HEWLETT PACKARD

TACHYON User's Manual

HPFC-5000

PTI 172419

TACHYON

User's Manual

HEWLETT® PACKARD

HP Manual Part No. 5965-1247E

Printed in USA May 1996

First Edition

PTI 172420

Notice

Copyright and Trademark Information

TACHYON is a trademark of the Hewlett-Packard Company.

Patents are pending for technical information presented in this document.

Warranty

The information contained in this document is subject to change without notice. Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Print History

Technical Reference Specification Revisions

Revision 1.2, December, 1993
Revision 1.4, January, 1994
Revision 1.5, January, 1994
Revision 2.0, April, 1994
Revision 2.1, April 25, 1994
Revision 3.0, July 15, 1994
Revision 4.0, December 16, 1994
Revision 5.0, Draft, April 26, 1995
Revision 5.0, May 25, 1995

User's Manual Revisions

1st Draft, August 1, 1995
2nd Draft, December 4, 1995
3rd Draft, February 14, 1996
4th Draft, May 13, 1996 (Internal Distribution Only)

First Edition, May 27, 1996 (Printed and Bounded)

PTI 172421

Contents

PTI 172428

Tables

PTI 172431

Figures

PTI 172433

Preface

Tachyon Errata

For Tachyon chip discrepancies and Tachyon User's Manual updates, please refer to the 3.3V Tachyon Errata. The Tachyon Errata is available by ordering publication number 5965-1433E.

Target Audience

This Tachyon User's Manual is targeted for a variety of readers. It is primarily intended for technical managers and hardware and firmware engineers who are evaluating the Tachyon Fibre Channel controller chip or designing it into their Fibre Channel products. Those readers will want a thorough understanding of the contents of this book. Readers wanting to understand Tachyon functionality without delving into implemention specifics can concentrate on Chapters 1 through 4. Readers wanting a general overview of Tachyon's capabilities can obtain that information from Chapters 1 through 2.

Notes, Cautions, and Warnings

Note	This is the format for the note callout. Notes highlight specific information for a particular section.

Caution	This is the format for the caution format. Cautions describe situations that should be avoided.

WARNING	This is the format for the warning callout. Warnings are primarily used to describe conditions that may cause data corruption.

Document Use

This document is not a Fibre Channel tutorial and does not contain general information about the Fibre Channel architecture. Readers are assumed to understand both the Fibre Channel technology and the content of the Fibre Channel documents listed in the at the end of this section. Refer to "Related Publications" on page xxi.

WARNING	Tachyon is meant to be used according to this Tachyon User's Manual. If it is not, Tachyon may enter an indeterminate state and/or produce unpredictable results. In addition, if any of the Fibre Channel specifications are violated, the host is responsible for all error recovery.

PTI 172434

Chapter Summaries

Chapter 1 Introduction

Chapter 1 provides introductory information about the Tachyon Fibre Channel controller chip and how it operates. Chapter 1 provides a high-level introduction to Tachyon and includes a list of its features and a block diagram of a generic Fibre Channel host-bus adapter board with Tachyon as the controller. It also provides general information about how the Fibre Channel protocol layers map onto Tachyon's functional blocks.

Chapter 2 Architectural Overview

Chapter 2 provides an architectural overview of Tachyon's transmit and receive processes including in order and out of order reception of multiframe sequences. Also, SCSI hardware assists with Tachyon acting both as an initiator and as a target are described. Block diagrams of Tachyon's functional components are used to illustrate all the overview concepts.

Chapter 3 Architectural Details

Chapter 3 provides detailed information about Tachyon's functionality. The chapter provides all the details about the transmit and receive operations for both non-SCSI and SCSI transactions, with Tachyon as the initiator and the target, and SCSI hardware assists, concepts which were introduced in Chapter 2. Both Tachyon and host-based data structures used to accomplish these processes are illustrated and described. ACK generation, messge completion, interrupt generation and avoidance, and exception conditions are described. Tachyon's arbitrated loop operation, TCP/UDP hardware assists, and network management features are also discussed.

Chapter 4 Initialization and Configuration

Chapter 4 provides Tachyon memory requirements, initialization and configuration information with examples, and login parameters.

Chapter 5 Registers

Chapter 5 presents detailed diagrams, tables, and descriptions of the Tachyon-based registers. Initialization, alignment information, and programming notes are included.

Chapter 6 Data Structures

Chapter 6 contains detailed diagrams, tables, and descriptions of the host-based and SCSI data structures needed to program Tachyon. Data structure initialization, alignment information, and programming notes are included.

Chapters 7, 8 and 9 Tachyon Signal, Electrical, and Mechanical Descriptions

Chapters 7, 8, and 9 are hardware-oriented chapters. Chapter 7 covers Tachyon's signal descriptions, logic symbol, pinout, Physical Link Module (PLM) interface, PLL external connections, JTAG scan test interface and instructions, the Tachyon System Interface (TSI), functional waveform diagrams, and TSI and GLM timing information. Chapter 8 contains Tachyon's electrical descriptions, including absolute maximum electrical ratings, recommended operating conditions, electrical parameters, pull-up and pull-down values, and external PLL components. Chapter 9 presents Tachyon's mechanical descriptions, including dimensions and thermal specifications.

PTI 172435

Chapter 10 Error Information

Chapter 10 covers error detection and recovery, such as information about the types of errors Tachyon can detect, tables of error conditions for various Fibre Channel configurations and classes of service, and detailed descriptions of possible errors, including Tachyon behavior and host recovery procedures.

Appendix A Application Notes

Appendix A consists of Application Notes, each of which discusses an aspect of Tachyon for a particular implementation, including PCB layout suggestions, implementing Tachyon with a little endian system, using Tachyon in limited airflow applications, suggestions for improving performance, and information about SCSI command forwarding.

Appendix B Host Driver Notes

Appendix B lists a number implementation tips for host-based drivers.

Appendix C Building Data Structures

Appendix D contains information for building data structures.

Appendix D 8B/10B Encoding/Decoding

Appendix D describes the 8B/10B Encoding/Decoding scheme performed by the FC-1 layer and built into Tachyon.

Appendix E Fibre Channel Frame Quick Reference

Appendix E is a Fibre Channel Frame Quick Reference sheet. This sheet shows specific field and bit information for the 24 bytes of the FC Frame Header. Information for this sheet was taken from FC-PH 4.3.

Terms and Abbreviations

The Terms and Abbreviations section contains a list of acronyms and abbreviations.

Glossary

The Glossary contains the definitions of key words and phrases.

Index

The Index entries point to the main page numbers for detailed information about that subject. Most of the entries do not contain the page number for every occurrence of that entry.

Document Conventions

Length of Queues

The length of a queue is actually the number of entries in a queue. Refer to "Producer and Consumer Indices" on page 19. Since a queue is zero-based, the programmed value for the length field of a queue length register must be the number of entries minus one. Refer to "Table 5.2 Length Registers Information" on page 113.

Reserved

All fields labeled "Reserved", "Reserve", "Rsvd", "Res" or "R" should not be used and should be initialized to zero.

Remote Node

Remote Node means the same as a destination node, a responder node, or a client.

Set and Clear

"Set" implies "set to binary one". "Clear" implies "clear to binary zero". When a bit is discussed without a "set" or "clear" indication, then the bit is set to one.

Word

A "word" is 4 bytes (32 bits) long.

PTI 172437

Related Publications

ANSI Documents

- Fibre Channel Physical and Signaling Interface (FC-PH)
  X3T11/Project 755D/Rev. 4.3
- Fibre Channel Arbitrated Loop (FC-AL)
  X3T11/Project 960D/Rev. 4.54

Fibre Channel System Initiative (FCSI) Documents

- GigaBaud Link Module Family (GLM)
  FCSI-301-Rev 1.0
- Common FC-PH Feature Sets Profiles
  FCSI-101-Rev. 3.1
- SCSI Profile
  FCSI-201-Rev. 2.2
- IP Profile
  FCSI-202-Rev. 2.1
- Fibre Channel Arbitrated Loop Direct Attach SCSI Profile (priviate loop)
  FCSI-???-Rev. 2.0

IEEE Document

- JTAG Boundary Scan
  IEEE Std 1149.1-1990

PTI 172438

Technical Support

Telephone: 1-800-TACHYON (1-800-822-4966).
E-mail: tachyon@hp.com

Homepage

View Tachyon World Wide Web homepage: http://tachyon.rose.hp.com/

Pricing and Delivery

Contact your local HP Component Sales Representive.

Tachyon C Code

Sample Tachyon C Code is available via anonymous ftp over the Internet. The C code is available in two formats:

1. UNIX (shar) Format

   Filename - ProgGuideCode.shar

   This is a UNIX shar file of the C code .c and .h files. Instructions for unpacking the sample code files are included at the beginning of the shar file.

2. ASCII Text Format

   Filename - ProgGuideCode

   This is a concatenated ASCII file of all the .c and .h sample files. Each file is delimited by:

   ```
   /*              End of C file              */
   ```

To access the files,

On your system, type:	ftp rosegarden.external.hp.com
You will get a connect notice.	
At the name prompt, type:	ftp
At the password prompt, type your Internet e-mail address, for example,	
	username@hostname.somewhere.companyname.com
At the ftp prompt, type:	cd pub/tachyon
At the ftp prompt, type:	get ProgGuideCode.shar (or get ProGuideCode)
At the ftp prompt, type:	quit

PTI 172439

1. Introduction

Tachyon provides a very highly integrated Fibre Channel (FC) interface controller in a single chip solution. It supports both networking and mass storage over a single host connection. In addition, Tachyon provides a high performance controller at an inexpensive cost due to its significant number of design features.

1.1 Tachyon Features

- Provides a single chip Fibre Channel solution.
- Supports 1063, 531, and 266 Mbaud link speeds.
- Supports Fibre Channel Class 1, 2, and 3 services.
- Supports Fibre Channel Arbitrated Loop (FC-AL), Point-to-Point, and Fabric (Crosspoint-Switched) topologies.
- Compliant with the Fibre Channel Physical and Signaling Interface(FC-PH).
- Complies to Fibre Channel System Initiative (FCSI) profiles.
- Provides complete support of both networking and mass storage connections.
- Complies to industry standard MIB-II network management.
- Supports up to 2-Kbyte frame payload size for all classes of service.
- Manages sequence segmentation and reassembly in the hardware.
- Generates automatic acknowledgment (ACK) frames.
- Maintains data structures and data buffers in host memory for a low cost, high performance adapter implementation.
- Allows chip transaction accesses to be kept at a minimum due to host-shared memory data structures.
- Reduces host software support overhead due to the architecture design. Specifically, Tachyon:
  - Manages interrupts to one or zero per sequence.
  - Manages the protocol for sending and receiving network sequences over Fibre Channel.
  - Manages the entire protocol for peripheral I/O transactions via SCSI encapsulation over Fibre Channel (FCP).
  - Computes exact checksums for outbound IP packets and inserts them in the data stream, thereby offloading the host of a very compute-intensive task.
  - Computes an approximate checksum for inbound IP packets that partially offloads the checksum task from the host.
  - Contains hardware header/data splitting for inbound SNAP/IP sequences, offloading the data copy portion from the host.
- Supports up to 16,384 concurrent SCSI I/O transactions.
- Provides FCP assists for SCSI initiators and targets.
- Processes inbound and outbound data simultaneously due to full duplex internal architecture.
- Allows time-critical messages to bypass the normal traffic waiting for various resources via a low latency, high priority path through the chip.
- Interfaces directly to industry and ANSI standard 10-bit and 20-bit FC-0 physical link modules.
- Supports FC-PH-2 ACK_Form F_CTL bits for ACK model control on reception.
- Supports parity protection on internal data paths.
- Supports open broadcast replicate transmission and reception of FC-AL.
- Supports reception of fabric-addressed frames for a Fibre Channel services server on FC-AL.

