message received, and message acknowledge. Each of these bits is repeated so that single bit corruption can be detected.

The Message Control Error Detected 1 and Message Control Error Detected 2 subfields are defined in 3.4.4.16.1 and 3.4.4.16.5 respectively of a four bit subfield, three bits being preassigned to logic 1, and a single bit subfield. The two subfields should have the same value, being set to logic 1 if there is a disparity between the transmitted MCFCS and the value calculated from the incoming data stream. The Message Acknowledge 1 (ACK1) and Message Acknowledged 2 (ACK2) subfields are defined in 3.4.4.16.2 and 3.4.4.16.6 respectively and consist, respectively of a single bit subfield and a two bit subfield, one of which is preassigned to logic 1. The two subfields should have the same value, being set to logic 0 if a Message frame has been received containing the station's address and the message could not be loaded into the station's buffer. The message acknowledge subfields are transmitted at logic 1 so that for a multicast message any station to which the message is addressed that cannot load the message can indicate a negative acknowledge to the originating station.

The Message Received 1 and Message Received 2 subfields are defined in 3.4.4.16.3 and 3.4.4.16.7 respectively. Each subfield consists of two bits, one of which is preassigned to logic 1. The two subfields should have the same value, being set to logic 1 if a Message frame has to be received containing the station's address, regardless of the presence of information errors.

The Information Error Detected 1 and Information Error Detected 2 subfields are defined in 3.4.4.16.4 and 3.4.4.16.8 respectively, and consist, respectively of a two bit subfield, one of which is preassigned to logic 1 and a single bit subfield. The two subfields should have the same value, being set to logic 1 if there is a disparity between the transmitted IFCS and the value calculated from the incoming data stream.

- 3.4.3.3 Beacon Frame: The Beacon frame is used to transmit control information during start up and reconfiguration. This information is used to configure a ring with the maximum number of operating stations on it. It is also used to determine which station will be the master station on the ring.

The Beacon frame consists of a BFSD field, a BCON field, a HKA field, a SC field, a BFCS field and a Beacon Frame Ending Delimiter field (BFED). Within this document, the contents of the Beacon frame can be traced from the Beacon frame definition in 3.4.3.3 to the field and subfield definitions in 3.4.4.18 to 3.4.4.22. The delimiters are defined as symbols and the remaining fields are defined as data bits.

- 3.4.3.3.1 BFSD Field: The BFSD field is defined in 3.4.4.17 and consists of a K symbol followed by a J symbol, which is uniquely recognizable sequence.
- 3.4.3.3.2 BCON Field: The BCON field is defined in 3.4.4.18 and its subfields in 3.4.4.18.1 through 3.4.4.18.2. The BCON field consists of a BT subfield and BPI subfield.

The BT subfield consists of three data bits which define the type of Beacon frame. There are seven types of beacons represented by binary 000 through 110. Binary 111 is reserved.

000 - Warm Start Beacon. Warm Start Beacons are initiated by Slave or Master Stations whose Loop Time Counter, Lost Token Starting Delimiter Counter or Lost Free Token Counter has expired and in the event of Token status errors. It is used to recover the ring by the Master Station issuing a new free token when reconfiguration of the ring may not be necessary.

001 - Warm Recover Beacon. Warm Recover Beacons are issued by a Master Station which has received a Warm Start Beacon. This is issued by the Master Station to the Slave Stations to return them to a normal operating state.

010 - Restart Beacon. A Restart Beacon is issued by a station to begin the reconfiguration process.

011 - Vie Beacon. A Vie Beacon is issued by a station after it initiates a Restart Beacon. The Vie Beacon is used to pass station addresses in the HKA field to determine which station will be the Master Station.

100 - Configure Ring 0 Beacon. A Configure Ring 0 Beacon is only issued by a station after it has determined it is the Master Station and Ring 0 is the desired ring. It is used by stations to determine the appropriate RIU states for the chosen configuration.

101 - Configure Ring 1 Beacon. Same as Configure Ring 0 beacon except for configuring on Ring 1.

110 - Configure Loopback Beacon. A Configure Loopback Beacon is issued by a station which has determined that it is the Master Station in a loopback ring.

111 - Reserved.

3.4.3.3.2 (Continued):

The BPI subfield consists of one data bit. It is used so that the receiving station knows whether the HKA data came exclusively from the received ring or if it reached the station after having gone through both rings.

3.4.3.3.3 HKA Field: The HKA field is defined in 3.4.4.19 and consists of eight data bits. The most significant bit is a binary zero and the seven least significant bits indicate the highest station address on the ring known to the transmitting station. It is used to determine which station will be the new Master Station on the ring. It is updated whenever a Vie Beacon is transmitted.

3.4.3.3.4 SC Field: The SC field is defined in 3.4.4.20 and consists of eight data bits. The most significant bit is a binary zero and the seven least significant bits represent the number of stations included in a complete ring and is updated when Vie Beacons are transmitted.

3.4.3.3.5 BFCS Field: The BFCS Field is defined in 3.4.4.21 and consists of 16 data bits. It is updated by the station transmitting the beacon and is used to check the validity of the BCON field.

3.4.3.3.6 BFED: The BFED is defined in 3.4.4.22 and consists of a T symbol.

3.4.4 Fields: In the standard, the frames are first defined and the field definitions follow in a hierarchical manner. For clarity the field descriptions in this Handbook follow on from the frame descriptions given above.

3.5 Protocols:

Within the standard, the protocols are presented in the two forms: A state machine definition and a verbal description. The state machine definition takes precedence over the verbal description, which is provided for ease of understanding. The following paragraphs are an explanation of the state machine definition.

3.5.1 Protocol State Definitions: The finite state definition of the HSRB protocol is given in 3.5.1. The states and transitions are defined at two levels: the primary states and transitions, as shown in Figure 3.5.1-1 of the standard, and the substates for each primary state, as shown in Figures 3.5.1.6-1, 3.5.1.7-1, and 3.5.1.10-1. The primary states are S1 Unconnected, S2 Reconfiguration, S3 Inactive Ring, S4 Repeater Ring and S5 Active Ring. The finite state definition serves to provide a formal definition of the standard; it does not imply a particular implementation.

In the hierarchical state definition provided an attempt has been made, for reasons of clarity, to limit the number of states per diagram and to maintain a consistent degree of abstraction (detail) within each level of the hierarchy.

Since it is possible to configure the HSRB as a simplex system, the finite state definition has been prepared in such a way that identical state machines can be applied to two RIUs in a dual redundant (counterrotating) HSRB. Interaction between the state machines is limited to the synchronization of certain transitions and a small number of status indications.

- 3.5.1.1 Unconnected State (S1): The S1 is defined in 3.5.1.1 and is entered when power is removed from the RIU, or when commanded to reset by the Host. Because S1 is the initial state on power up, there are transitions to this state from all other states on loss of power. The only transitions out of S1, from S1 to S2 (Reconfiguration), occurs when valid signals are detected from another station. The state S1 thus serves to transition the RIU from an unpowered or reset condition to a condition where it is synchronized with other stations and may attempt to enter the ring.
- 3.5.1.2 Reconfiguration State (S2): The S2 is defined in 3.5.1.2 and may be entered from S1 as described, or from any other primary state. A transition back to S1 occurs under two conditions: if power is lost (an effective transition), or if the RIU is reset by the Host. A self-loop occurs when state S2 is to be re-entered because the reconfiguration process cannot be completed successfully due to insufficient connectivity to its neighbors. This self-loop transition also allows for the synchronization of state machines of a counterrotating pair due to a newly powered station. A transition to S3 (Inactive Ring) occurs if at the end of reconfiguration, a complete ring has been achieved and the state machine in the counterrotating pair is not on the ring to be used for data transmission. A transition to S4 Repeater Ring occurs if the system has been configured as a looped back ring and the state machine in question is in the side of a station in a loopback path rather than in the actively monitored path. The side of a station in the loopback path merely repeats information from its input to its output.

A transition to S5 (Active Ring) occurs if, at the end of reconfiguration, the system has been configured such that the state machine in question is on the side of a station in the main data transmission path, i.e., designated to receive and transmit messages. The state machine may enter S5 either as the Master Station or as a Slave Station, this part of protocol being defined by the substates for S2.

- 3.5.1.3 Inactive Ring State (S3): The S3 is defined in 3.5.1.3 and is entered only from state S2 at the end of reconfiguration if the RIU configures as part of a ring system and this is the Inactive ring. Data received on the Active Ring is transmitted on the Inactive ring output. Data received on the Inactive ring is monitored but not repeated. The only valid transitions from S3 are to either state S1 or S2.

A transition to S1 occurs under two conditions: if power is lost (an effective transition), or if the RIU is reset by the Host. A transition back to state S2 occurs under four conditions: if the Host commands the ring system to reconfigure, if a valid Restart Beacon has been received on the Inactive ring input, if there has been a change in the line state of the Inactive ring input, or if the state machine defining the operation of the Active ring indicates a transition to state S2. The state machine defining the operation of the Active ring will also enter S2 for the cases where a valid Restart Beacon has been received or a line state change has been detected on the Inactive ring input.

- 3.5.1.4 Repeater Ring State (S4): The S4 is defined in 3.5.1.4 and is entered when the state machine defining the operation on the counterrotating ring is in the Active Ring state (S5) and the station is in the middle of a looped back ring. Data received on the repeater ring input is transmitted on the repeater ring output without being made available to the Host. State S4 is entered only from state S2 at the end of reconfiguration if the RIU configures as part of a looped back ring system and this is the Repeater Ring. The only valid transitions from S4 are to either state S1 or S2.

A transition to S1 occurs under two conditions: if power is lost (an effective transition), or if the RIU is reset by the Host. A transition back to state S2 occurs under four conditions: if the Host commands the ring system to reconfigure, if a valid Restart Beacon has been received on the Repeater ring input, if there has been a change in the line state of the Repeater ring input, or if the state machine defining the operation of the Active ring indicates a transition to state S2. The state machine defining the operation of the Active ring will also enter S2 for the cases where a valid Restart Beacon has been received or a line state change has been detected on the Repeater ring input.

- 3.5.1.5 Active Ring State (S5): The S5 is defined in 3.5.1.5 and is entered when the ring has been designated the ring on which data are to be received. State S5 is entered only from state S2 at the end of reconfiguration if the RIU configures as part of a ring system and this is the Active ring. The only valid transitions from S5 are to either state S1 or S2. When in the Active ring state, the state machine defining the operation of the counterrotating ring, if it is present, will be in either state S3 or S4, depending on the ring configuration.

A transition to S1 occurs under two conditions: if power is lost (an effective transition), or if the RIU is reset by the Host. A transition back to state S2 occurs under five conditions: if the Host commands the ring system to reconfigure, if a valid Restart Beacon has been received on the Active ring input, if there has been a change in the line state of the Active ring input, if the state machine defining the operation of the counterrotating ring indicates a transition to state S2, or if the warm start protocol fails. The state machine defining the operation of the counterrotating ring will also enter S2 for the cases where a valid Restart Beacon has been received or a line state change has been detected.