1.2 Host Bus Adapter Board

A generic Fibre Channel host bus adapter board contains the following components:

1. A Backplane Connector, which connects the backplane interface chip to the host bus
2. A Backplane Interface Chip, which enables connection to PCI, EISA, a propriety bus, etc.
3. Tachyon, the Fibre Channel interface controller chip
4. A Physical Link Module (PLM), for example, a 20-bit Gigabit Link Module (GLM)

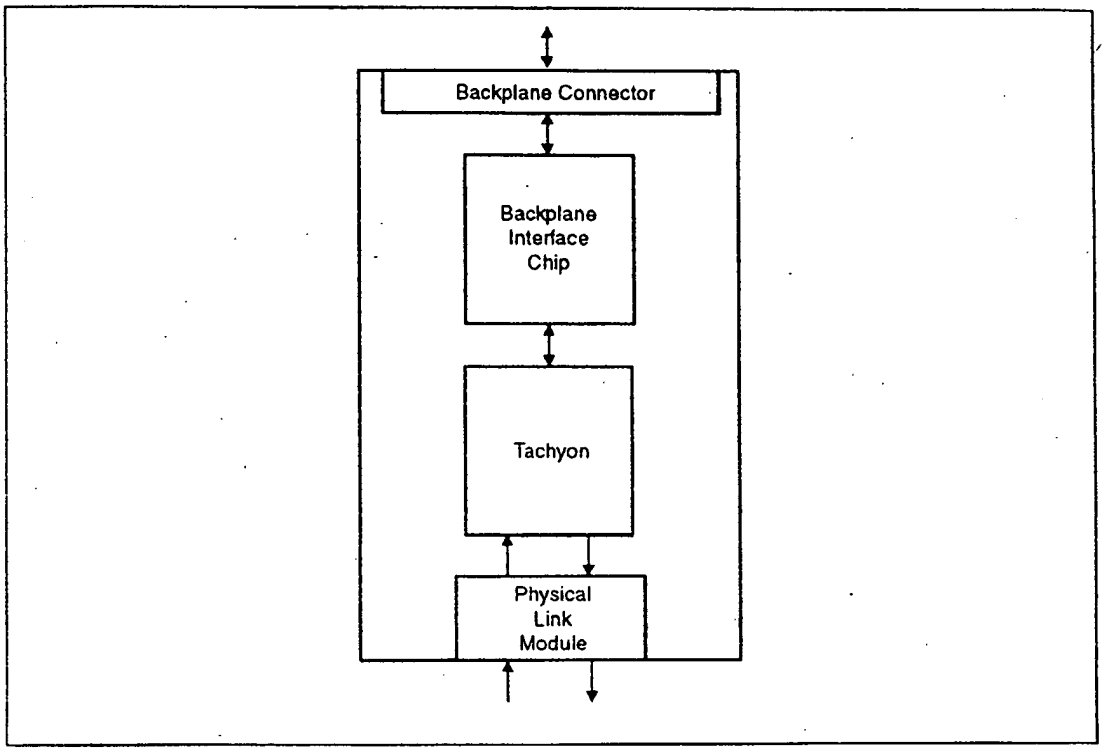

Figure 1.1 Typical Host Bus Adapter Board Block Diagram

PTI 172441

1.3 Fibre Channel General Information

Fibre Channel supports both networking and mass storage interconnections. The Fibre Channel standard defines a number of encapsulation methods. These methods enable both networking and mass storage protocols to be transported over a single Fibre Channel physical interface. This Fibre Channel General Information section gives an overview of the Fibre Channel data hierarchy and describes how Tachyon manages networking and mass storage (particularly SCSI) devices.

1.3.1 Fibre Channel Data Hierarchy

Fibre Channel technology defines a hierarchical structure for data transactions. The data terminology of this hierarchy is defined below:

Exchange

The exchange is the highest level of the data transaction hierarchy. The exchange defines a conversation occurring between two nodes. This conversation can be a long or short term interchange of information. An exchange consists of one or more sequences. An exchange is also known as an I/O transaction.

Sequence

The sequence in Fibre Channel protocol correlates to a data packet or datagram in networking protocols. A sequence consists of a set of one or more frames. Each frame of a particular sequence contains a common Sequence_ID. A node sends frames of a sequence unidirectionally to the remote node. The sequence includes the corresponding Link Control frames (e.g., ACKs, BSYs, RJTs, etc.) that may be sent from the remote node.

A sequence can be any length, but to satisfy various timing and complexity constraints, a node may not be able to send the entire sequence all at once. The sequence can be divided into chunks, or frames, of data.

Frame

The frame is the smallest unit of information interchange between two nodes. A frame must meet many stringent constraints. Various fabrics and nodes may impose their own constraints, so some nodes may have to send frames of varying sizes in order to satisfy these constraints. Frames have the following FC frame format:

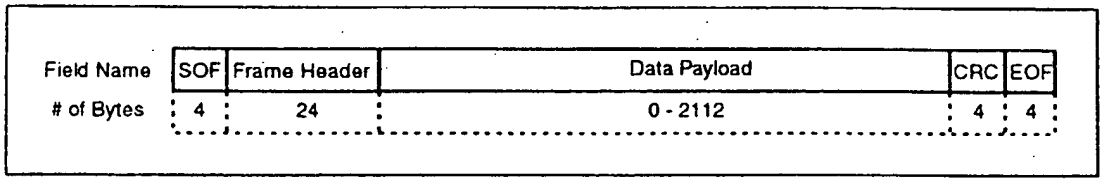

Figure 1.2 Fibre Channel Frame Format

The Fibre Channel standard supports a frame payload up to 2112 bytes; however, Tachyon only supports an outbound frame payload up to 2048 bytes. This outbound frame payload includes optional headers.

PTI 172442

The following figure shows an exchange example.

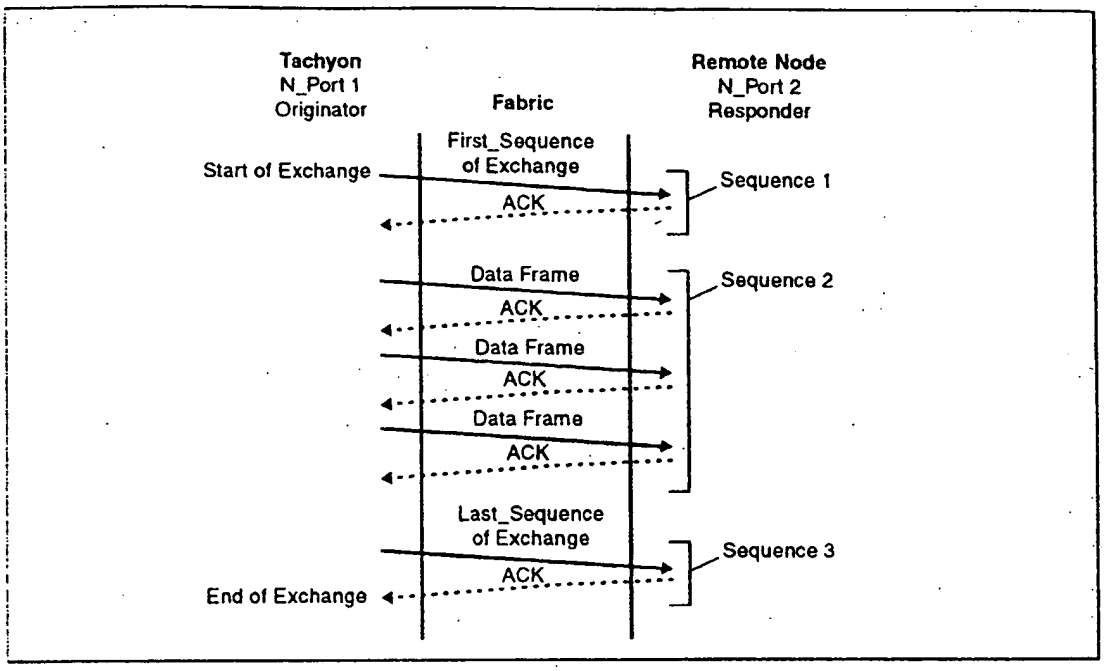

Figure 1.3 Exchange Example

PTI 172443

1.3.2 Networking Encapsulation

Fibre Channel supports any known networking protocol. This is possible because Fibre Channel has defined FC-4, which specifies the mapping of Upper Level Protocols (ULPs) to the lower Fibre Channel levels, FC-3, FC-2, FC-1, and FC-0. Examples of ULPs include SCSI, IP, IPI, etc. The Fibre Channel Physical Interface comprises Fibre Channel protocol levels FC-2, FC-1, and FC-0.

The following figure shows a block diagram of Tachyon with the associated Fibre Channel protocol layers.

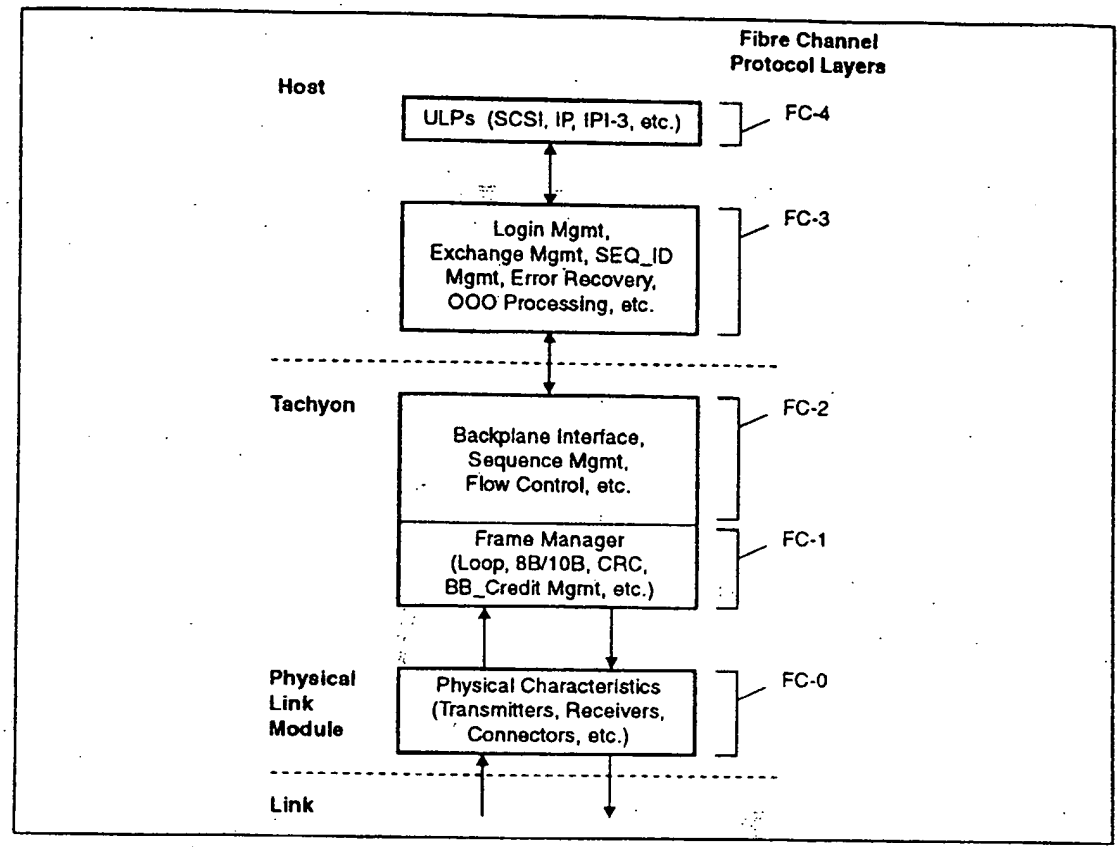

Figure 1.4 Fibre Channel Protocol Block Diagram

Using the Fibre Channel Physical Interface, FC-4 mapping protocols can transfer ULP data. For example, in most networking protocol stacks, the concept of a data packet or datagram exists. The networking data packet or datagram is analogous to the Fibre Channel sequence. To send a sequence, the Originator node segments the sequence into frames. The ULP data packet is written to the data payload field of the Fibre Channel frame(s). Then the SOF, header, CRC, and EOF information is added. The Originator node transmits these encapsulated frames across the Fibre Channel interconnection fabric. The Responder node receives and reassembles the sequence. When the data transmission is complete, the ULP manages the remainder of the data processing.

Mass storage Fibre Channel exchanges map directly to device I/O transactions. However, networking exchanges are not as easily defined and networking exchanges can occur in many different ways. One way is a "long term" exchange. This is where a node sends data to another node, establishes an exchange between the nodes as a virtual link, and keeps the exchange active. Another way is a "short term" exchange, where a TCP connection is mapped onto an exchange. An exchange is created when an application connects. An exchange is terminated when an application disconnects.