The operational state of all possible state machine pairs and the associated active ring path for a counterrotating ring system is shown in Figure 3.5-1. The solid line indicates the path of the configured ring. The broken line indicates a path of data transfer which is not part of the configured ring. For example, in Figure 3.5-1(a) the configured ring path passes through the Active RIU state machine only. Data, however, are also looped back to the Inactive RIU state machine and passed back to the station's left neighbor. That station (assuming Ring 0 or 1 is configured) will also be in the configuration as indicated in H3.5-1(a) and will ignore data entering its Inactive ring input (from the right side).

These apparently unused paths exist for two reasons: to maintain synchronization between all stations, and to simplify the reconfiguration process by using the same station configuration for more than one ring configuration. An example of this second reason is seen by examining Figures H3.5-1(a) and (c). The same station configuration is used for two different ring configurations.




FIGURE 3.5-1 - Major State Configurations for a Station in a Counterrotating Ring System

3.5.1.5 (Continued):

If a complete physical ring (0 or 1) is good, then all stations should be configured as shown in Figure 3.5-1(a). If a loopback ring is formed, the stations at the right and left end will be configured as shown in Figures H3.5-1(c) and (d) respectively. The stations in the middle of the loopback ring will be configured as shown in Figure 3.5-1(b).

3.5.1.6 Unconnected State Substates: The function of the Unconnected State is to transition the RIU from an unpowered or reset condition to a condition where it is synchronized with other stations and may attempt to enter the ring (reconfiguration). Under normal power-up, the RIU will enter the following substates in succession before transitioning to S2: S1.1, S1.2, S1.4, S1.5. If the self-test fails the RIU will remain in S1.3.

3.5.1.6.1 S1.1 Unpowered: The RIU is in this substate when no power is applied to the RIU. One RIU of a station may transition to S1.2 Powered even if the other side of the station (RIU) remains unpowered.

3.5.1.6.2 S1.2 Powered: The RIU must be bypassed when in this substate because it is not yet ready or authorized to enter the ring. Failure of the state machines is detected by use of self-test in this substate. If failure of the opposite ring's state machine is detected, the working state machine will ignore the failed state machine.

3.5.1.6.3 S1.3 Quiescent: This substate is entered by an RIU if self-test is not passed. The other RIU (if present) will operate normally if it passes self-test. This substate can only be exited by a reset by host or upon power loss.

3.5.1.6.4 S1.4 Wait: Idle symbols are transmitted when in this substate to allow downstream receivers to synchronize to this station, after the RIM is engaged. The wait substate is included to allow time for the RIM to engage.

3.5.1.6.5 S1.5 Synchronization: Synchronization should be achieved within 1024 symbol times with clock recovery and synchronization tolerances as specified in the HSRB document. Once this station is able to receive valid symbols from a neighbor it will enter S2 reconfiguration.

3.5.1.7 Reconfiguration State Substates: A station will enter the Reconfiguration state after exiting the Unconnected state in order to configure the ring(s) and begin normal operation. Under normal conditions, stations in this substate will enter state S2.1, S2.2 and then either S2.3 or S2.4 before going to one of the normal operating states S3, S4, or S5. The station will go to S2.3, from S2.2, if it is determined to have the HKA and be the master station, else the station will go to substate S2.4. Under conditions of insufficient link connectivity, the station will go to substate S2.5.

3.5.1.7.1 S2.1 Reconfiguration Initialization: When Reconfiguration is entered, there is no knowledge of ring conditions, synchronizations, or timing between stations. This substate is entered to initialize the ring by resetting the BLT and transmitting Restart Beacons. The Restart Beacons serve to notify all stations that initialization is to begin. All stations will thus be initialized within one BLT. Restart Beacons are ignored during the first BLT of reconfiguration in order to prevent continuous restarting of stations.

- 3.5.1.7.2 S2.2 Vie: Stations in the Vie substate will continuously transmit Vie Beacons having their most recent HKA and BPI information. By noting the HKA and the BPI information from the beacon in which the HKA has been received, stations can determine if the HKA station can communicate to this station over a single ring alone. The best ring configuration can then be determined distributively in each station. The station with the HKA (Master Station) will know if a ring is good by receiving beacons with its own Beacon address and with the BPI bit set to 1, on that ring. Otherwise a loopback ring must be formed. In this case each station will independently determine if it is at the end of the loopback ring by knowing on which ring the HKA was received, the relation of the received HKA to the stations HKA and the Receiver Status of each input. The Vie process ends after 2 BLT's have expired. This allows time for all stations information to reach all other stations.

Vie Beacons are transmitted continuously while in the Vie substate (S2.2) every 16 symbol times. The Vie Beacons length is 12 symbols; this leaves 4 symbols of IDLE between beacon transmissions. IDLES are transmitted to maintain bit synchronization. The idle interval between beacons is used to allow nonsynchronized transmission of beacons and ease hardware timing requirements. Specifying the time for propagation of updated beacon information through a station allows the worst case ring delay for the BPI and HKA (the BLT) to be calculated. 36 symbol times, as specified, is a reasonably implemented delay.

- 3.5.1.7.3 S2.3 Master Configuration: A station which, after the Vie process has ended, has not received an address higher than its own station address will become the Master Station. This station will determine the best ring configuration, send a Configure Beacon, and notify Slave Stations of the ring configuration; ring 0 good, ring 1 good or loopback ring. The Configure Beacon is followed by one BLT of IDLE symbols to allow stations to switch from the 36 symbol delay during reconfiguration to the 4 bit delay during normal operation. The free Token will be issued by the active state machine of the Master Station. The station will then go to normal operation.
- 3.5.1.7.4 S2.4 Slave Configuration: A station which is not the Master Station will enter this substate when a Configure Beacon is received. The Configure Beacon indicates which state (S3, S4 or S5) each side of the station should transition to. The configure rules defined in Table 3.5.15-2 ensure that only one active side of a station exists, and that the other side is appropriately configured.
- 3.5.1.7.5 S2.5 Reconfiguration Idle: There are instances in which a station can receive valid signals from its neighbor(s) but should not be part of the configured ring. Thus may happen, for instance, when the station is bypassed (due to bypass failure) on the only complete ring or when the station (because of link failures) cannot join a loopback ring. In these cases the station should not retry reconfiguration else it will continue to restart a ring which it cannot join. Thus, this substate specifies that only IDLE symbols be transmitted while waiting for an indication, such as line state change, receipt of a Restart Beacon, or Host command, that it may try to restart the ring.

- 3.5.1.8 Inactive Ring State Substates: A station side in the Inactive ring state will be in a station at the end of a loopback ring or in a station on an unlooped ring with the other ring active. The Inactive ring state side will be receiving data on a link which is not part of the (logical) ring. The Inactive ring side transmits data from the active ring in order to complete the loopback function. The Inactive ring state will also be entered as the opposite side from an active station on a complete ring of a counterrotating pair. In this case active data can only be transmitted on one ring even if both rings are good.
- 3.5.1.9 Repeater Ring State Substates: The Repeater ring state is entered by one state machine of a station that is in the middle of a looped back ring. The other station will be in the Active state. In this manner exactly one active connection to a looped back ring will exist. The repeater side (state machine) of the station serves to complete the loopback media connection only. It may not participate in normal protocols. The active side has been arbitrarily chosen to be Ring 0.
- 3.5.1.10 Active Ring State Substates: The active ring side of a station is used to receive and transmit messages and perform normal protocol operations. A station which may not, or does not wish to, send a message and receives a free token will enter the following substates under normal operation; S5.1, S5.2, S5.3 to S5.1. A station which receives a free token and sends its message will enter the following substates normally; S5.1, S5.2, S5.5, S5.6, S5.7, S5.8, to S5.1. A station receiving a claimed token will enter the following substates normally; S5.1, S5.2, S5.3, S5.4 to S5.1. Substate S5.9 is entered only when a BFSD is received and will usually transition to S2, Reconfiguration. Substate S5.10 is the error recovery state, Warm Start.
- 3.5.1.10.1 S5.1 Await TSD: An RIU is in state S5.1 when it is not originating a message and is awaiting receipt of a TSD. Upon receipt of a valid TSD the RIU will transition to state S5.2.
- 3.5.1.10.2 S5.2 Check Token Priority: In this state the RIU reinitializes the token counters and evaluates the Token Frame Priority field. If the RIU doesn't have a message to transmit, or if the station's message is of lower priority than that contained in the Token's priority field, or if the Token has been previously claimed, the RIU transitions to state S5.3, the Reserve Token Priority state. If the RIU has a message to transmit that is of an equal or higher priority level than that in the Token Priority field within the received free Token, the RIU will transition to state S5.5, the Token Claim state.
- 3.5.1.10.3 S5.3 Reserve Token Priority: When an RIU determines that it can't claim the received Token, it will enter this state upon exiting from state S5.2, the Check Token Priority state. If this RIU has a message to transmit and the Token has been previously claimed, it will change the Reservation field in the claimed Token to its priority level if it is of higher priority than that presently residing within the Reservation field. This Reservation field manipulation is required to enable the highest priority message to reserve the next available free Token. The dynamics of the Reservation field manipulation allows the Reservation field modification be done on a bit-by-bit basis to assure minimal impact on the station delay. This state also provides for verification of the Token format and for monitoring the logical agreement of the Token Status subfields, TS1 and TS2.

- 3.5.1.10.4 S5.4 Repeat Message: An RIU, after receiving a Claimed Token, will transition to this state to repeat a message, to check for errors in the message, or to receive the INFO field when it is determined that the DA field in the message header matches the station address and that the message is error free. While repeating the message, the RIU monitors and verifies that the protocol and error check bits are correct and that the Message frame is valid. FCS calculations are also performed and compared with the received FCS fields to check for Message frame errors. If Message frame errors are detected the appropriate bits are set in the FS field. Similarly, message receipt and acknowledge are indicated by setting the appropriate bits in the Frame Status field. After repeating a message the station transitions to state S5.1. Await TSD. The RIU will transition to a different state upon direction from the Host, loss of power, detection of a line state change detection of a Restart Beacon, or detection of an incorrect MFED or MFSD.
- 3.5.1.10.5 S5.5 Token Claim: When the RIU has a message to transmit and is ready to claim a free Token it will transition to this state upon receiving a free Token, provided that the priority of the message is greater than or equal to the priority of the Token. The TS1 and TS2 bits will be set to be claimed, the SMC value is saved and the Reservation field set to ones. A IFA, which is used to facilitate synchronization, is now inserted after the transmitted Claimed Token. For normal operation the RIU transitions to state S5.6, Originate Message.
- 3.5.1.10.6 S5.6 Originate Message: This state is entered after the RIU has claimed the Token. All received symbols are stripped from the ring and the properly framed and formatted Claimed Token and message is transmitted. The value in the Reservation field of the returning Claimed Token is extracted and made available to the RIU. Upon receiving a TSD the RIU checks the returning Claimed Token and Message frame for proper format, FCS, station address, etc., for errors. Upon completion of the MT, the station transitions to state S5.7, Issue Token.
- 3.5.1.10.7 S5.7 Issue Token: This state is entered after completion of a MT and controls the issuance of a free Token. The previously transmitted message may be either of two types, a long message or a short message. A long message has occurred if the returning Claimed Token arrives before MT is complete. Otherwise a short message has occurred. If a long message has been transmitted, a free Token is appended to the transmitted message with a priority level equal to the Reservation field level in the returned Claimed Token. Also, the SMC field is set to zero. If this station has transmitted a short message with the multiple message option not invoked or if the multiple message option was invoked but the SMC is at its maximum value (15), this station issues IDLES until it's Claimed Token is received. When the Claimed Token returns, a free Token with a priority level set to the value in the returning Reservation field and the SMC set to zero, will be transmitted. For a short message with the multiple message option invoked and the SMC less than its maximum value, this station will issue a free Token immediately upon completion of its MT. The SMC field will be incremented and the priority set at its lowest value. This will give low priority messages (at other stations) an opportunity to be transmitted. After the free Token is issued, this station transitions to state S5.8, Strip Message.

- 3.5.1.10.8 S5.8 Strip Message: A station enters this state, after issuing a free Token, to complete the removal of its message from the ring. As the station removes the message the format and check bits are verified. If the message is determined to have been transferred in error and it is to be retried a retry sequence is initiated. The message state is always made available to the Host, regardless of message error status. Upon completion of message stripping and if no errors have been detected, the station transitions to state S5.1 Await Token Starting Delimiter.
- 3.5.1.10.9 S5.9 Check Beacon: This state is entered when a valid BFSD is detected. If it is determined that a Restart Beacon has been received, the station transitions to state S2, Reconfiguration. If it is determined that a Warm Start or a Warm Recover Beacon was received, the station transitions to state S5.10, Warm Start. If the received beacon is invalid or is a Vie or Configure Beacon, the station will also transition to S5.10, Warm Start, since these conditions are abnormal. Receipt of a Restart Beacon on the Inactive Ring causes the Active Ring to enter reconfiguration.
- 3.5.1.10.10 S5.10 Warm Start: This state is entered upon receipt of a Beacon other than a Restart Beacon, or upon expiration of the Loop Time Counter or the Token counters. All data, if present, is flushed from the ring and all stations will reinitialize their respective Loop Time, Lost Token Starting Delimiter, and Lost Free Token Counters. After completion of this process the station that has been previously designated as the Master Station issues a free Token to resume normal operation. If the Warm Start procedure fails, the Reconfiguration state is initiated. Typically, a Warm Start failure is indicated when the Warm Recover Beacon is not received within two loop times.

4. APPLICATION ISSUES:

4.1 Topology:

Issues which may arise in the application of a ring topology are addressed in the following paragraphs. The term topology should not be confused with architecture. Topology is strictly the network configuration of terminals or stations connected by a transfer medium. Architecture at the system level defines the logical connection of message generating and consuming entities.

- 4.1.1 Ring Redundancy: While the HSRB protocol has been designed such that simplex systems can be implemented directly, two specific redundancy mechanisms are incorporated. A bypass mechanism is defined as part of each slash sheet to allow failed or off-line stations to be excluded from the ring. In addition, counterrotating rings are supported directly by the protocol, offering both static and dynamic redundancy. Static redundancy arises simply as a result of the presence of a complete alternate path, and use of this path is the primary mode for the recovery from faults. Dynamic redundancy is achieved by using operational elements from both rings in a loopback configuration which forms a single data path from the two paths available. This feature, which is particularly useful in recovering from common mode media failure, is fully automatic.

4.1.1 (Continued):

Should further redundancy be required it can take on two possible forms:

- a. Modular redundancy
- b. Additional links

While further modular redundancy (i.e., the replication of complete HSRB systems) is not supported directly by the HSRB protocol, there are no features of the protocol which preclude it. This was a conscious decision during the development of the HSRB protocol, made on the basis that applications with very high reliability requirements would need to demonstrate segregation of redundant elements and coupling of these elements at the protocol level would be counterproductive. Higher levels of modular redundancy would, therefore, involve independent counterrotating rings, the host (or hosts) handling the configuration aspects (standby redundancy or quadruple redundancy, etc.).

- 4.1.2 Topology Design: It is clear that the use of additional modular redundancy as discussed in 4.1.1 is under the control of the system designer and not strictly an issue for the data communications system itself. Design of a data communications topology using the HSRB is concerned principally with the configuration and powering of stations and the use of additional redundant links.

If the bypass specified in the particular slash sheet called up by the system specification has a restriction on the number of bypassed stations which can exist in succession, it may be necessary to configure the power supplies to stations with such a restriction in mind. Alternatively, additional links can be incorporated which bypass complete sections of a ring. This approach is particularly useful where the HSRB is used for communications within a cabinet and on a system wide basis.

- 4.1.2.1 Group Bypass Issues: Group bypassing can be defined as a bypass mechanism in which more than one terminal is bypassed by a single bypass unit. Such a bypass could be needed for partial powerup situations, to increase reliability, or to superimpose bypass options reflecting the natural structure of the network (bypassing an entire wing or pod).

There are several issues to be considered with group bypassing:

- a. The control of the group bypass, since it is no longer associated with a single terminal.
- b. In the case of bleed-off bypasses, the power budget.
- c. The algorithm used to invoke the bypass in general and in particular when members of the group begin to fail.

The implementation of group bypasses is a subject for individual system design.

4.2 Protocol:

The following paragraphs indicate some of the system level requirements that influenced the development of the protocols for the HSRB standard. Also addressed are system design issues related to the HSRB protocol which allow the system designer to implement efficient designs taking full advantage of the HSRB protocol.

- 4.2.1 General System Considerations: The HSRB is intended to support the inter-process communication needs for integrated fully distributed real-time control applications, that is, both inter and intra-cabinet communications (see AIR4271 Handbook of System Data Communications). This concept translates into a set of facilities in the media access protocol that supports the distribution of execution control, for instance task initiation and remote interrupts, and the transmission of data on the same communications medium. In general, the distribution of execution control imposes stricter requirements on the media access protocol than do the data communications. These execution control requirements include low message latency, low periodic traffic jitter, accuracy, and order of arrival.

The low message latency requirement is self-evident; any scheduled data transfer must be initiated and completed within the time allowable.

Periodic traffic on any shared medium cannot be guaranteed to be exactly periodic; the jitter addressed here is the variance in periodicity. This variance defines the degree of precision required for time distribution, process execution control, and sensor/effector data integrity. Sensor data integrity, for example, requires that in order to obtain coherent time series from a set of sensors that are physically distributed, the sensors must be sampled at the same time.

Accuracy is the differential reception time. It is the time delta between reception at the first and last destinations of a message transmitted simultaneously to more than one destination.

Allowable message latency, periodic jitter, and accuracy are application and system design dependent. For data driven distributed applications, guaranteed message order of arrival is essential for subsets of message categories. Other distributed applications require guaranteed message order of arrival during degraded system states for anomaly handling when making best effort decisions on partial data sets.

Fault tolerance and automatic recovery were additional system level requirements that influenced the development of the HSRB protocols, specifically the need for speedy reconfiguration algorithms.

- 4.2.2 Priority: Conceptually, a communications system can be viewed as a system resource that in a real-time application requires scheduling not unlike a processor. Like a processor, this system resource may experience transient overloads. The access protocol to the communications medium must provide a method that allows deterministic media scheduling in order to enable deadlines to be met while providing for the sharing of the system resource.

4.2.2 (Continued):

The HSRB access protocol provides message based priority. Priority access is implemented through priority reservation. The priority operation is a mechanism that allows the system designer to schedule the ring traffic on a per message basis. This mechanism provides for flexibility in the selection of specific scheduling algorithm. When assigning priority to message categories, the system designer must not confuse priority with the level of importance of a message or service category. Priority assignment is strictly a means for ring scheduling and message latency control.

A system design may utilize a periodic control message broadcast (or multicast) to sample a distributed sensor set. If the sampling frequency is high, in relation to the change of data between samples, an occasional loss of this message will not severely affect system operation; the individual message is not 'important'. However, it is essential that it be delivered with little jitter in the period, that is, with low variance in the periodicity particularly if subsequent processing operations are dependent on its arrival.

Periodic traffic jitter for the HSRB can be controlled by priority assignment. It is also a function of the maximum allowed ML in the system design, or the specific configuration ring length, whichever is larger. For example, if the highest priority is only allocated to the time distribution task, different levels of precision can be achieved by controlling the ML as illustrated in Table 4.2.2 for an 80 Mb/s HSRB.

TABLE 4.2.2 - Periodic Traffic Jitter for Highest Priority Message
at 80 Mb/s as Function of Maximum Message Size

Max Msg Length at 80 Mb/s Number of Bytes	Max Synchronization Jitter Priority Order Selected	Max Synchronization Jitter Short Message Option
2048	0.40 ms	0.45 ms
1024	0.20 ms	0.35 ms
512	0.10 ms	0.30 ms
32	0.03 ms	0.25 ms

The second set of data is when priority order is not selected and multiple short messages are allowed. If there is a second periodic message that is assigned next highest priority, the jitter for this second message will be roughly one and a half times the values of the first. If they are both of the highest priority, the jitter for both will be equal to one and a half times the values in the table.

Jitter for the highest priority message when Priority Order is selected, and there are no other messages pending of this priority, can be approximated to be equal to two times the maximum ML allowed, if this is longer than or equal to the ring length. Else, it is equal to two times the ring length. When the Short Message Option is used, jitter for the highest priority message is either equal to two times the maximum ML allowed, or equal to one maximum message allowed plus sixteen times the RRT, whichever is larger.

4.2.2 (Continued):

The above calculations were made assuming a RRT of 15.3 Fs, which roughly corresponds to a system with 90 stations on a 2 km ring. Note that for the 32 byte value and priority order selected, ring length rather than message size drove the upper bound on the jitter. The numbers shown are only intended as an example. Please refer to the calculations defined in Chapter 7 for specific system configurations.

The example shows how the system designer can use the priority for deterministic scheduling of the traffic on the ring. When assigning message priorities it is essential that the method used is an integrated approach to the total system scheduling. Scheduling of task execution will require data with different periods and deadlines. Message priority assignment must be accomplished in light of the specific task scheduling algorithm used in the system design; for example, rate monotonic or dynamic deadline scheduling.

4.2.3 Latency Control: The message based priority mechanism implemented in the HSRB not only allows the system designer to schedule the traffic on the ring but also to control message latency. Latency here is defined as the sum of the time it takes for a message to gain access and the actual transmission time of the message on the ring; upper layer protocols are not considered.