1.3.3 Mass Storage Encapsulation

Fibre Channel was originally intended as a follow-on technology to lower performance mass storage interconnects. In particular, Fibre Channel focused on improving SCSI and IPI technologies. The Fibre Channel standards groups defined encapsulation methods for both of these protocols. Tachyon supports both SCSI and IPI protocols, and provides special hardware assists for the SCSI protocol. These SCSI hardware assists execute the encapsulation protocol in hardware to provide high performance and low latency I/O transactions. The exchange originator transmits these encapsulated SCSI frames across the Fibre Channel interconnection fabric. The responder receives and reassembles the fibre channel sequence. When the sequence is reassembled, the remainder of the data processing is managed using upper layer SCSI protocol.

1.3.4 Fibre Channel Protocol for SCSI

This section includes Fibre Channel Protocol for SCSI (FCP) information and how a SCSI I/O transaction maps onto the Fibre Channel transaction hierarchy.

Each SCSI I/O transaction maps on a separate independent Fibre Channel exchange. For FCP, an exchange splits into three distinct phases: the Command Phase, the Data Phase, and the Status Phase. Each phase involves the transmission and/or reception of one or more Fibre Channel sequences. These sequences provide a mechanism of flow control and notification of impending transfers between the initiator and the target device. The Command Phase, Data Phase, and Status Phase are defined below:

Command Phase

The FCP Initiator uses the Command Phase to inform the target device of its intention to perform an I/O transaction. The Command Phase consists of a single frame SCSI Command Sequence (FCP_CMND) that carries information to the target device(s). This information includes the direction of transfer, the length of transfer, and the SCSI address. After the target device processes this SCSI Command Sequence, the transaction enters the Data Phase. If a Data Phase does not exist, the transaction will enter the Status Phase.

Data Phase

The Data Phase consists of a two step process:

1. The target transmits or receives data via the Data Sequence (FCP_DATA). The direction of transfer is identified in the original Command Sequence. This data sequence may consist of one or more frames.
2. The target device sends a Transfer Ready (FCP_XFER_RDY) sequence to inform the initiator that it has the resources ready for a data transfer. The FCP_XFER_RDY sequence is an optional step for an FCP Initiator Read Exchange. When the initiator receives the FCP_XFER_RDY, the process moves on to the second step.

Status Phase

The Status Phase begins when the last FCP_DATA completes. In the Status Phase, the target device sends a single frame SCSI Status Sequence (FCP_RSP) describing how well the transaction proceeded. Once the initiator processes this sequence, the SCSI I/O transaction is complete.

1.3.5 Profile Support

Tachyon supports the FCSI profiles for FC-IP and FCP. Tachyon does not provide hardware assists for any other profiles, such as IPI-3. By using the generic sequence transmit and receive capabilities of Tachyon, the host can support any known profiles. Hewlett-Packard will only test FCSI profiles for FC-IP and FCP.

2. Architectural Overview

This section provides an architectural overview of the transmit, receive, and SCSI hardware assists processes of Tachyon.

2.1 Transmit Process Overview

The following figure illustrates the transmit process overview.

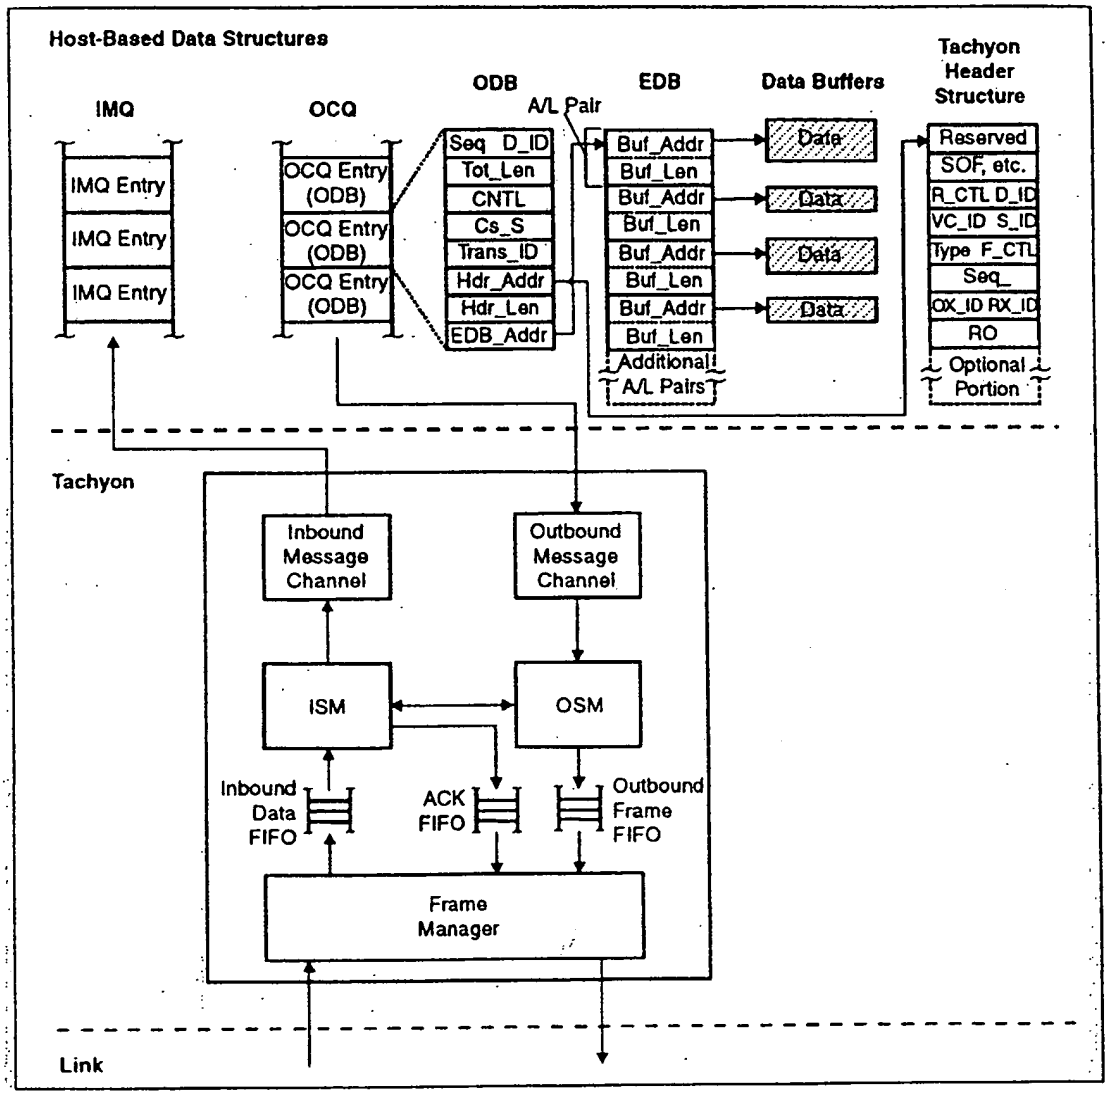

Figure 2.1 Transmit Process Overview

PTI 172447

A normal transmit transaction begins when the host creates buffers of data to be sent. These data buffers are stored in host memory. The host creates a Tachyon Header Structure which contains Fibre Channel header information and resides in host memory. The host creates an Outbound Descriptor Block (ODB) host data structure that defines the sequence of data to be transmitted. The ODB contains Fibre Channel information (e.g., maximum frame size, class of service, exchange identifier, sequence identifier, etc.) that is obtained from the host managed login parameters. The ODB contains a pointer (Hdr_Addr) to the Tachyon Header Structure. Tachyon uses the Fibre Channel information in the ODB and the Tachyon Header Structure to construct the Fibre Channel frame headers. The ODB also contains a pointer (EDB_Addr) to the Extended Descriptor Block (EDB). The EDB contains Address/Length pairs (A/L pairs) that define where the data buffers are located in host memory.

The Outbound Command Queue (OCQ) is an area in host memory consisting of an array of OCQ entries. The producer index of the OCQ, known as the OCQ Producer Index, points to the next available OCQ entry in which an ODB can be created. Once the host creates the ODB, it writes the index of that OCQ entry (ODB) to the OCQ Producer Index register. This notifies Tachyon that a new valid ODB exists. The OCQ Producer Index is incremented to the next available OCQ entry for the following valid ODB.

If Tachyon's Outbound Sequence Manager (OSM) is not currently transmitting a sequence, the OSM performs a DMA operation to move the ODB via the Outbound Message Channel from host memory into an internal Tachyon resource. When the OSM receives the ODB, the OSM has all the information needed to transmit the sequence. It has all of the Fibre Channel information to construct the Fibre Channel frame headers and all of the A/L pairs that point to the data in host memory.

The OSM retrieves data from the Outbound Message Channel in frame size packets for transmission. The OSM DMAs the first frame from the data buffer in host memory to the Outbound Frame FIFO. Once the entire frame is in the Outbound Frame FIFO, the OSM notifies the Frame Manager to begin transmitting the frame onto the link. When the first word of the frame is transmitted onto the link, the OSM is notified. The OSM then begins to move the second frame from host memory to the Outbound Frame FIFO. The first frame is being moved from the Outbound Frame FIFO onto the link during the same time that the second frame is being moved from host memory into the Outbound Frame FIFO. This operation continues until the entire sequence has been transmitted.

For a normal Class 1 (dedicated connection with ACKs) or Class 2 (connectionless transaction with ACKs) transmit transaction, the remote node returns ACKs to Tachyon when the remote node receives frames. When Tachyon's Inbound Sequence Manager (ISM) receives a frame, the ISM first determines if it is an ACK frame. If the frame is an ACK frame, the ISM passes the ACK to the OSM. The OSM verifies that the ACK is associated to the current outbound sequence and increments an ACK counter. The OSM maintains a count of the number of frames transmitted for the sequence and the number of ACK frames received. When all frames of the sequence have been transmitted and the number of ACK frames equals the number of transmitted frames, then the sequence has been transmitted successfully to the remote node. For Class 3, when the remote node receives the frames, it does not return ACKs to Tachyon.

When the OSM transmits the sequence successfully, it notifies the ISM to generate an outbound_completion message. The ISM sends this completion message to the host as an entry in the Inbound Message Queue (IMQ). Completion messages contain information on the status of the transfer and any information needed by the host for queue maintenance. In this case, the outbound_completion message notifies the host that the sequence has been transmitted successfully.

2.2 Receive Process Overview

The ISM manages receive processes for four types of sequences: 1) Single Frame Sequences, 2) Multiframe Sequences, In Order, 3) Multiframe Sequences, Out of Order and 4) Multiframe Sequences, Deferred P_BSY Mode.

The following figure illustrates the SFS and MFS receive processes:

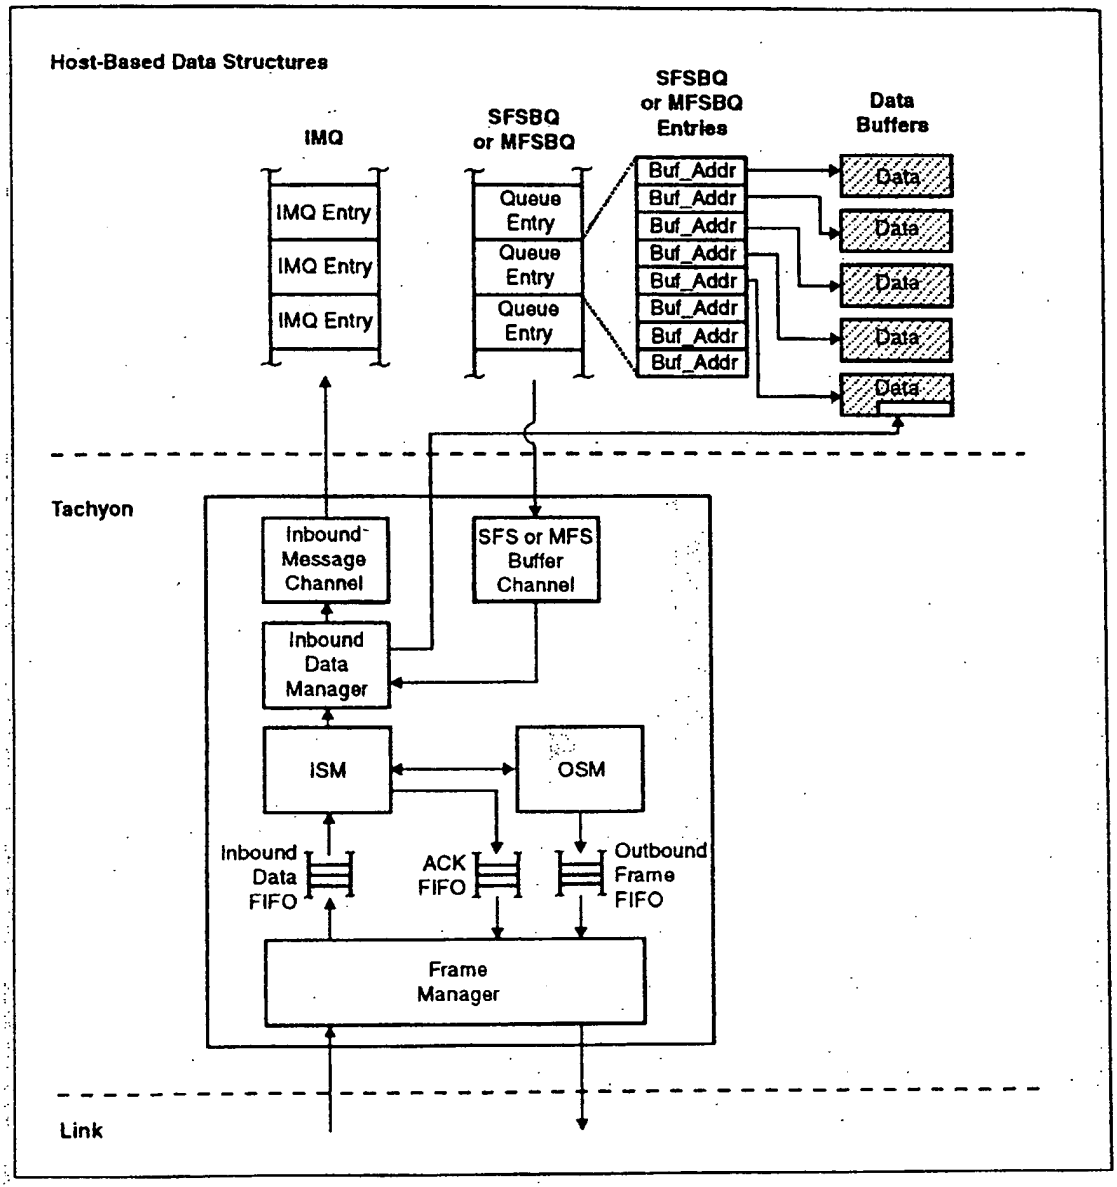

Figure 2.2 Receive Process Overview

PTI 172449

2.2.1 Single Frame Sequence Reception

As Tachyon receives a frame, it stores the frame in the Inbound Data FIFO while it verifies the Cyclic Redundancy Check (CRC). If the CRC fails, Tachyon discards the frame. If the CRC passes, Tachyon notifies the ISM of the received frame.

The ISM checks the frame to determine if it is a single frame sequence (SFS). If it is, Tachyon uses the next available buffer in the Single Frame Sequence Buffer Queue (SFSBQ) and DMAs the SFS (Fibre Channel Header and data payload) to the host via the Inbound Data Manager and the SFS Buffer Channel. If the entire SFS does not fit into one buffer, Tachyon continues packing the remaining data into the next available buffer until it becomes full. If the buffer becomes full, the ISM provides a new buffer and the process continues until the SFS is totally stored. When the DMA operation completes, the ISM transmits the ACK frame (if the frame is Class 1 or Class 2) that is associated with the SFS.

Next, the ISM generates an inbound_sfs_completion message and passes it through the Inbound Data Manager and Inbound Message Channel as an entry in the IMQ. After Tachyon sends the completion message, Tachyon then generates an interrupt to signal the host to process the received sequence. As its default, Tachyon generates interrupts for inbound sequences. For the host to avoid unnecessary interrupts, interrupt avoidance techniques can be used when any type of completion message is passed as an entry into the IMQ. Refer to "3.3.8 Interrupt Avoidance Techniques" on page 33.

2.2.2 Multiframe Sequence, In Order Reception

As Tachyon receives a frame, it is stored in the Inbound Data FIFO while the CRC is verified. If the CRC fails, then the frame is discarded. If the CRC passes, the ISM is notified. The ISM first checks the frame to determine if it is a SFS, which in this case, it is not.

Next, the ISM checks if it is currently reassembling a multiframe sequence (MFS). If the ISM is already reassembling a multiframe sequence, the ISM checks the frame to determine whether the frame belongs to the current multiframe sequence. If it does not, the ISM generates a busy response (P_BSY) to be sent to the remote node and the frame is discarded. The P_BSY informs the remote node that Tachyon does not have the resources available to reassemble another sequence at this time. If Tachyon is in the Deferred P_BSY mode, Tachyon does not automatically send a P_BSY to the remote node. Refer to "2.2.4 Multiframe Sequence, Deferred P_BSY Mode" on page 13.

If the ISM receives a frame for a MFS when it is not currently reassembling a MFS, Tachyon DMAs the frame to host memory. The host uses buffers from the Multiframe Sequence Buffer Queue (MFSBQ) to reassemble a MFS. When Tachyon receives the first frame of the MFS, a new buffer is provided from the MFSBQ and the Fibre Channel Header information of the frame is moved into that buffer. Data payload of the frame is packed into the following buffer. As subsequent frames arrive, the ISM checks if each frame is the next expected frame of the sequence. Assuming that Tachyon receives the sequence in order, Tachyon discards the Fibre Channel Header and passes the remaining data payload to the host to be packed into buffers. As each new frame arrives, Tachyon processes it the same way. Tachyon packs the data payload into the current buffer until it becomes full. Once it becomes full, the ISM provides a new buffer and the process continues. As Tachyon receives each frame and moves it into host memory, the ISM sends the ACK frame (if the frame is Class 1 or Class 2) that is associated with the MFS.

When the last frame arrives, various fields in the Fibre Channel Header are saved in internal registers. Tachyon passes the last frame, with the Fibre Channel Header discarded, to the host. The ISM generates an inbound_mfs completion message and passes it as an entry on the IMQ, and updates the IMQ Producer Index. After Tachyon sends the completion message to the host, Tachyon generates an interrupt to signal the host to process the received sequence. As its default, Tachyon generates interrupts for inbound sequences. To avoid unnecessary interrupts,the host can use interrupt avoidance techniques when any type of completion message is passed as an entry into the IMQ. Refer to "3.3.8 Interrupt Avoidance Techniques" on page 33.

PTI 172450

2.2.3 Multiframe Sequence, Out Of Order (OOO) Reception

The Out of Order Multiframe Reception operates the same as the in order multiframe reception until the first OOO frame arrives. When the ISM determines that the frame is OOO, it generates an inbound_ooo_completion message and passes it as an entry on the IMQ. Then the ISM restarts the reception by copying the entire OOO frame, including the Fibre Channel Header, and sending it to the host using new data buffers. Then the ISM processes the remaining frames in the sequence, assuming that the subsequent frames arrive in order. The ISM does not generate an interrupt until the entire sequence has arrived.

At each point of discontinuity, that is, when a new OOO frame arrives, Tachyon generates an inbound_ooo completion message and uses a new buffer to re-start the reassembly.

2.2.4 Multiframe Sequence, Deferred P_BSY Mode

Tachyon is designed to reassemble one MFS at a time. When Tachyon receives a frame for a new MFS while a current MFS is being reassembled, Tachyon generates a N_Port busy response (P_BSY) for the frame(s) of that new MFS. The remote node that receives the P_BSY can retry the sequence later. Hopefully, when the remote node retries sending the sequence, Tachyon has completed the previous MFS so that it can accept the new MFS.

This process works well when Tachyon communicates with only one other remote node, e.g., a client in a client-server configuration. However, a server in a client-server configuration may be receiving MFSs from many clients (remote nodes) at the same time. When Tachyon receives several new MFS frames, other than the current MFS being reassembled, the new frames are P_BSY'd. The remote nodes re-send the frames causing even more interference with the current MFS being reassembled. These excessive retries may cause the remote nodes to stop sending their MFS, which leads to upper layer error recovery overhead. This overhead can be reduced by using the Deferred P_BSY mode.

In Deferred P_BSY mode, the host sets the Disable AUTO P_BSY bit in the Tachyon Configuration register so that Tachyon does not automatically send a P_BSY to the remote node. When Tachyon receives a new MFS in Deferred P_BSY mode, Tachyon sends the new MFS frame to the host via the SFSBQ. Tachyon generates an inbound_busied_frame completion message. The host stores this frame in a Deferred P_BSY queue.

When Tachyon receives all of the frames for the current MFS it is reassembling, it sets the Deferred ACK bit in the inbound_mfs_completion message and defers the final ACK generation to the host. Tachyon does this so that the host can send a P_BSY to one of the waiting remote nodes. This operation allows one of the waiting remote nodes to retry its MFS before the current remote node can start a new MFS. In Deferred P_BSY mode, Tachyon attempts to process MFSs in a way that all remote nodes have fair access.

2.3 SCSI Hardware Assists Overview

The ability to reassemble only one multiframe sequence may be sufficient for client-based or request reply networking traffic, but is not adequate for I/O traffic. A typical scenario is a host that is connected to many SCSI I/O devices. Because SCSI I/O devices are typically mechanical devices and are relatively slow compared to the host, a host may have hundreds or even thousands of active SCSI I/O transactions. Since it is possible to have many active SCSI I/O transactions, it is necessary to reassemble more than one inbound sequence at a time. By using Tachyon's SCSI hardware assists, the host is capable of concurrently reassembling 16,384 SCSI-assisted sequences.

2.3.1 FCP Read for Tachyon as an Initiator

An FCP Read transaction is considered "inbound" for Tachyon as an initiator. When the initiator host wants to perform a Read transaction, it creates an entry in the host-based SCSI Exchange State Table (SEST). Each used entry in the SEST contains exchange state information for one SCSI I/O transaction. An Inbound SEST Entry contains a pointer to the SCSI Descriptor Block (SDB). The SDB is a list of pointers to empty buffers in host memory. The host pre-allocates these empty buffers to receive all the Read data for the exchange.

The initiator host sends an FCP_CMND for a Read to the target . The FCP_CMND requests the target to send the SCSI read data.

When the target is ready to transmit the SCSI read data to the initiator, it may send an FCP_XFER_RDY, which is an optional step for an FCP Read Exchange. When the initiator Tachyon receives the FCP_XFER_RDY, it discards it. When the initiator Tachyon receives a frame for the Read data sequence from the target, the initiator Tachyon operates in one of two modes, 1) Out of Order Reassembly mode or 2) In Order Reassembly mode. When the initiator host enables OOO Reassembly, then frames for a sequence from the target can arrive in any order. When the host enables In Order Reassembly, then frames must arrive in order.

For each frame received, Tachyon uses the Originator Exchange Identifier (OX_ID) in the Fibre Channel Header as an index into the initiator host SEST. Once Tachyon has identified the apppropriate Inbound SEST Entry, the ISM DMAs the SCSI read data to the host buffers. Tachyon continues receiving the data until all of the data is packed into host buffers.

In the final phase of the FCP Read, the initiator Tachyon receives a FCP_RSP from the target. When the ISM receives the FCP_RSP, it passes the FCP_RSP to host memory via the SFSBQ. This FCP_RSP notifies the host that the FCP Read transaction is complete. Then, the ISM generates an inbound_scsi_status completion message and passes this completion message as an entry on the IMQ. The ISM interrupts the host. The interrupt avoidance rules apply here as well. Refer to "3.3.8 Interrupt Avoidance Techniques" on page 33.