In the following explanations it is assumed that the Short Message Option is not exercised. Access to the ring for a pending message depends on the number of messages pending of equal or higher priority. If there is no message pending of equal or higher priority the maximum time it takes for the message to be scheduled and gain access is dependent on the length of the messages being transmitted. The worst case access occurs when the message arrives to be scheduled when the priority reservation field for the next free token has just passed by the station, forcing it to reserve the following free token instead. This means the upper limit of best case access (no equal or high priority pending) is the time it takes for the transmission of two maximum length messages or two RRT, whichever is larger. If there are messages pending of higher priority, they will all be transmitted before the one in question. This will increase the latency and delay the access for the message in question by the transmission time for the higher priority messages. Messages pending of equal priority may or may not be transmitted before the one in question depending on the physical locations of the stations on the ring. When more than one message is pending of the same priority, access reverts to a round-robin; i.e., they will gain access one after the other in the order they are physically connected on the ring.

Deterministic message latency control can be achieved by the assignment of priorities on a system level. For a given message the latency is equal to the sum of the transmission and the access times. Maximum access time is equal to the sum of the transmission time for all messages pending of higher and equal priority plus the transmission time for two maximum length messages. Best case access is the transmission time of just the higher priority messages. Higher degrees of precision in latency control can be achieved by limiting the maximum length of messages allowed on a system level basis.

4.2.3 (Continued):

When the Multiple Short Message Option is chosen, the limits of latency control are not those described. The upper limit of access when no messages of higher or equal priority are pending is the transmission time for one maximum length message plus sixteen token rotation times. (See Section 7).

4.2.4 Short Message Option: The HSRB access protocol (message based priority reservation) insures strict priority order; higher priority messages are always transmitted before those of lower priority. In addition, a modification to this access protocol based on multiple concurrent short messages is provided as a dynamically selectable design option.

When a message is short in relationship to the ring circumference, transmission may be complete before the beginning of the message (with the priority reservation in the claimed token) returns to the originating station for issuance of the free token. In this case waiting for the return of the claimed token before issuance of a free token will waste available bandwidth. Therefore an option is provided to the strict priority based protocol allowing the simultaneous presence of more than one short message. Selecting this option increases the throughput for short messages at the cost of priority order and an increase in latency (see Section 7). As a compromise between throughput and latency only sixteen consecutive short messages are permitted before priority order is enforced again. Use of the short multiple message option does not affect priority operation for long messages.

Some applications require guaranteed message order of arrival in all system states and others only during certain system states. Therefore, the multiple short message protocol was made a dynamically selectable option. In this manner, access through priority reservation can be preserved when the system operation requires strict priority order.

Short message operation is defined as an option. However, implementers are strongly recommended to provide the short message operation facility as a user selectable option in order that their implementations may be used in any system as required during system design. Disabling of the short message option can be achieved at the state machine level by continuously setting the SMC to 15. This forces any message sent by this station to be a long message and utilizes the returning reserved priority level. It is that the enabling or disabling of the short message option be fixed at system design time. Although stations with and without the option can coexist on one ring, if it is required to preserve priority operation for every message, the short message option must not be enabled in any station on the ring.

4.2.5 Message Retry: This document provides for a single retransmission of a message in the event of a message failure. This message failure detection and retransmission is performed at the HSRB protocol level to avoid the need to notify the host of the error and have it resend the message to the RIU. This saves host processing time as well as minimizing the latency of retransmission.

4.2.5 (Continued):

Messages may be designated for retry by the host on a per message basis. For example, the system designer may choose to select control oriented messages to have automatic retry. These messages are generally aperiodic in nature and often time critical. The additional delay that would be incurred by notifying the host to resend the message could result in a system fault. Other types of messages, such as data inputs from various sensors, tend to be periodic, providing data values sampled over time. The system designer may choose not to select automatic retry for these types of messages. The rationale is that the loss of any one message is not critical; the missed sample will appear as a brief instance of noise, which generally will be filtered out by the system, and another sample will soon follow.

Messages that are designated by the host for retry on the occurrence of a message failure will be retried at the same priority level as the original message. It will become the first message of its priority to be transmitted. As indicated in this document, the retry message should not pre-empt a queued message of a higher priority.

4.2.6 Addressing Modes: Selection of the ring station or stations is provided by the DA field. There are two addressing modes provided which are mutually exclusive within a message. A separate single bit field, the Logical-Physical (LP) subfield, is used to select the addressing mode. A Physical Addressing Mode is selected when the LP subfield is "0". To select the Logical Addressing mode set the LP subfield to "1".

The Physical Addressing of 2^{**7} station addresses with 2^{**9} station subaddresses is provided in the DAP subfield. Selection of the ring station is provided by bits 9 through 15, inclusively, of the DAP subfield. When used for Physical Addressing these bits are referenced as DP9 to DP15. DP15 is the most significant bit of this subfield. Bits 0 through 8, inclusively, of the DAP are used for subaddress selection of the destination address. These bits are denoted as DPO to DP8 with DP8 as the most significant bit. These two subfields enable the selection of up to 128 physical addresses, each with up to 512 subaddresses.

The Logical Addressing of messages, identified by message labels or message content IDs, has provision for extension. The extent of the logical address is controlled by the AWC 2 bit subfield. These two bits, AWC1 and AWC0, are used to signal a logical addressing space of:

- a. AWC1, AWC0 = 00: 2^{**16} logical address space
- b. AWC1, AWC0 = 01: 2^{**32} logical address space
- c. AWC1, AWC0 = 10: 2^{**48} logical address space
- d. AWC1, AWC0 = 11: 2^{**64} logical address space.

To provide for this range of logical address space the DA field is allowed to be of one to four words in length. The implied DA field word length directly corresponds with the logical addressing space of the AWC subfield. For physical addressing the AWC field must be set to "00".

4.2.6 (Continued):

There are several reasons for provided the capability to extend to a rather large logical address space. Designers of smaller systems, e.g., single ring systems with few dozen stations, may find a 2^{16} logical address space adequate. However, the HSRB can support up to 128 stations per ring on each of four rings, for a total of 512 stations. Some system designers provided input on applications where there would be many software tasks per station (50 - 100), with each task generating numerous messages (10-20).

This translates into approximately $512 * 100 * 20$ or 1024000 messages, requiring approximately 20 bits of logical address. This leads to two logical address words which provides a 2^{32} logical address space. Along with this, interest was expressed to not preclude future standardization of the logical address space with respect to the content labels for messages. It was envisioned that some form of segmentation of the logical address space would be required, e.g., to provide subfields for country codes, branch of service, platform, system, subsystem etc. This could easily require 15 to 20 additional bits. This leads to three logical address words which provides a 2^{48} logical address space.

While not specifically called out as an option, e.g., like the Short Message Option described in 4.2.4, the selection of physical or logical address mode and size of logical address space is a system design decision. It is important to note, however, the HSRB standard does not provide the system designer with the option or capability to directly set the logical address tables over the ring. For reasons of integrity, the station host, which is presumed to have some intelligence, must be involved in setting up or changing the logical address tables.

- 4.2.7 Loopback: In a dual ring HSRB network each RIU provides for a data path from the primary ring to the secondary ring. In the event of faults on both rings the dual ring HSRB will route data from one ring to the other ring, through the RIU in those stations that are determined to be at the ends of the loopback ring. Stations in the middle of the reformed loopback ring operate with one RIU state machine in the active ring state and the other in the repeater state. That is, data received at a ring input to the RIM is retransmitted at the corresponding ring output, but only the active RIU state machine of the station passes data with the host. Refer to 3.5.1.5 and Figure 3.5-1 for more details on the loopback protocol. This procedure enables the reformed loopback ring to exclude media faults on both the primary and secondary ring and thereby provide fault tolerance to multiple fault media failures. Multiple fault media failures may result in the formation of one or more subrings.

4.3 Bridges:

Bridges are used to pass messages between two identical types of networks using a store and forward technique. The bridge acts as the destination station on the ring which originates the message. The bridge, if it recognizes a DA for the other network to which it is connected, copies the message from the first network and completes the protocol. Once the message is completely copied, the bridge will retransmit it on the other network at the first available opportunity. On this network, the bridge will be the originating station.

4.3 (Continued):

Different techniques can be used in designing bridges and network systems containing bridges. The "intelligent bridge" approach requires the bridge to have knowledge of the system configuration, but the originating stations do not. The originating station transmits a message on the network addressed to logical or physical task without knowing whether or not the message must pass through one or more bridges. The bridge must receive and forward each message received on its primary network. In addition, it must check the address contained within the message, whether logical or physical, and determine if it must copy it and forward it across to its secondary network. Hence the bridge must contain knowledge or tables about the system network and the destination of all logical or physical addresses. The bridge information must be kept current and reconfiguration of any network in the system affects it. The "intelligent bridge" is the most complex of the bridge types in both design and system use.

The "header bridging" technique is the simplest of the bridge types. It requires that the originating station knows the destination network for the message and the bridges, if any, through which the message must pass. Address words are added to the message header for each bridge through which the message must pass. The address words identify each bridge which must copy and retransmit it on its secondary bus. The bridge address words must appear in the header in the same order in which the bridges will be encountered on the network system. Each bridge removes its address word from the header before retransmission on the secondary bus. Hence, after the last bridge, only the DA is left in the message. This technique can be used in any size system with any number of bridges. The adding of bridge address words to the message header does reduce the system efficiency.

A third bridge type is the "protocol bridge". This type utilizes a field in the message protocol or message header to specify bridging. The originating station must know the destination network for the message and specify it in the header information. The bridge only needs to be aware, have tables, of the system configuration, not specific addresses. A bridge receives and forwards all messages on its primary network. It also checks the destination network field. If it specifies a network which requires the use of the bridge, it will be copied and retransmitted on the secondary network. The bridge may be required to modify some field in the header before it bridges, retransmits, the message. This type of bridging can be performed by a simple bridge and with little effect on HSRB efficiency.

4.3.1 Bridging on the HSRB: While any of the three types of bridging discussed above can be used with the HSRB, there are provisions in the protocol for simple "protocol bridging". The HSRB protocol supports up to four rings connected by bridges in a variety of configurations. Both physical and logical addressing across the bridge is supported. The following principles apply to HSRB bridge addressing:

- a. The HSRBs are labeled 0 through 3.
- b. All stations have the same view of the rings. The addressing of a station or task is independent of the ring of the originating station.

4.3.1 (Continued):

- c. A station communicating with another station on the same ring does not need to know which ring they are on.
- d. A station originating a physically addressed message to a station on another ring does need to know which ring it is on.
- e. A station originating a logically addressed message to a station(s) on another ring(s) does need to know which ring(s) it is on.
- f. Bridges need to know the complete ring interconnect system to know which messages to bridge and what modification of the global addressing fields are required on bridged messages.

4.3.2 Global Addressing: Global addressing, addressing across bridges, is controlled by the 8-bit AC field in the header of a message as shown in Figure 4.3-1.