The following figure illustrates the FCP Read process for Tachyon as an initiator:

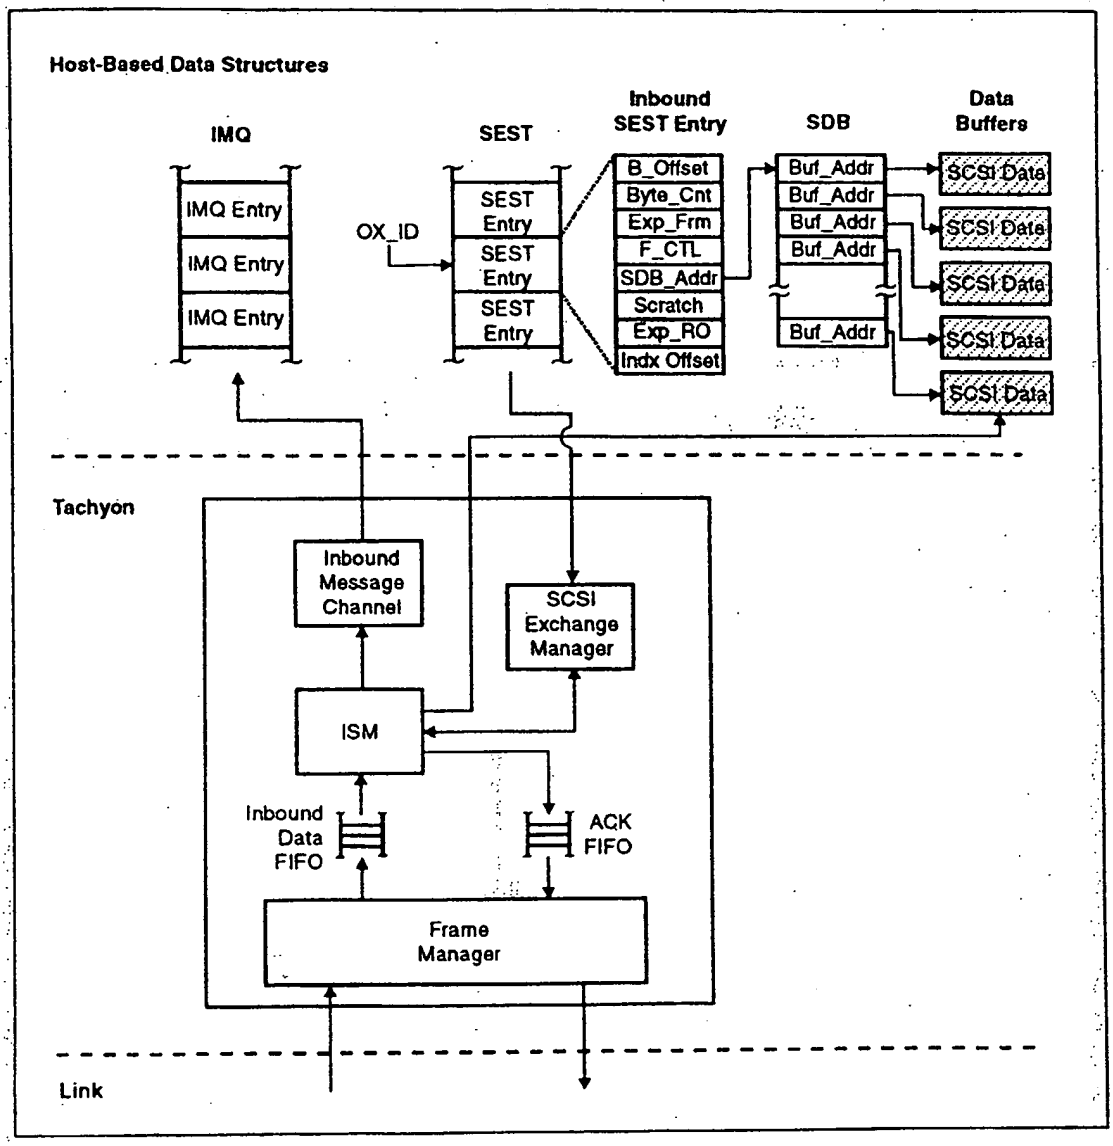

Figure 2.3 FCP Read for Tachyon as an Initiator

PTI 172453

2.3.2 FCP Read for Tachyon as a Target

An FCP Read exchange is considered "outbound" for Tachyon as a target. For FCP Read transactions for the target Tachyon, SCSI hardware assists are not used.

The target Tachyon receives an FCP_CMND for an FCP Read from the initiator and passes it to the target host. The target host builds an EDB that defines where the Read data is located in the target host memory. The target host may program the target Tachyon to send an FCP_XFER_RDY via the OCQ to the initiator indicating that it is ready to send the Read data. The target host then programs the target Tachyon to send the requested SCSI read data via the OCQ to the initiator. When it is finished, the target host programs the target Tachyon to send an FCP_RSP indicating that the FCP Read transaction is complete.

2.3.3 FCP Write for Tachyon as an Initiator

An FCP Write transaction is considered "outbound" for Tachyon as an initiator. When the initiator host wants to perform a Write transaction, it creates an Outbound SEST Entry. The Outbound SEST Entry contains a pointer to an EDB that points to the SCSI write data.

The initiator host sends an FCP_CMND for a Write to the target. When the initiator Tachyon receives a FCP_XFER_RDY from the target, it checks the value of the DATA_RO field. If the DATA_RO is zero, the initiator Tachyon manages the Write transaction. If the DATA_RO field is non-zero, Tachyon passes this FCP_XFER_RDY to the host. In this case, the initiator host is responsible for managing the data transfer.

When the initiator Tachyon sends all the Write data, it waits to receive an FCP_RSP from the target. Initiator Tachyon passes the FCP_RSP and an inbound_scsi_status_completion message to the initiator host. This informs the host that the exchange completed.

2.3.4 FCP Write for Tachyon as a Target

An FCP Write exchange is considered "inbound" for Tachyon as a target. The target Tachyon receives an FCP_CMND for an FCP Write from the initiator. The target host must create an SDB that points to an empty buffer that receives the Write data from the initiator. When the target host has allocated enough buffers and is ready to receive the data, it sends an FCP_XFER_RDY with the DATA_RO field equal to zero to the initiator. If the target host is not ready to manage the entire requested Write transaction, it sends multiple FCP_XFER_RDYs to complete the transaction.

When the target Tachyon receives the SCSI write data, it operates in one of two modes, 1) Out of Order Reassembly mode or 2) In Order Reassembly mode. When the target Tachyon receives all of the data, the target Tachyon sends an inbound_scsi_data completion message to the IMQ of the target host. The target Tachyon then sends an FCP_RSP to the initiator indicating that the exchange has completed.

3. Architectural Details

3.1 Tachyon Internal Block Diagram

The Tachyon Internal Block Diagram shows the high level chip architecture.

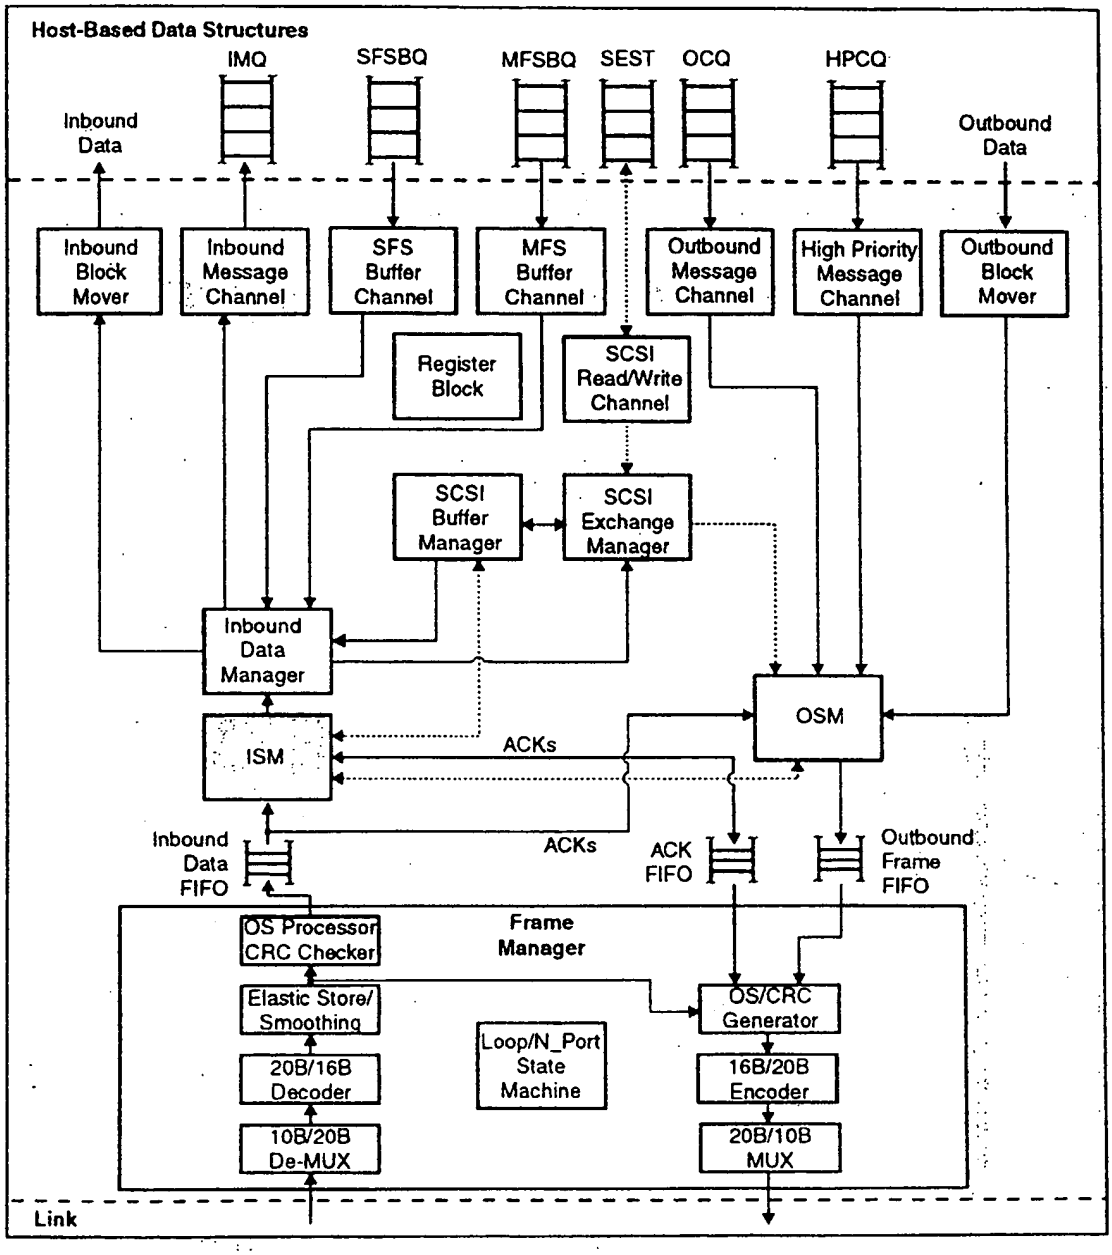

Figure 3.1 Tachyon Internal Block Diagram

PTI 172455

3.2 Circular Queues

3.2.1 Overview

Tachyon and the host use five host-based circular queues to pass messages and memory descriptors.

Circular queues that are located in host memory:

1. Outbound Command Queue (OCQ)
2. High-Priority Command Queue (HPCQ)
3. Inbound Message Queue (IMQ)
4. Single Frame Sequence Buffer Queue (SFSBQ)
5. Multiframe Sequence Buffer Queue (MFSBQ)

Circular queues consist of four main parts:

1. The queue itself
2. The queue length
3. The producer index
4. The consumer index

Circular queues contain entries that describe:

1. Things that have been done
2. Things to do
3. Lists of resources that can be used by Tachyon

Each circular queue is a contiguous area in host memory that can be thought of as an array [0 to (n-1)] of entries (assuming 'n' is the number of entries). All entries in all of the queues are 32 bytes in length. Queues must contain a minimum of two entries, i.e., have a length of two, except for the IMQ which must contain a minimum length of four entries. Base addresses of the queues must be aligned on sizeof (queue) boundary, i.e., 32 multiplied by the number of entries of the queue.

3.2.2 Producer and Consumer Indices

Producer and consumer indices for the queues are decribed in the following three sub-sections: the OCQ and HPCQ, the IMQ, and lastly, the SFSBQ and MFSBQ:

1. OCQ and HPCQ

For the two outbound queues, the OCQ and the HPCQ, the host is the producer and Tachyon is the consumer. As the producer, the host fills in the OCQ or HPCQ entries, i.e, it creates ODBs or HPDBs. After the host fills in an entry, it increments and writes to the OCQ or HPCQ Producer Index register. These producer indices point to the next empty entry that the host can use to create ODBs or HPDBs. The OCQ and HPCQ Producer Indices reside in Tachyon registers.

As the consumer, Tachyon processes the ODBs or HPDBs so that data can be transmitted. Once the ODB or HPDB is processed, the entry is considered empty. When Tachyon processes the ODB or HPDB, it increments and writes to the OCQ or HPCQ Consumer Index. These consumer indices point to the next ODB or HPDB that Tachyon will process. The OCQ and HPCQ Consumer Indices reside in host memory.

For the OCQ and HPCQ, the producer index exists in consumer space and the consumer index exists in producer space. This eliminates the need for any Read operations of indices across the backplane interface, which could reduce the performance of the host I/O bus.

Refer to the following OCQ Producer and Consumer Index Example figure, which is applicable for the HPCQ, as well.

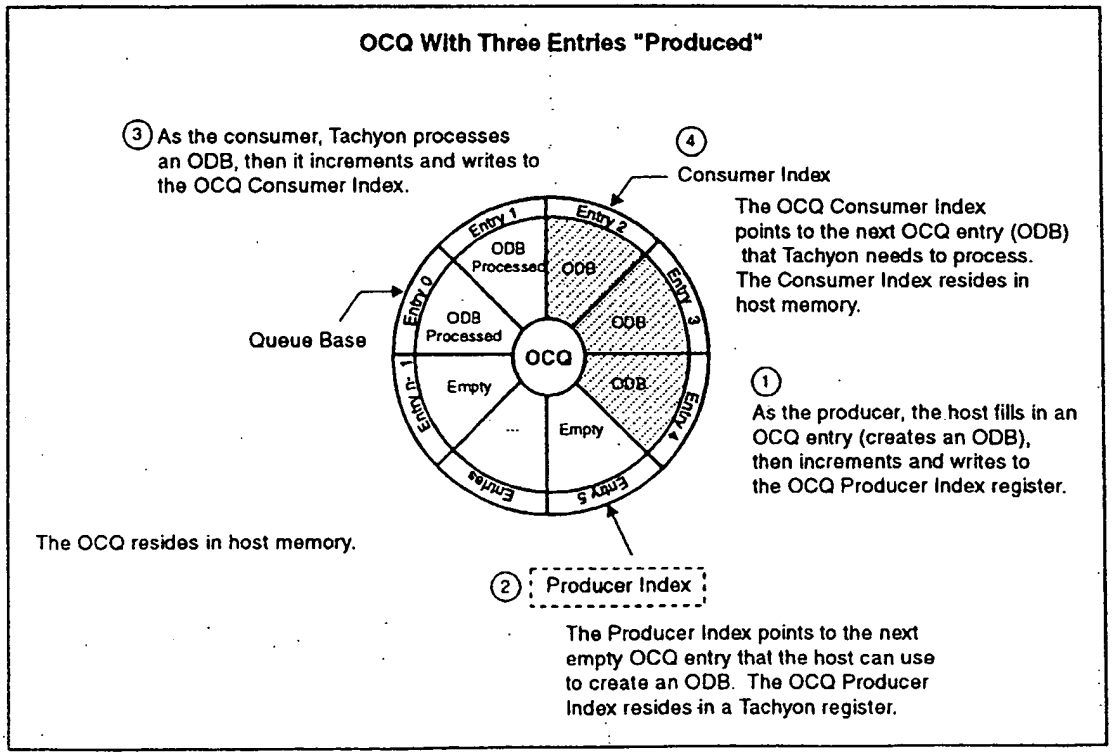

Figure 3.2 OCQ Producer and Consumer Index Example

PTI 172457

When an index points past the end of the queue (the index equals the size of the queue), it wraps to the beginning of the queue.

The OCQ and HPCQ is said to be "full" when the host has filled in n-1 entries, i.e., only one entry remains. The OCQ and HPCQ is "empty" when the consumer has processed all of the ODBs or HPDBs and the consumer index equals the producer index.

Refer to the following OCQ Full and Empty Example figure, which is applicable for the HPCQ, as well.

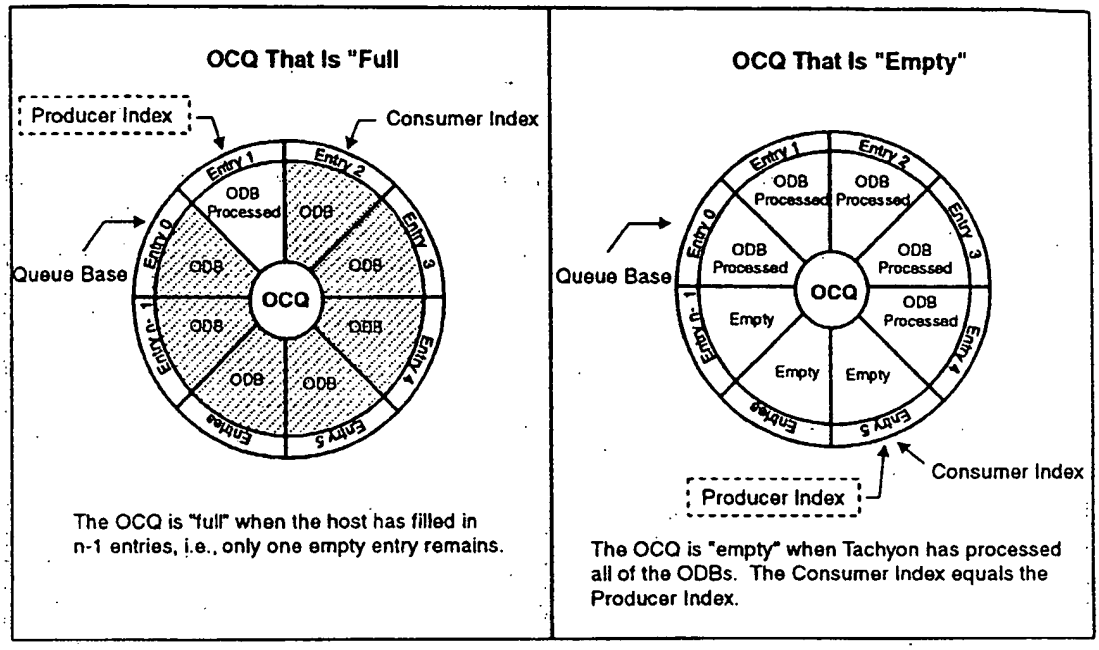

Figure 3.3 OCQ Full and Empty Example

PTI 172458

2. IMQ

For the IMQ, Tachyon is the producer and the host is the consumer. As the producer, Tachyon fills in IMQ entries, i.e., it posts completion messages. After it posts a completion message, Tachyon increments and writes to the IMQ Producer Index. This index points to the next empty entry that Tachyon can use to create a completion message. The IMQ Producer Index resides in host memory.

As the consumer, the host processes completion messages to know the status of sequences. Once the completion message has been processed, the entry is considered empty. When the host processes a completion message, it increments and writes the IMQ Consumer Index. This consumer index points to the next completion message that needs to be processed. The IMQ Consumer Index resides in a Tachyon register.

For the IMQ, the producer index exists in consumer space and the consumer index exists in producer space. This eliminates the need for any Read operations of indices across the backplane interface, which could reduce the performance of the host I/O bus.

Refer to the following IMQ Producer and Consumer Index Example figure.

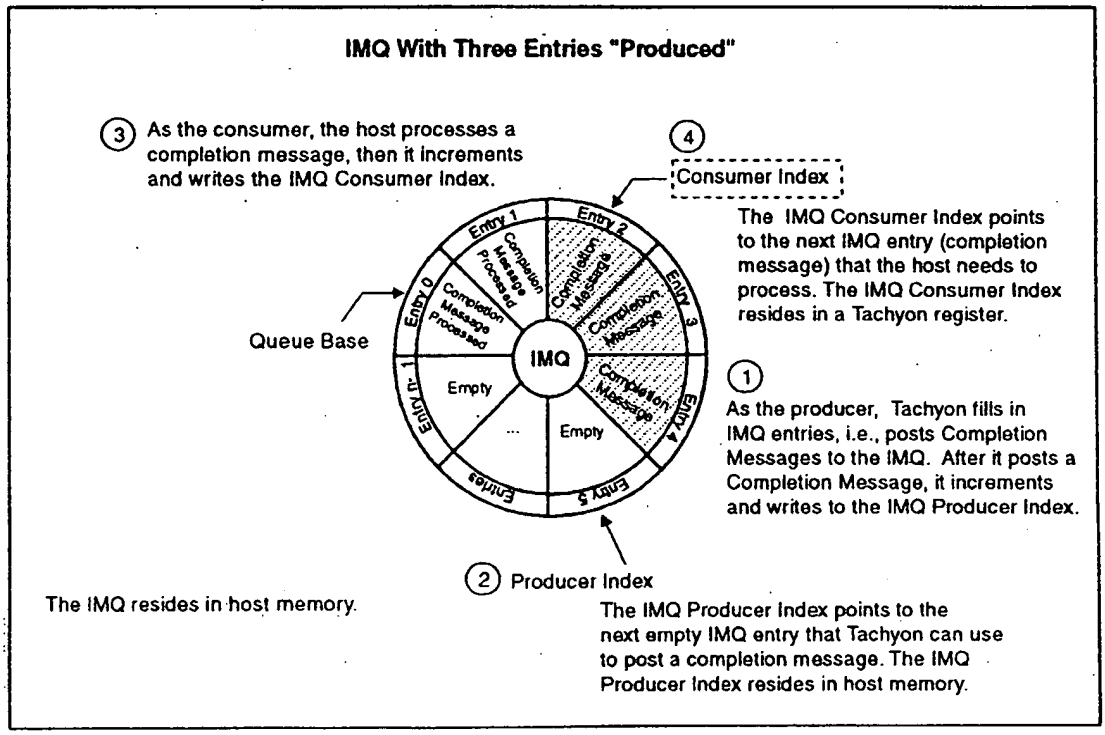

Figure 3.4 IMQ Producer and Consumer Index Example

PTI 172459

When an index points past the end of the queue (the index equals the size of the queue), the index wraps to the beginning of the queue.

The IMQ is "full" when Tachyon has filled in n-2 entries, i.e., only two empty entries remain. One of these entries is reserved for Tachyon to post an IMQ Buffer Warning completion message. This message notifies the host to provide more empty IMQ entries. All inbound and outbound processing stops until the host processes the completion messages and provides empty IMQ entries. The IMQ is "empty" when the consumer has processed all of the completion messages and the consumer index equals the producer index.

Refer to the following IMQ Full and Empty Example figure.

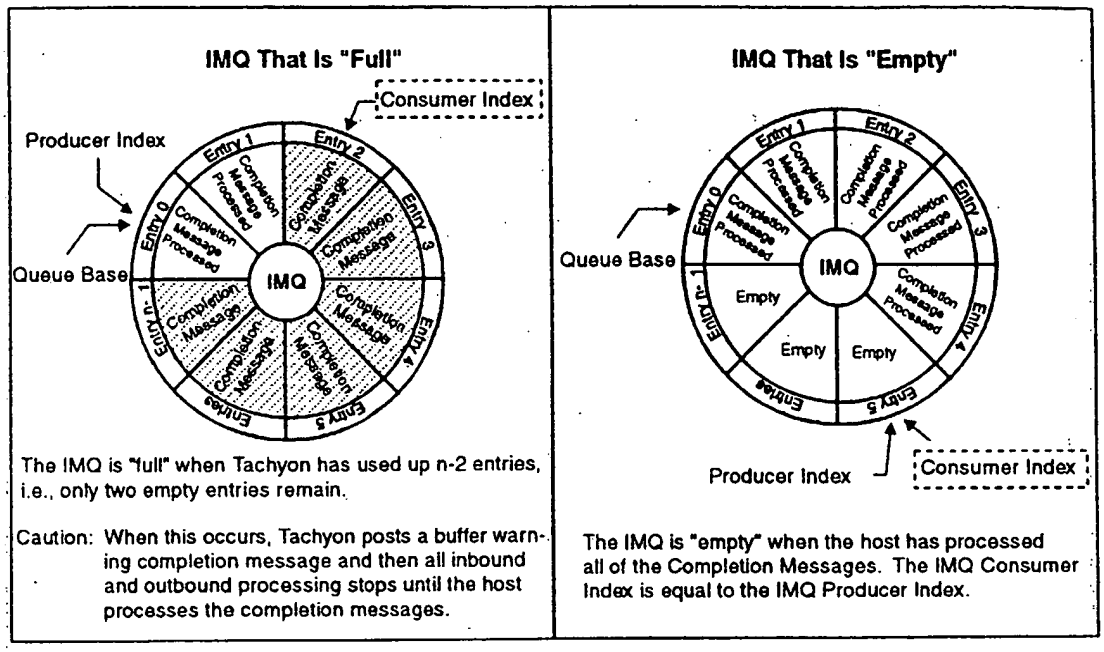

Figure 3.5 IMQ Full and Empty Example

PTI 172460

3. SFSBQ and MFSBQ

For the two inbound data buffer queues, the SFSBQ and the MFSBQ, the host is the producer and Tachyon is the consumer. As the producer, the host fills in entries with pointers to empty data buffers. These data buffers are used by Tachyon to store the incoming sequences. After the host fills in the eight pointers per entry, it increments and writes to the SFSBQ or MFSBQ Producer Index register. These producer indices point to the current end of the queue. The SFSBQ and MFSBQ Producer indices reside in Tachyon registers.