Logical Physical Subfield		Address Word Count Subfield		Global Address Physical Global Address Logical			
0	LP = 0 LP = 1	AW(1)	AW(0)	BA GL(3)	0 GL(2)	GP(1) GL(1)	GP(0) GL(0)

FIGURE 4.3-1 : Address Control Field

The LP bit of the LP subfield is used to specify the addressing mode for both bridged and unbridged Message frames. When the LP bit is a one, the logical addressing mode is specified. When it is a zero, the physical addressing mode is specified.

The GAL subfield is used for logical addressing. Each bit in this field is used to specify a ring on which the message is to be transmitted. For example, if GL(2) is a one then the message is to be transmitted on ring 2. If GL(2) is a zero, then the message is not transmitted on ring 2 unless it is required to reach another specified ring. Multiple bits may be set to specify the transmitting of the message on the multiple rings. The bridge must interpret these bits in relation to the HSRB interconnect system to determine whether or not to bridge the message. The nonbridge stations do not decode these bits as part of the logical DA. They are not required to be zero in order for a station to be permitted to receive the message.

The GAP subfield is used for physical addressing. The BA and GAP GP(n) bits being a zero specifies that the current ring is the destination ring and bridging is not required. When the BA is a one, then the GP(n) bits specify the physical address of the destination ring. These bits, the ring number, become an extension of the physical DA. The bridges must interpret the ring number in relation to the HSRB interconnect system to determine whether or not to bridge the message. The bridge which transmits the message onto the destination ring must set the GAP subfield to zero. A nonbridge station requires this subfield to be zero before it can decode the physical DA.

4.3.2 (Continued):

Bridging is performed on the HSRB with a store and forward technique. The bridge is the destination station on the primary ring, the ring which originates the messages. It copies the message and completes the protocol as any destination station would. The bridge then performs as the originating station for the message on the secondary ring, the ring onto which the message is being bridged. The bridge will place its address in the Source Address field and transmit the message at the first available free token of the proper priority level. The Message Priority subfield of the original message specifies the priority level to be used during bridging. The bridge uses this priority level in making reservations and claiming free tokens.

4.3.3 Four Rings with Serial Bridging: The first bridging example to be discussed is a system of four rings bridged in series utilizing three bridges as shown in Figure 4.3-2. In this example, each bridge must either pass a message containing a global address to its connected bus or not pass it. If required, each bridge can have an associated standby to provide a redundant path. Address translation may be required of a bridge which passes the message.

4.3.3.1 Physical Addressing: Each bridge which receives a message containing a global physical address must do one of three things:

- a. Do not bridge it.
- b. Bridge it to its attached ring with the GA field unchanged.
- c. Bridge it to its attached ring and set the GA field to zero.

The direction of passing (bridging) messages with global physical and logical addressing is shown in Figure 4.3-2. Figure 4.3-3 presents the same information in a different format and also indicates what operation the bridge performs on the GAP subfield.

4.3.3.2 Logical Addressing: Each bridge which receives a message containing a global logical address must do one of two things:

- a. Do not bridge it.
- b. Bridge it to its attached ring with its global address bit and those of other HSRBs not in that path set to zero.

The direction of passing (bridging) messages with global logical addressing is shown in Figure 4.3-2. Figure 4.3-4 presents the same information in a different format and also indicates what operation the bridge performs on the GA field. When a bridge passes a message, it zeros the global address bit of both rings to which it is attached.

In addition, it also zeros those bit for any rings that are not in the direction of the passed message. Note that a message originating on ring 1 or 2 may be required to be passed in both directions.




FIGURE 4.3-2 - Serial Bridge Configuration With Three Bridges

PHYSICAL ADDRESS \ BRIDGE RING	BRIDGE 0:1		BRIDGE 1:2		BRIDGE 2:3	
	0 to 1	1 to 0	1 to 2	2 to 1	2 to 3	3 to 2
1000	No Pass	Pass As 0000	No Pass	Pass Unchanged	No Pass	Pass Unchanged
1001	Pass As 0000	No Pass	No Pass	Pass As 0000	No Pass	Pass Unchanged
1010	Pass Unchanged	No Pass	Pass As 0000	No Pass	No Pass	Pass As 0000
1011	Pass Unchanged	No Pass	Pass Unchanged	No Pass	Pass As 0000	No Pass

FIGURE 4.3-3 - Physical Address Bridging Algorithm

LOGICAL ADDRESS \ BRIDGE RING	BRIDGE 0:1		BRIDGE 1:2		BRIDGE 2:3	
	0 to 1	1 to 0	1 to 2	2 to 1	2 to 3	3 to 2
XXX1	No Pass	Pass As 0000	No Pass	Pass As 0001	No Pass	Pass As 00X1
XX1X	Pass As XX00	No Pass	No Pass	Pass As 000X	No Pass	Pass As 001X
X1XX	Pass As X100	No Pass	Pass As X000	No Pass	No Pass	Pass As 00XX
1XXX	Pass As 1X00	No Pass	Pass As 1000	No Pass	Pass As 0000	No Pass

FIGURE 4.3-4 - Logical Address Bridging Algorithm

- 4.3.4 Four Rings with Parallel Bridging: The second bridging example to be discussed is a system of four rings with parallel bridging as shown in Figure 4.3-5. This network allows a message to be sent from a station on one ring to a station on any other ring with the use of only one bridge. This will minimize message latency since the store and forwarding of the message will only be done once. A bridge would be configured to pass a message only if the destination ring of the message is the secondary ring on the bridge.




FIGURE 4.3-5 - Four Rings With Parallel Bridging

This network has a high degree of fault tolerance. If one or more bridges fail, the remaining can be reconfigured to pass messages around the failed bridges. A message may now have to be passed through two or three bridges to reach its destination ring, but it would get there.

In this example, if there are no failed bridges, each bridge must either pass a message containing a global address to its secondary ring or not pass it. The algorithm of whether or not to pass a message is simple; if the secondary HSRB of the bridge is specified in the global address, the message is bridged. Otherwise, it is not bridged. Likewise, the address translation is also simple; if the message is bridged, the global address is set to zero.

- 4.3.4.1 Physical Addressing: Figure 4.3-6 exemplifies the direction of passing (bridging) messages with global physical addresses and indicates what operation the bridge performs on the GA field.

- 4.3.4.2 Logical Addressing: Figure 4.3-7 exemplifies the direction of passing (bridging) messages with global logical addressing and indicates what operation the bridge performs on the GA field. Note that the GA field of a bridged message should be set to zero when it is desired to prevent another bridge on the secondary ring from passing it again. This is necessary because of the multiple paths between individual HSRBs. The multiple paths would result in a multicast message being transmitted on a HSRB more than once.

WARNING: Each topology requires its own bridging algorithm. Many different bridging topologies are possible - this is but one example. A bridge failure in Figure 4.3-5 creates a new topology and, thus, a new algorithm would be required; for example, Figure 4.3-6 and Figure 4.3-7 would change.

- 4.3.5 Message RETRY - Through a Bridge: The Message frame retry option should be implemented in all stations, including bridges and is a feature selectable by the user.

Bridges are used to pass messages between two identical types of networks using a store and forward technique. The bridge acts as the destination station on the ring which originates the message. The bridge, if it recognizes a DA for the other network to which it is connected, copies the message from the first network and completes the protocol. Once the message is completely copied, the bridge will retransmit it on the other network at the first available opportunity. On this network, the bridge will be the originating station.

The HSRB protocol defines message acknowledgement bits within the Frame Status field (FS). All station's, including bridges are required to finish the acknowledgement. This acknowledgement is valid, however, only on the ring on which the message was passed, not between rings. In the case of bridging over multiple rings, the HSRB protocol provides an acknowledgement on a single ring rather than on end-to-end acknowledgement. That is, receiving a valid acknowledgement from an intermediate bridge does not indicate that the message was received at the final destination station.

The automatic retry option based on this acknowledgement remains a useful option. However, the user is reminded that this retry is not an end-to-end retry; it is on a per ring basis. An end-to-end acknowledgement and associated retry can be built into protocol layers above the HSRB protocol.

4.4 Global Time Reference:

The fully deterministic access protocol of the HSRB is well suited for distribution of a global time reference. Thus, no specific protocols have been included in the HSRB for this function since it can be easily implemented by the system designer using the standard HSRB capability should it be desired for an application. For instance, a specific logical address can be allocated to the time distribution task and an interrupt generated from this address. Precision can be controlled as described earlier under Priority (4.2.2) and Latency control (4.2.3).

PHYSICAL ADDRESS \ BRIDGE RING	BRIDGE 0:1		BRIDGE 0:2		BRIDGE 0:3	
	0 to 1	1 to 0	0 to 2	2 to 0	0 to 3	3 to 0
1000	No Pass	Pass As 0000	No Pass	Pass As 0000	No Pass	Pass As 0000
1001	Pass As 0000	No Pass	No Pass	No Pass	No Pass	No Pass
1010	No Pass	No Pass	Pass As 0000	No Pass	No Pass	No Pass
1011	No Pass	No Pass	No Pass	No Pass	Pass As 0000	No Pass

PHYSICAL ADDRESS \ BRIDGE RING	BRIDGE 1:2		BRIDGE 1:3		BRIDGE 2:3	
	1 to 2	2 to 1	1 to 3	3 to 1	2 to 3	3 to 2
1000	No Pass	No Pass	No Pass	No Pass	No Pass	No Pass
1001	No Pass	Pass As 0000	No Pass	Pass As 0000	No Pass	No Pass
1010	Pass As 0000	No Pass	No Pass	No Pass	No Pass	Pass As 0000
1011	No Pass	No Pass	Pass As 0000	No Pass	Pass As 0000	No Pass

FIGURE 4.3-6 - Physical Address Bridging Algorithm

LOGICAL ADDRESS	BRIDGE RING	0 to 1	1 to 0	0 to 2	2 to 0	0 to 3	3 to 0	1 to 2	2 to 1	1 to 3	3 to 1	2 to 3	3 to 2
0001		NA	P	NA	P	NA	P	NP	NP	NP	NP	NP	NP
0010		P	NA	NP	NP	NP	NP	NA	P	NA	P	NP	NP
0100		NP	NP	P	NA	NP	NP	P	NA	NP	NP	NA	P
1000		NP	NP	NP	NP	P	NA	NP	NP	P	NA	P	NA
0011		NA	NA	NA	P	NA	P	NA	P	NA	P	NP	NP
0101		NA	P	NA	NA	NA	P	P	NA	NP	NP	NA	P
1001		NA	P	NA	P	NA	NA	NP	NP	P	NA	P	NA
0110		P	NA	P	NA	NP	NP	NA	NA	NA	P	NA	P
1010		P	NA	NP	NP	P	NA	NA	P	NA	NA	P	NA
0111		NA	NA	NA	NA	NA	P	NA	NA	NA	P	NA	P
1100		NP	NP	P	NA	P	NA	P	NA	P	NA	NA	NA
1011		NA	NA	NA	P	NA	NA	NA	P	NA	NA	P	NA
1110		P	NA	P	NA	P	NA	NA	NA	NA	NA	NA	NA
1101		NA	P	NA	NA	NA	NA	P	NA	P	NA	NA	NA

NA = NOT APPLICABLE

NP = NO PASS

P = PASS

FIGURE 4.3-7 - Bridging With Logical Addressing

5. IMPLEMENTATION ISSUES:

5.1 RIU Architecture:

The RIU architecture has to be viewed in the larger context of the total system design. One of the main issues to be considered when developing an HSRB ring interface card, in addition to meeting the standard, is the system level functional partitioning. The desired level of integration of upper layer communication protocols will drive the design of this card as well as define the application/host interface.