Instead of filling the entries incrementally, the host may choose to fill in all the entries in the queue at initialization. If the host chooses this method, the producer index must point to Entry n-2. For example, if the queue contains 8 entries labeled Entry 0 through Entry 7, then the producer index must point to Entry 6. When this occurs, the queue is considered "full". Refer to the following figure and the following page for more information about a "full" SFSBQ or MFSBQ.

As the consumer, Tachyon uses the pointers in the entries to store the incoming SFSs or MFSs in the designated data buffers. When Tachyon finishes storing the SFS or MFS, it writes to the SFSBQ or MFSBQ Consumer Index register with the index of the SFSBQ or MFSBQ entry that it is currently using to store the SFS or MFS. The SFSBQ and MFSBQ Consumer indices reside in Tachyon registers.

Refer to the following SFSBQ Producer and Consumer Index Example figure, which is applicable to the MFSBQ, as well.

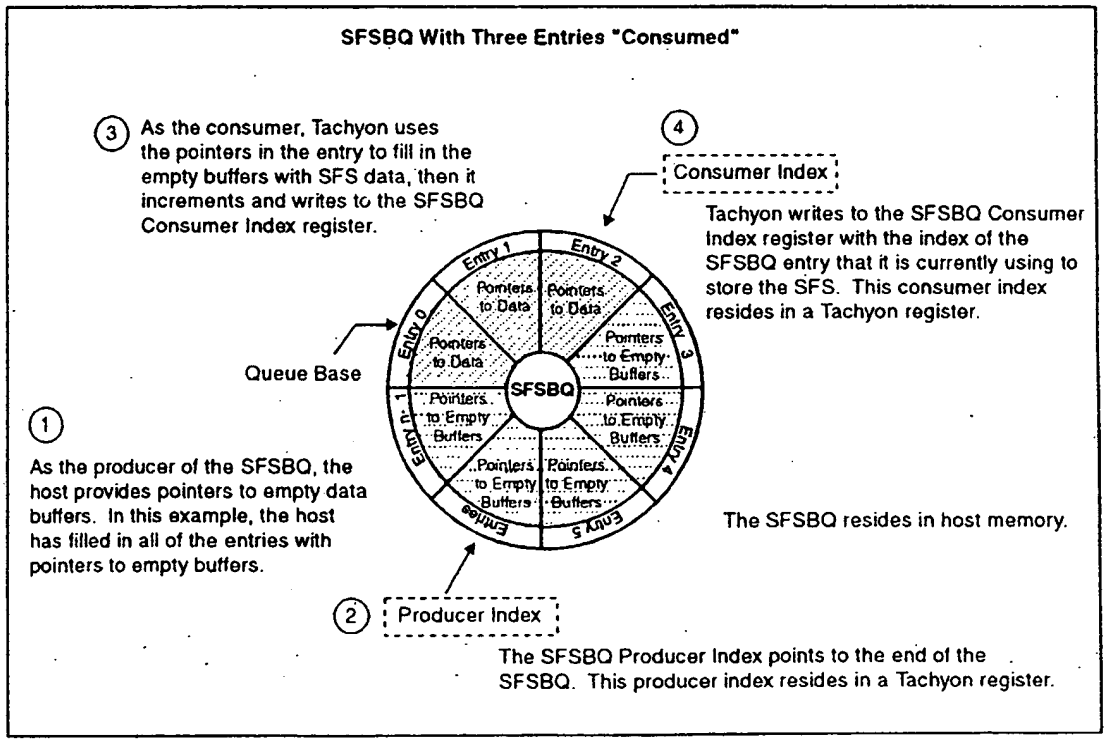

Figure 3.6 SFSBQ Producer and Consumer Index Example

PTI 172461

When an index points past the end of the queue (the index equals the size of the queue), the index wraps to the beginning of the queue.

The SFSBQ and MFSBQ are "full" when the host has filled in n entries with pointers. With a "full" queue, the consumer index points to the first entry, Entry 0, and the producer index points to the n-1 entry, Entry n-2. When the queues are "full", the data buffers of the pointers do not have any data in them.

The SFSBQ and MFSBQ are "empty" when Tachyon has used all of the available data buffers to store the SFSs or MFSs. When this occurs, the consumer index equals the producer index. Before the consumer index equals the producer index, Tachyon posts a buffer warning completion message to inform the host that it needs to provide more pointers for more empty buffers. All inbound and outbound processing stops until the host provides more pointers to empty buffers.

Refer to the following SFSBQ Full and Empty Example figure, which is applicable to the MFSBQ, as well.

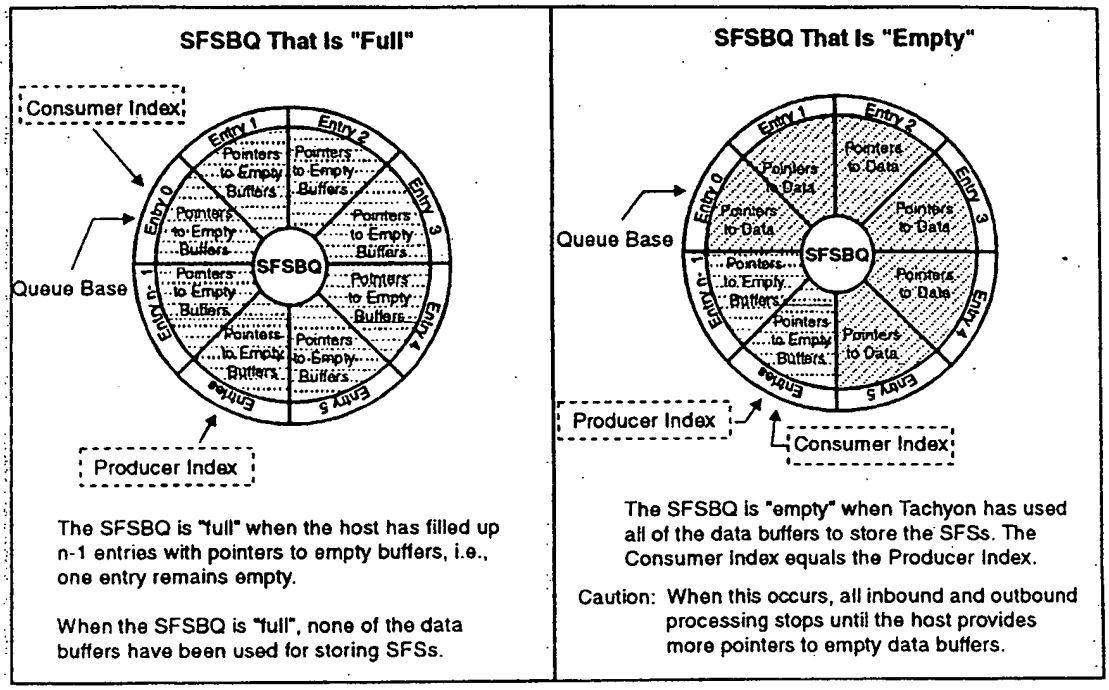

Figure 3.7 SFSBQ Full and Empty Example

PTI 172462

This table summarizes the queues' producer and consumer indices.

Circular Queue	Producer	Producer Index Resides In...	Consumer	Consumer Index Resides In...
OCQ	host	Tachyon register	Tachyon	host memory
HPCQ	host	Tachyon register	Tachyon	host memory
IMQ	Tachyon	host memory	host	Tachyon register
SFSBQ	host	Tachyon register	Tachyon	Tachyon register
MFSBQ	host	Tachyon register	Tachyon	Tachyon register

Table 3.1 Producer and Consumer Index Summary

PTI 172463

3.3 Completion Messages

Completion messages are 32-byte messages that Tachyon passes to the host via the Inbound Message Channel as an entry in the IMQ. The first word (4 bytes) of a IMQ entry indicates the type of completion message. Depending on the type of completion message, the remaining 28 bytes of the entries contain additional information or pad words.

The additional information indicates to the host details of a sequence, status of a sequence, and error information. The pad words are added to the end of all completion messages so that the completion message is 8 words (32 bytes) long. The values of these pad words are undefined.

The following table indicates the types of completion messages.

Completion Message Name	Description
outbound_	Outbound completion message
outbound_i	Outbound completion message with interrupt
out_hi_pri	Outbound High Priority completion message
out_hi_pri_i	Outbound High Priority completion message with interrupt
inbound_mfs	Inbound MFS completion message
inbound_ooo	Inbound Out of Order MFS completion message
inbound_sfs	Inbound SFS completion message
inbound_unknown_frame_i	Inbound Unknown Frame completion message with interrupt
inbound_busied_frame	Inbound Busied Frame completion message
inbound_c1_timeout	Inbound Class 1 Timeout completion message
sfs_buf_warn	SFS Buffer Warning completion message
mfs_buf_warn	MFS Buffer Warning completion message
imq_buf_warn	IMQ Buffer Warning completion message
inbound_scsi_data	Inbound SCSI Data completion message
inbound_scsi_command	Inbound SCSI Command completion message
bad_scsi_frame	Inbound Bad SCSI completion message
inbound_scsi_status	Inbound SCSI Status completion message
frame_mgr_interrupt	Frame Manager Interrupt completion message
read_status	Read Status completion message

Table 3.2 Types of Completion Messages

PTI 172464

3.3.1 Outbound Completion Messages

Completion Message Name	Description
outbound_	Outbound completion message
outbound_i	Outbound completion message with interrupt
out_hi_pri	Outbound High Priority completion message
out_hi_pri_i	Outbound High Priority completion message with interrupt

Table 3.3 Outbound Completion Messages

Description

Outbound completion messages indicate that an outbound sequence from the OCQ has been processed or has been interrupted by the occurrence of an error condition. Outbound high priority completion messages indicate that an outbound sequence (or frame) from the HPCQ has been processed or has been interrupted by an error condition.

Refer to "3.3.7 Completion Messages with Interrupts" on page 32.

3.3.2 Inbound Completion Messages

Completion Message Name	Description
inbound_mfs_	Inbound MFS completion message
inbound_ooo	Inbound Out of Order MFS completion message
inbound_sfs_	Inbound SFS completion message
inbound_unknown_frame_i	Inbound Unknown Frame completion message with interrupt
inbound_busied_frame	Inbound Busied Frame completion message
inbound_c1_timeout	Inbound Class 1 Timeout completion message

Table 3.4 Inbound Completion Messages

Description

Inbound completion messages indicate that Tachyon has received an entire sequence, a single unexpected frame, a portion of an out of order sequence, or a Class 1 time-out. The inbound completion message contains information on the success or failure of the reception, information on the buffers that store the sequence, and the information required to reassemble an OOO sequence.

To indicate a successful reception for Class 1 or Class 2, Tachyon ACKs certain received frames. Tachyon ACKs frames that have an inbound completion type of inbound_mfs_completion, inbound_ooo_completion, or inbound_sfs_completion. Tachyon does not ACK frames with an inbound completion type of inbound_unknown_frame or inbound_busied_frame.

PTI 172465

Unknown frames are frames that have FC Header fields that are not recognized by Tachyon. When Tachyon receives an unknown frame, invalid frame, or a RJT, Tachyon passes the frame to the host via the SFSBQ without processing it. These frames are passed to the host along with an inbound_unknown_frame_completion_i message.