This is analogous to existing commercial LAN card implementations. For instance, IEEE standard Ethernet (802.3), Token Bus (802.4) and Token Ring (802.5) are available in what is usually referred to as either shallow or deep card versions. The shallow version typically implements just the LAN standard, while the deeper cards also contain LLC (Logical Link Control) functions and present an IEEE 802.2 interface to the host. There are also LAN cards available that go still further in the upper layer protocol integration by implementing the whole OSI protocol stack or the full TCP/IP suite outboard.

A system developer will have to make similar design decisions with regard to how deep to make the HSRB interface card, that is, how much communications protocol to integrate outboard on the HSRB card. Figure 5.1-1 shows an example of a high level system functional partitioning. It consists of a user application, upper layer communication protocols (user specific and real time) and the RIU. HSRB card implementations may include just the RIU function or integrate part of the full suite of the real-time communication protocols (RTCP), resulting in a hardware interface range as shown in Figure 5.1-1.

The scope of this handbook is limited to the RIU. Therefore, references to the host on this document include all upper layer communication functions as illustrated in Figure 5.1-1. It should be understood, however, that many if not all of these host functions that interface to the RIU may execute on the HSRB interface card rather than on the applications processor.




FIGURE 5.1-1 - High Level System Functional Partitioning

- 5.1.1 RIU Example Implementation: Figure 5.1-2 shows an example of a shallow HSRB interface design. The card contains only the RIU Function and no upper layer protocols.




FIGURE 5.1-2 - Ring Interface Unit Functional Diagram

For this example, the host-specific Interface transfers data between the host and the Message Queue. The data are passed in 16-bit words at 1/20 the serial data rate found on the ring.

The Message Queue buffers incoming and outgoing messages. The Queue Controller places both incoming and outgoing messages in priority order. The Message Queue and Queue Controller may be implemented with a single RAM bank and one controller that handles both incoming and outgoing messages. Alternatively, two RAM banks may be used (one for incoming messages, the other for outgoing messages) with separate controllers.

5.1.1 (Continued):

The RIUC (Ring Interface Unit Controller) provides the full duplex parallel data path to one of the two RIUCs. The RIUC that is on the Active ring will pass the outgoing message data from, and will supply the incoming message data to, the Message Queue. The RIUC Interface may also perform some station management functions including: RIU self-test, initialization of the RIUC, and connection of the RIU to the ring via the RIM.

Connections between two RIUCs of the same station in a counterrotating ring configuration are used to pass status between the two RIUCs, or sides, of a station during power-up and reconfiguration, and pass data during normal operations. Data are passed from the active RIUC to the inactive RIUC to either complete a loopback ring or maintain synchronization on an inactive ring.

The RIUC in this example performs the token passing protocol, error detection, error recovery and reconfiguration. The RIUC receives the serial bit stream from the optical receiver, modifies it as necessary, such as claiming a token, and passes the bit stream to the optical transmitter. The RIUC also passes messages to the Message Queue through the RIUC Interface.

The Look-Up Table (LUT) is used to check logical addresses. The LUT could be implemented with a 64k by 1-bit RAM for a 16-bit logical address space. During the start-up routine, the RAM would be initialized with a logic '1' at RAM addresses corresponding to the logical addresses that are valid for this station. Alternatively, a ROM could be used as a static LUT.

The optical transmitters and receivers transfer data to and from the ring and the RIUC. The Voltage Controller Crystal Oscillator (VCXO) is used in the phase-locked loop of the clock recovery circuitry for the serial paths at each receiver. The fixed clock is used for serial data transmission.

5.2 RIU Storage Requirements:

The HSRB has an extremely efficient priority handling system. However, if the systems using the facilities of the HSRB are to gain full use of the priority system provided it is important to ensure that high priority messages are not queued behind lower priority messages in buffers in the path either to or from the ring.

A subsystem that uses a HSRB RIU may be a collection of equipments which share a common terminal. It may be that this is the only thing that they have in common. Even if the subsystem is just one computer, that computer may be running a number of data producing tasks that are somewhat unrelated. For these reasons it may be very difficult for the subsystem to order all of its messages in priority order before sending them to the RIU. It is, thus, likely that the subsystem will send messages to an RIU that are in no particular priority order. (The designer of a subsystem is cautioned to ensure that such a method of operation does not impose an unacceptable delay to higher priority messages).

An HSRB RIU is one entity. It can have the ability to know the priority level of every message that it contains. However, as pointed out, the messages that it contains are unlikely to be arranged in priority order when it receives them.

5.2 (Continued):

In order that higher priority messages are not held up behind lower priority messages the RIU will have to handle the arranging of the message queue. The suggested set of rules for handling such message queues are as follows:

- a. If an RIU has number of messages that are of the same priority then the RIU should pass these messages to the ring in the order that it received them.
- b. If a message is one that requires acknowledgement then the message passed to the ring must be stored until it has been acknowledged so that it can be repeated in the case of a failed transmission.
- c. If an RIU contains messages of different priorities, then it should order these messages and send the highest priority message (that it contains) out first. This applies equally as well to messages for transmission, where the RIU is sending out to the subsystem.

In the case of messages that require acknowledgement, in the event of a nonacknowledged message then the nonacknowledged message should be retried at the next transmission opportunity unless by that time there exists in the RIU a message of higher priority. In this case the higher priority message should be sent out first and the previously nonacknowledged message returned to its place in the message queue.

It is the responsibility of the designer specifying the RIU to ensure that the RIU contains sufficient storage so as to ensure that a situation does not arise whereby messages are lost. The amount of storage required will depend on the number of priority levels, the ML and the re-try philosophy. In many cases the amount of storage required will not be inconsiderable.

5.3 Synchronization and Elastic Buffer Operation:

- 5.3.1 Background of Token Ring Synchronization: A token ring LAN is comprised of a set of stations connected by serial point-to-point links in a closed-loop topology. Synchronization of the stations in a token ring LAN can be accomplished in a number of ways. Independent link synchronization uses independent receiver and transmitter clocks in each and every station in the ring. This eliminates a potential problem with accumulated jitter, reduces bit error rate of the system due to accumulated phase jitter, and simplifies phase-locked loop (PLL) design. The HSRB employs an improved form of Independent Link Synchronization. The conventional and improved Independent Link Synchronization techniques are described next.

- 5.3.2 Conventional Independent Link Synchronization: The asynchronous token ring employs independent link synchronization where an elastic buffer is placed between the receiver and transmitter within each station. The receiver clock recovery device is locked onto the adjacent upstream transmitter's fixed clock source as shown in Figure 5.3-1. The receiver clock rate can differ from the transmitter clock rate in each station. If the receiver clock rate is faster than the transmitter clock rate, then the number of bits in the elastic buffer will continue to increase. If the receiver clock rate is slower than the transmitter clock rate, then the number of bits in the elastic buffer will continue to decrease. Obviously, the number of bits utilized within the elastic buffer will vary. But the elastic buffer size cannot increase or decrease indefinitely without bound; otherwise, an infinite buffer size with infinite delay is required.




FIGURE 5.3-1 - Asynchronous Repeater Functional Diagram

5.3.2 (Continued):

One solution for preventing buffer overflow or underflow is for the protocol to provide a special field during which the elastic buffer for each station can be adjusted to the middle of its range. The adjustment field is placed between or within every message and/or as defined in the standard. The adjustment of the elastic buffer results in bits being either inserted or deleted from the adjustment field. The elastic buffer size and adjustment field must be large enough to accommodate worst case. It is assumed for the worst case that the receiver and transmitter clocks are permitted to drift apart by a maximum of D bits between adjustment opportunities. The maximum value of D can be computed as a function of the maximum number of bits (L) between adjustment opportunities and the maximum clock rate error (E) between the receiver and transmitter clocks within a station. The relationship is $D = L * E$. Since the specified clock tolerance permits the receiver and transmitter clocks to drift apart no more than D bits, the elastic buffer must be capable of either gaining or losing D bits between every adjustment opportunity. This implies that the buffer size B must have a range of at least $2 * D$ bits. When the adjustment occurs, the elastic buffer is set to the nominal delay of D bits. If the elastic buffer is initialized at a value less than D bits during an adjustment opportunity, then there is a chance that the buffer may underflow prior to the next adjustment opportunity. The buffer underflow may occur when the receiver clock rate is slower than the transmitter clock rate. To prevent underflow, bits will be inserted into the adjustment field during the adjustment period. Conversely, bits will be deleted from the adjustment field to prevent the buffer overflow when the receiver clock rate is faster than the transmitter clock rate. If there are N stations in the ring, then it can be shown that the adjustment field must be originated with at least I bits, where $I = D + N$ bits. In describing the elastic buffer operation, it was assumed that the nominal buffer delay is one half its range, and the adjustments, by either inserting or deleting bits from the adjustment field, are symmetric operations. It will be shown next that an asymmetric elastic buffer technique can decrease the buffer delay by skewing the nominal delay to a minimum.

- 5.3.3 Improved Independent Link Synchronization: The efficiency of the protocol for token ring is dependent on the adjustment field size and the average number of bits utilized in the elastic buffer. The inefficiencies in the previous scheme can be minimized either by reducing the adjustment field size, or by reducing the nominal elastic buffer delay, or by reducing both, while maintaining the same time between adjustment field occurrences. This new elastic buffer technique permits a significant reduction in the size of the adjustment field while maintaining a smaller elastic buffer delay. The nominal elastic buffer delay has no effect on the required initial adjustment field size since the elastic buffer underflow can always be avoided by inserting bits into the adjustment field. Since increasing the nominal delay does not decrease the average adjustment field size, the nominal elastic buffer delay should be minimized to improve the protocol efficiency. The minimal nominal elastic buffer delay should be 1 bit. Since the receiver and transmitter clocks in a station are asynchronous, the nominal setting of the elastic buffer delay is theoretically greater than or equal to 1 bit, but less than 2 bits, and the average elastic buffer delay is 1.5 bits.

5.3.3 (Continued):

Although the initial size of the adjustment field is not dependent on the nominal buffer delay, it is dependent on the maximum delay or bit-space that the elastic buffer can provide since the buffer overflow must be avoided by deleting bits from the adjustment field. If the initial adjustment field size is too small, the elastic buffer may overflow. One method for prevention of buffer overflow is by postponing the adjustments which decrease the size of the adjustment field if the adjustment field size has diminished to zero. This postponement is permitted only if there is enough additional delay in the elastic buffer to accommodate it.

To understand the relationship between the elastic buffer size B and the initial adjustment field size I more clearly, assume that the transmitter clock frequencies of N successive stations are decreasing, where the frequency of the first station is as high as the tolerance permits and the frequency of last station is as low as the tolerance permits. Also, assume that, at the beginning of a particular adjustment opportunity, the elastic buffers in each station simultaneously require an adjustment to reduce the delay by 1 bit toward the nominal delay. If the adjustment field is $D+N$ bits long, then all stations can be accommodated at every adjustment opportunity, and B is 1 bit plus the maximum initial elastic buffer delay. This event, however, cannot occur again until at least N adjustment opportunities later because of the small average clock drift between adjacent stations (the sum of the drifts of all N stations cannot exceed D). The initial adjustment field size can be reduced by increasing B or the number of bits in the elastic buffer of each station. This will allow a station to postpone an adjustment if it detects that its present adjustment field size has diminished to zero due to previous bit deletions. The more B is increased, the longer an adjustment can be postponed. The difference between B and the maximum initial elastic buffer delay is equal to the number of postponed adjustment opportunities before it must adjust its elastic buffer. Hence, it can be shown that the adjustment field size is reduced from $I = 1+N$ (for the case $D=1$) to $I = 1+N/(B-2)$.

For example, if $D=1$ bit per adjustment opportunity and the maximum number of stations $N=128$, then the conventional method would require that the elastic buffer be 3 bits long with a nominal delay of between 1 and 2 bits. The adjustment field must be originated with 129 bits. In the new method, if the elastic buffer is 8 bits long with a nominal delay of between 1 and 2 bits, then the adjustment field need be only 23 bits long instead of 129 bits. This reduction in the required size of the adjustment field is significant in improving the token ring throughput and latency performance.

- 5.3.4 Clock Tolerance: The clock tolerance required is a function of the maximum frequency error between two adjacent stations (twice the tolerance) and the maximum number of bits between adjustment opportunities. The maximum spacing between adjustment fields actually occurs between the first IFA immediately preceding the MFSD, and the next IFA immediately following the FS. The worst case is when the message size is 256 words. The number of bits over this span is $190 + 256 * 20 = 5310$ bits. This requires a clock accuracy between adjacent stations of $1/5310 = 0.01883\%$, or a tolerance from a nominal frequency of 0.0094% . The required tolerance is rounded to 0.009% . There should be sufficient clock phase margin within the elastic buffer of each RIU to accommodate some additional clock phase noise.

5.4 Media:

5.4.1 Optical Media: The section discusses some aspects of the optical media for the HSRB.

5.4.1.1 Some Definitions: Optical fibers have several basic characteristics which make them more or less suitable for HSRB applications.

FIBER SIZE: Common core fiber sizes are 50, 62.5, 100, 150, and 200 μm . Cladding adds another 40-100 μm .

MODE OF OPERATION: For the above fiber sizes, multimode operation is implied. Single mode operation is permissible by the appropriate selection of smaller fiber sizes.

Single mode fiber provides a single wave propagation mode. Resulting in wave transmission with less signal distortion and higher bandwidth. Multiple mode fiber has multiple propagation paths causing signal distortion and lower bandwidth performance.

DISPERSION/CHROMATIC DISPERSION (OR BANDWIDTH): These are characteristics of the cable which appear to "disperse" or slow a rise time at the end of the cable. The dispersion results from components of light taking different amounts of time to propagate down the cable. It effectively limits the bandwidth of the fiber.

INDEX: In multimode fiber there are two kinds of index between the core and cladding.

In Step Index fiber, the boundary between core and cladding is a step function, with light at certain angles reflecting back into the core and light at other angles leaving the core.

In Graded index the index of the core is bell shaped in value and light propagates in a curved path down the fiber instead of reflecting sharply.

5.4.1.2 Tradeoffs:

- a. Single mode fiber has much less loss per kilometer than multimode.
- b. Thinner fiber tends to yield significantly poorer light coupling than larger fiber (since area available for light capture is proportional to fiber area).
- c. Thicker fiber is more prone to thermal problems and bend radius limitations than thinner fiber.
- d. Graded index fiber has less dispersion (higher bandwidth) than step index but coupling efficiency is less.

5.4.1.3 Fiber and the HSRB Ring: Without going into exorbitant detail the effect of all these factors on ring design include the following:

- a. Mechanically, with mechanical bypasses taken into account, thin fibers (particularly single mode operation) are less attractive than larger multimode fibers. Although (on paper) there is plenty of operating margin in the single ended ring structure, the bypass is a limiting factor and connector losses will need to be minimized.
- b. With kilometers of fiber, fiber loss itself becomes significant, as does the chromatic dispersion with LED sources. With aircraft-like lengths and 100 Mbaud operation, these are minor factors.

TABLE 5.4.1.3 - Table of Representative Characteristics

Core/Cladding (Microns)	Type	Numerical Aperture	Attenuation
100/140	Step	0.3	5-12 dB/km (850 nm)
50/125	Graded	0.2	3 dB (850 nm) 0.6 dB (1300 nm)
60.5/125	Graded	0.29	1 (1300 nm)
9/125	Single Mode	0.1	0.5 dB (1300 nm) 0.25 dB (1500 nm)

5.4.2 Electrical Media: This section discusses some aspects of the use of electrical media for the HSRB.

Because of the requirements for shielding of cable, triaxial cable (center conductor, inner and outer shields) is recommended for the HSRB operating at 50 Mbaud.

The signal is transmitted and received differentially on the center conductor and inner shield with the outer shield generally tied to chassis ground at both ends.

Some cable characteristics germane to the HSRB are:

- a. SIZE: Outer jacket size is usually the significant parameter.
- b. IMPEDANCE: Choices range from 50 to 100 ohm(s) cable.
- c. LOSS: Cable losses can be 10 dB/km and up.

5.4.2.1 Tradeoffs:

- a. Higher impedance cables have connectors which are difficult (more expensive) to build and are commercially less available.
- b. Lower impedance cables take more power to drive and tend to be higher loss for the same size cable.
- c. Lower impedance cables tend to be heavier.
- d. Lower loss cable is generally heavier and larger diameter.

5.4.2.2 Cable and the HSRB Ring: The effect of all these factors on ring design include the following:

- a. While less power is required by 75 Ω cable and there may be some attenuation advantages, the easy availability of 50 Ω connectors may sway the decision to 50 Ω triax (RG-8/U triax, RG-58/U triax), choosing the lightest weight cable to meet attenuation requirements.
- b. Where attenuation is critical (very long cables), the cable chosen is driven by the attenuation specification and depending on other factors, may be 50 or 75 Ω .
- c. "Premium" cable (matching impedances of lots to within 2 Ω) may be important in systems with long cables or otherwise operating on the margins.

TABLE 5.4.2.2 - Table of Representative Characteristics

Cable	Belden	Impedance	Outer Diameter	Attenuation 100 m/50 MHz
RG-58A/U	9222	50 ohms	0.24 in	10.8 dB
RG-59/U	9267	75 ohms	0.36 in	2.8 dB
RG-11/U	9192	75 ohms	0.52 in	3.3 dB
RG-11/U	8233	75 ohms	0.475 in	3.3 dB
RG-8/U	9888	50 ohms	0.48 in	3.9 dB
RG-59/U	8232	75 ohms	0.315 in	5.9 dB
RG-11/U	9232	75 ohms	0.52 in	3.3 dB

5.5 Bypasses:

The function of the bypass in a ring bus is to divert signals which normally go through a station around that station per Figure 5.5-1. The bypass is necessary primarily to accommodate station failure when other means of maintaining ring integrity fail.

Two major approaches to the bypass are the relay model (Figure 5.5-2) and the bleeder model (Figure 5.5-3). The relay model bypass consists of a DPDT switch configured in normally closed condition. When power to the bypass is absent, data goes through the bypass without going through the station and the output of the station is connected to its input.

5.5 (Continued):

The bleeder model is applicable to systems with dynamic gain control on the receivers. The bypass adds a small amount of the incoming signal to the output of the station. If the station is transmitting, its signal is strong and the receiver is not affected by the small superimposed bleeding signal. If the station is not transmitting, the bleed signal is the only one present and the gain control on the receiver adjusts to this lower level.

Bypasses are needed for both electrical and optical systems, but each technology has unique problems which are discussed in the following paragraphs.

- 5.5.1 Electrical Bypassing: Major design issues of the bypass for the electrical ring are the shock and vibration limits a switch must operate through and the high bandwidth it must transmit with limited loss. The shock environment calls for a nonmechanical mechanism such as a semiconductor switch, while the frequency of operation calls for a direct electrical connection such as a relay.


Both mechanizations of the switch are conceptually straightforward. Figure 5.5-2 shows the block diagram of the function in either case. For the SAE ring, a terminal is bypassed if power is removed from the bypass, or if the terminal via a control signal causes itself to be bypassed.

- 5.5.1.1 Electromechanical: The electromechanical bypass is the straightforward relay approach using ruggedized double-pole, double-throw relays.


A problem with an electromechanical relay is that given enough force, a relay will either open when it is required to be closed, or will be unable to switch when commanded. One vendor quotes 30 g to 3000 Hz for vibration, 1000 g for 6 ms for shock, and 75 g for acceleration, all without chattering. Tradeoffs to get more shock at the expense of vibration resistance are possible, but whether a mechanical relay is at all suitable is dependent on the specific requirements.

Typical numbers for long-term reliability are quoted in terms of operations and specifications in the tens of millions of operations and typical losses for a bypass operating at 50 Mbaud are 0.5 dB total.

Another aspect of Electromechanical switches that should be considered is their switching characteristics. Detailed behavior of a relay during switching is not usually specified by vendors. A guaranteed stable switching time is specified, but how a relay switches under shock and vibration affects how the ring reconfigures. Total switching time and chattering during a transition is critical to the performance of the ring since the ring reconfiguration response time is of the same order of magnitude as the relay switch time. A slow bouncing switch could cause a reconfiguration around the bouncing terminal, only to reconfigure again when the bypass is finally connected the terminal.



A) RING TERMINALS IN UNBYPASSED STATE



B) RING TERMINAL T IN BYPASSED STATE

FIGURE 5.5-1 - Bypass Function in a Ring System




FIGURE 5.5-2 - General Bypass Configuration: Relay Model




FIGURE 5.5-3 - General Bypass Configuration: Bleeder Model

5.5.1.2 Semiconductor: The semiconductor relay approach uses low on-resistance FETs. Very low ($<1\ \Omega$) drops in FETs are possible, but at the expense of capacitance (1000's of pf.) Currently, there is no good choice for a practical semiconductor bypass at 50 MHz. There is no equivalent to the low resistance, low attenuation at high frequency of a mechanical relay. If there are applications in which some sort of semiconductor bypass is possible, the other attributes of the approach are good.

Assuming a low capacitance semiconductor switch could be fabricated, the switching time would probably be submicrosecond, or an order of magnitude less than an electromechanical switch. Also a semiconductor bypass could be constructed to the same shock and vibration specification as other semiconductor modules.

The voltages do not approach the limits of FET technology, so reliability should be very good. Semiconductor switches may need a bias voltage for bipolar electrical operation.

5.5.1.2 (Continued):

Losses will be large but dependent upon design. A low ohmic loss device with 1000 pf capacitance will exhibit an impedance of $3\ \Omega$ to ground, which in a $50\ \Omega$ system is generally unacceptable. Reducing the capacitance means increasing series resistance. A semiconductor bypass optimized for 50 MHz and minimum total loss will probably still exhibit a large loss, limiting its application to very special cases.

5.5.1.3 Electrical Regenerative Bypass: A solution to the electromechanical and semiconductor relay limitations is some sort of regenerative bypass. Like the terminal itself it would consist of logic which does a limited receive, decode, re-encode and retransmit function. It would add no more to the system delay than the terminal itself.

Its drawbacks are the distribution of power to the bypass, its complexity and power dissipation, and cost.

However the integrated circuit used to implement the terminal could be used in a bypass with some modifications for switching. Except for the need to supply its power, it has the characteristics of an ideal ring bypass.

5.5.2 Optical Bypasses: Optical Bypasses either resemble Figure 5.5-2 a relay model, or Figure 5.5-3, a bleeder approach.

5.5.2.1 Relay Model Bypasses: Relay models are mechanical or one of various nonmechanical types. Mechanical optical bypasses have been on the market for several years, but like mechanical electrical relays they need development to meet the shock and vibration requirements of the military.

The various kinds of optical switches mentioned in the literature include such terms as Mechanical Mirrors, Electro-optic Effect, Calcite Liquid Crystal, Thermo-Optic Deflection, Electrowetting, Electro-optic Lithium Tantalate, Lithium Niobate, and Magneto-optics. Brief descriptions are appended to this section.

5.5.2.2 Bleeder Bypasses: The operation of the bleeder relies on a receiver being able to reject a small bleed signal in the presence of a strong primary signal, and accept the small bleed signal in the absence of a strong primary signal.

Table 5.5-1 gives an example of the kind of analysis which must take place to determine bleeder characteristics and transmitter receiver characteristics. Here a bleeder which attenuates the through light by -9 to -11 dB is assumed. A six connector path (1 from the terminal, 2 each on two bypass units, and 1 to the terminal) is also assumed. At the bottom of the table three parameters are isolated: the bleed signal to primary signal ratio (worst case), the worst case (minimum) power of the bleed signal into the receiver, and the worst case (maximum) primary power into the receiver. In this particular example, the receiver requires a 20 dB operating range which is achievable but requires that the primary signal must be received properly with bleed signal "noise" attenuated only 4 dB riding on top. Attenuating the bleed signal more requires a larger operating range by this analysis.

TABLE 5.5-1 - Sample Bleeder Bypass Analysis

Parameter	Value	Notes
launch power min	-6 dBm	
launch power max	-3 dBm	
loss/connector min	0.5 dB	
loss/connector max	1.5 dB	
receiver sensitivity min	-32 dBm	
receiver sensitivity max	-10 dBm	
input to atten from T-1 min	-9 dBm	assume two connectors, no cable loss
input to atten from T-1 max	-4 dBm	assume two connectors, no cable loss
input to atten from T min	-9 dBm	assume two connectors, no cable loss
input to atten from T max	-4 dBm	assume two connectors, no cable loss
attenuated bled input mix	-20 dBm	assume -11 dB bleeder attenuation
attenuated bled input max	-13 dBm	assume -9 dB bleeder attenuation
bled to primary ratio max	-4 dBm	max bled power-min primary power
bled to primary ratio min	-16 dBm	min bled power-max primary power
bled power at rcvr min	-26 dBm	-20 dBm through 4 more connectors
primary power at rcvr max	-6 dBm	max launched through 6 connectors
Critical Parameters		
bled to primary ratio max	-4 dB	
bled power at rcvr min	-26 dBm	
primary powr at rcvr max	-6 dBm	

5.5.2.2 (Continued):

With receiver operating ranges limited to a 20 to 30 dB range a Bleeder model bypass can only allow for one station being down.

5.5.2.3 Descriptions of Some Optical Switch Approaches: For completeness all major approaches are addressed here. A brief description of some of the approaches is given.

5.5.2.3.1 Mechanical (Mirrors): The bypass is accomplished by placing a mechanically driven mirror in the optical path.

5.5.2.3.2 Liquid Crystal: Light is guided by the refractive index of the liquid crystal which is controlled by the applied E-field. The switching angle is 14°. The switch consists of a waveguide with liquid crystal at one interface. The applied field causes a change in the index of refraction of the liquid crystal, causing the incoming light to exceed the critical angle for total internal reflection.

5.5.2.3.3 Electro-optic Effect: Light is passed through a crystal which changes the polarization as a function of the applied E-field. This is utilized in a bypass switch by placing a Wollaston Prism (W.P.) at the input and output. The input W.P. forces the input light from parts 1 and 2 into 2 linear polarizations. Changes in polarization are caused by the applied voltage to the LiNbO_3 (Lithium Niobate) crystal. The exit beams are then separated by another W.P. with orthogonal linear polarization.

5.5.2.3.4 Calcite Liquid Crystal: The switch uses the calcite crystal as:

- a. A polarization separator
- b. An optical beam deflector
- c. A polarization combiner

The LC is in the form of a 90° twisted nematic liquid crystal (TNLC) which causes a 90° rotation of the polarization with no voltage applied (bypass state), or left unchanged with 6 V applied (on state). Because the calcite is birefringent the S and P polarizations are separated by C-2 resulting in a different exit point for different input polarizations.

5.5.2.3.5 Thermo-optic Deflection: The temperature dependence of the refractive index of glass is used to deflect light to the desired output port. Switching is accomplished by applying voltage to a small film heater deposited on glass.

5.5.2.3.6 Electrowetting: Electrowetting is a term used to describe electrically induced surface tension effects. The surface tension is used to force a mercury slug into or out of the path of an optical light beam.

5.5.2.3.7 Electro-optic LiTaO (Lithium Tantalate): The switch is based on Pockels effect (a linear shift in refractive index with applied voltage) with the input beam steered by the refractive index of the LiTaO. Coupling is via SELFOC lenses. The switch operates with multimode fibers.

5.5.2.3.8 Lithium Niobate Integrated Optical Switch: Waveguides are formed in the Lithium Niobate. The index of refraction of the wall of the waveguide is controlled by the applied voltage. The curvature of the waveguide and the distance for which the coupled waveguides are parallel determines the coupling efficiency.

5.5.2.3.9 Magneto-optic: Light is deflected by the Magneto-optic effect in a Garnet crystal. The stripe domains of the Garnet are controlled by the applied magnetic field. These stripe domains act as a diffraction grating, deflecting the light to any point on a defined circular pattern. The magnetic field is generated electrically using a coil in close proximity to the garnet crystal.

- 5.5.2.4 Active Bypasses: An active bypass can be defined as a device which performs the bypass function described in 5.5, but has an optical receiver and transmitter in the bypass. That is, it has active devices, as opposed to the passive bypass which has at most relay drive current or bias voltage applied.

The primary advantage of the active bypass is the ability to regenerate the optical power and allow an indefinite number of stations to be bypassed. Passive schemes are limited to as few as one bypass in series.

The primary disadvantage of the active bypass is that it is another active, power-dissipating, (and expensive) item in what is otherwise a "simple" passive system.

However, realistically, the active bypass can add to the overall ring reliability and feasibility since the limitations of passive devices is overcome. In addition, the implementation of group bypasses can be simplified with "smart" bypasses.

- 5.5.2.4.1 Degrees of Active Bypassing: An active bypass contains at least one transmitter/receiver pair and can have several possible designs.

The simplest possible is a design which, per Figure 5.5-4, receives and transmits the bypassed signal under the control of the control signal.

While simple, this probably does not work at the frequencies projected for the HSRB, because of pulse width distortions introduced by both optical components. After one or two bypasses, the pulse widths could be distorted by several nanoseconds and the error rate could become unacceptably high. The advantage of the active bypass is really not achieved.




FIGURE 5.5-4 - Simple Active Bypass Design

5.5.2.4.1 (Continued):

Figure 5.5-5 illustrates a second design, in which a nondecoding retrimming element is introduced to restore edge transitions. The design of this element takes into account the known legal pulse widths and maximum allowed jitter. It is a device which, with a delay of the order of 1 bit time, inputs slightly distorted electrical signals and puts out the same signal with edges restored to near-perfect. When these are transmitted through the optical transmitter, the original optical distortion specifications are again met and the bypass becomes optically transparent. In principal, an indefinite number of bypasses can be invoked.




FIGURE 5.5-5 - Active Bypass With Retiming

The problem with this approach is that, while pulses are being manipulated, no error checking occurs and system test and failure tracking becomes difficult. Technologically, such retiming systems operate successfully on Manchester at 50 MHz and could be assumed to be practical at 100 MHz 4B5B and above.

Figure 5.5-6 illustrates a third design, in which a decoder and encoder replaces the retiming circuit. The advantage to this is the error detection now present in the bypass. The disadvantage is more cost and complication and a longer delay. Since the delay is the same as the station delay this really doesn't affect the system. So the system becomes somewhat smarter but needs a new layer of software management to make it work. The (limited) intelligence of the decoder is useful only if it can be used in diagnosis of itself or upstream transmission system. It thus requires a status line to the host terminal.

Figure 5.5-7 shows a variation of the previous in which the bypass receives the terminal's transmitted signal on a second receiver and thereby can a bad terminal transmitter (which the terminal itself, to this point, could not do directly).




FIGURE 5.5-6 - Active Bypass With 4B5B Decoder




FIGURE 5.5-7 - Active Bypass W/Decoder and Additional Receiver

5.5.2.4.1 (Continued):

None of these designs so far allow loopback testing of the terminal. In an active design such a function requires expensive additional transceivers or some mechanical crossover net which reintroduces some of the problems avoided with the active approach.

But 5.5-7 actually allows a form of self-test capable of detecting bad drivers and receivers with minimal hardware. The somewhat intelligent bypass can detect and report absence of terminal optical output. When installed in a system, a bad terminal receiver can be diagnosed by the bypass reporting activity on its receiver while the terminal notes no activity on its receiver.

However, no design so far allows the full stand-alone loopback capability of the generalized bypass.

Figure 5.5-8 illustrates what this would take. This requires the largest number of optical elements.




FIGURE 5.5-8 - Active Bypass W/Loopback

Once the decision is made to make the bypass active, it is inexpensive to add system functions which simplify such things as group bypassing. As an example, in a group bypass, simple logical functions which take in information from a group of terminals can be embedded and allows the group bypass to invoke itself, rather than be dependent on a single terminal.