Tachyon sends an inbound_unknown_frame_completion_i message if one of the following conditions occur:

1. If the ACK Disable bit in the Tachyon Configuration register is set. The ACK Disable bit should only be set for debug purposes.
2. If a frame with an EOFdti (invalid delimiter) is received.
3. If the SOF type of the frame is not recognized.
4. If the link control frame is not recognized (i.e., R_CTL [7..4] = link control, R_CTL [3..0] >= 0111, or the R_CTL field is not DEV_DATA, VIDEO_DATA, BSC_LNK_SERVICE, or XTD_LNK_SERVICE, or LINK_CNTL).
5. If a Class 1 data frame is received while no inbound Class 1 connection is opened.
6. If the S_ID or D_ID does not match the ID of the connected node, for Class 1.
7. If a Class 1 data frame is received while an outbound Class 1 connection is opened.
8. If an SOFc1 is received while a Class 1 connection is already open.
9. If an SOFc1 is received while Tachyon is in loop mode.
10. If a LNK_CNTL frame with incorrect Sequence Context (FCTL [22]) is received.
11. If a BSY is received that is not a BSY for the first frame of the sequence.
12. If a RJT for an outbound frame is received.
13. If an ACK is received that does not match the outbound sequence, arrives when OSM is in an unexpected state, or is Class 1 and arrives when an inbound Class 1 connection is open.
14. If a frame that contains an optional Fibre Channel Expiration Header is received (DF_CTL [6] is set to EXP_SEC_HDR).

For more information about the inbound_unknown_frame_completion_i message, refer to "3.3.7 Completion Messages with Interrupts" on page 32.

Tachyon generates an inbound_busied_frame to indicate that it has received an MFS frame while a previous MFS reassembly or Class 1 connection is in progress. No response frame is transmitted. Tachyon sets the Deferred ACK bit in the inbound_mfs_completion message indicating that the host is responsible for transmitting the final ACK. The host sets the Clear Deferred P_BSY bit in the Tachyon Control register to return to normal operation.

Tachyon generates an inbound_C1_timeout completion message to indicate that a Class 1 connection has been open and inactive for ED_TOV. Tachyon does not take further action.

Q_Index and Offset Information for Inbound Completions Messages

The second word of an inbound completion message contains a Circular Queue Index (Q_Index) and Queue Entry Offset (Offset). This word indicates the end of the current SFS or MFS. The Q_Index references the buffer queue entry that contains the address of the last data buffer that was used for this sequence or fragment. The Offset references the position of this address within this buffer queue entry. Refer to "Figure 3.8 Q_Index and Offset References for Inbound Completion Messages" on page 29. Tachyon requires that the host keeps track of the number and the order of the buffers allocated to Tachyon.

The host keeps track of the buffers that contain data for the current sequence by saving the following information:

1. The last data buffer from the previous sequence
2. The number and the order of the data buffers provided to Tachyon for the current sequence
3. The last data buffer for the current sequence.

PTI 172466

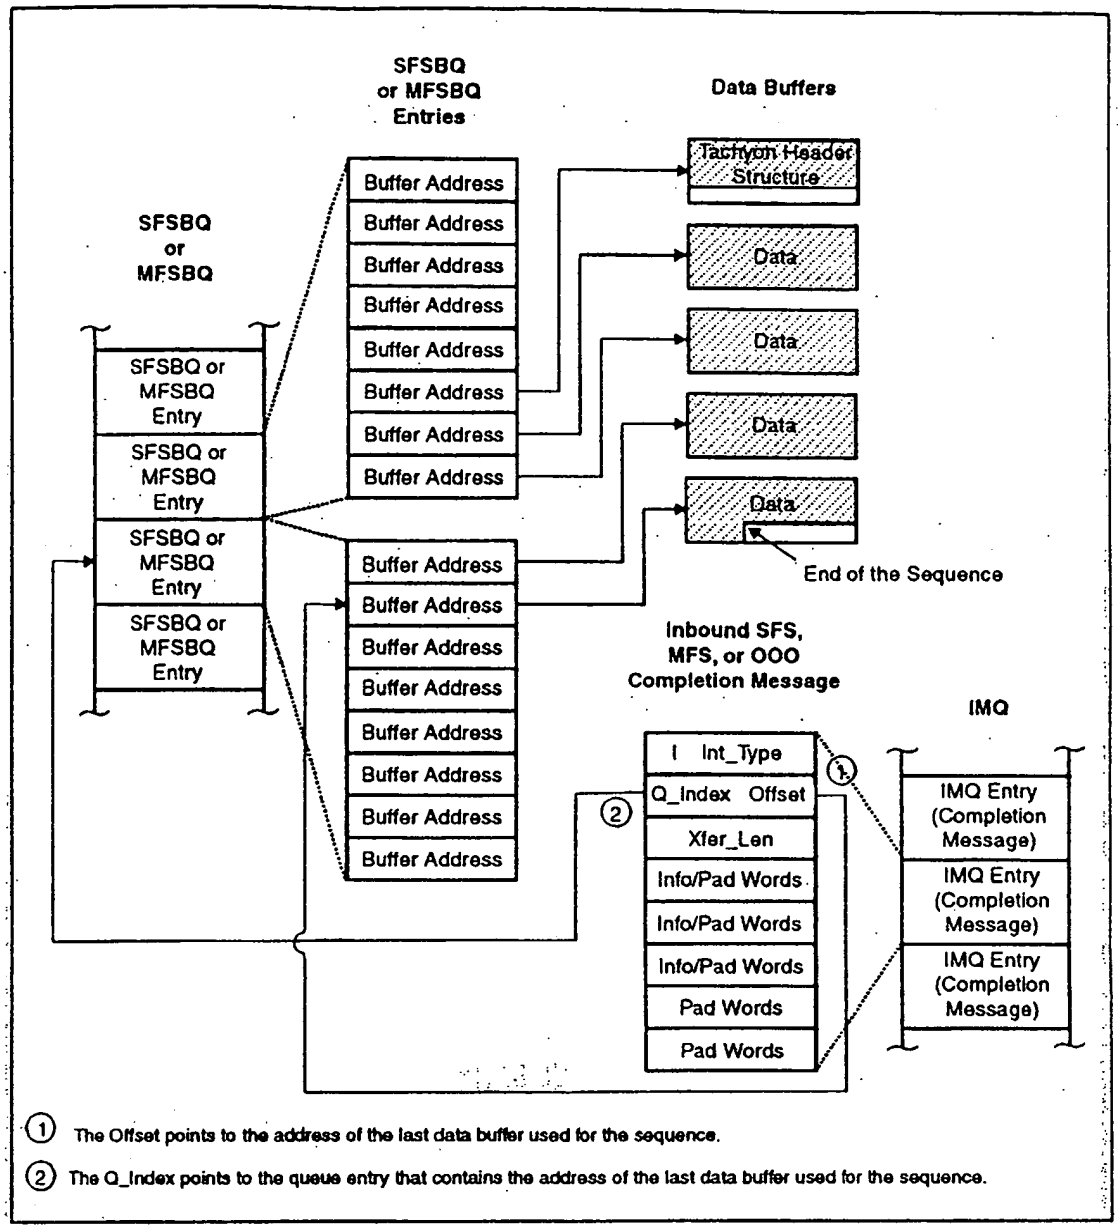

Figure 3.8 Q_Index and Offset References for Inbound Completion Messages

Tachyon generates an inbound_ooo_completion message for fragments of out of order sequences. As Tachyon receives each in order fragment of an OOO sequence, Tachyon must inform the host of the buffers that were used and the starting and stopping position of this fragment within the sequence. Tachyon does not generate an interrupt for each fragment of the sequence but waits until the last fragment of the sequence is received before generating an interrupt.

PTI 172467

3.3.3 Buffer Warning Completion Messages

Completion Message Name	Description
sfs_buf_warn	SFS Buffer Warning completion message
mfs_buf_warn	MFS Buffer Warning completion message
imq_buf_warn	IMQ Buffer Warning completion message

Table 3.5 Buffer Warning Completion Messages

Description

Tachyon generates an sfs_buf_warn or mfs_buf_warn completion message when only one SFSBQ or MFSBQ entry is available to receive data. This completion message informs the host to provide more data buffers for the incoming sequences. Because Tachyon generates the completion message when the last entry is read, but not necessarily processed, the host has some time to respond to the completion message. If Tachyon requires a data buffer and none are available, Tachyon does not send another interrupt to the host; rather, the ISM freezes, ACK generation processing stalls, and inbound and outbound sequence processing stops. To prevent this from occurring, the host must provide data buffers and write the producer index as soon as an sfs_buf_warn or mfs_buf_warn is received.

Tachyon uses the imq_buf_warn completion message to help the host manage inbound sequences. For instance, a combination of inbound OOO completion messages and outbound completion messages without interrupts could fill all the entries in the IMQ. If all of the IMQ entries contain unprocessed completion messages, then Tachyon cannot generate an interrupt to allow the host to process the completion messages. Therefore, when only two empty entries are left in the IMQ, Tachyon sends an imq_buf_warn completion message indicating that the IMQ is "full". The host must process all of the completion messages and provide empty entries, so that Tachyon can post more completion messages. The OOO completion messages are not complete sequences, and the host must temporarily save the completion information to allow it to provide the IMQ entries to Tachyon. When all of the IMQ entries are filled with unprocessed completion messages, the ISM freezes, ACK generation processing stalls, and inbound and outbound sequence processing stops until the host commits resources to empty the IMQ.

3.3.4 SCSI Completion Messages

Completion Message Name	Description
inbound_scsi_data	Inbound SCSI Data completion message
iinbound_scsi_command	Inbound SCSI Command completion message
ibad_scsi_frame	Inbound Bad SCSI completion message
iinbound_scsi_status	Inbound SCSI Status completion message

Table 3.6 SCSI Completion Messages

Description

As each phase of the Fibre Channel SCSI protocol completes, Tachyon can generate completion messages. However, it is recommended that the host use the Completion Message Disable bit in the Command Sequence and in the Outbound SEST entry. Refer to "3.3.8 Interrupt Avoidance Techniques" on page 33. This bit prevents Tachyon from generating a completion message until the target sends the Status sequence. When Tachyon sees the Status sequence, it passes the sequence and the completion message to the host, like any normal sequence reception.

Tachyon sends an inbound_scsi_command completion message with an interrupt to the host when an unsolicited command SFS (the R_CTL field in the Tachyon Header Structure is set to 0x06) is received and SCSI assists are enabled. Tachyon automatically ACKs frames that have a completion type of inbound_scsi_command completion message only if the SCSI Command Auto ACK bit in the Tachyon Configuration register is set to one.

Tachyon sends an inbound_scsi_data completion message with interrupt to the host when Tachyon is configured as a SCSI target and all frames for the SCSI sequence have been received.

Tachyon sends an inbound_scsi_status completion message when a SCSI inbound status frame (header R_CTL = 0x07) has been received and processed. The Circular Queue Index and Queue Entry Offset fields in the inbound_scsi_status completion message point to the SFS buffer that the SCSI status frame was placed into. The SFS buffer contains the FC Header which contains the OX_ID value that the host uses to locate the SEST entry associated with the SCSI transaction. Tachyon automatically ACKs frames that have a completion message type of inb_scsi_status_completion.

A bad SCSI frame is an unrecognized or unknown SCSI frame. Tachyon passes bad SCSI frames to the host via the SFSBQ. Tachyon automatically ACKs frames that have a completion message type of bad_scsi_frame only if the Bad SCSI Auto ACK bit is set in the Tachyon Configuration register.

Tachyon sends a bad_scsi_frame completion message if one of the following conditions occur:

1. Valid bit in the Inbound or Outbound SEST Entry is not set to one.
2. If Tachyon is an initiator and the SCSI Direction bit in the Outbound SEST Entry does not match the Inbound bit in the OX_ID field of the Tachyon Header Structure.
3. If Tachyon is a target and the SCSI Direction bit in the Inbound SEST Entry does not match the Inbound bit in the RX_ID field of the Tachyon Header Structure.
4. Relative Offset present bit in the FCTL field is not set to one.
5. For In Order Reassembly mode, the Expected Relative Offset in the Inbound SEST Entry does not match the received relative offset. (Tachyon receives an out of order SCSI frame during an In Order Reassembly.)
6. If Tachyon is an initiator and receives an inbound SCSI frame with an out-of-range SCSI OX_ID (SCSI OX_ID greater than the length of the SEST).
7. If Tachyon is a target and receives an inbound SCSI frame with an out-of-range SCSI RX_ID (SCSI RX_ID greater than the length of the SEST).

PTI 172469

3.3.5 Frame Manager Interrupt Completion Message

The Frame Manager may generate frame_mgr_interrupt completion messages to the host when certain link configuration changes or errors occur. The frame_mgr_interrupt completion message informs the host that it should read the Frame Manager Status register to determine any configuration changes or errors that have occurred. Depending on the situation, the host may need to write to the Frame Manager Control register to release an error or to force the Frame Manager to change states. If needed, the host should write to the Frame Manager Control register to allow transition from link failure. This interrupt process happens as part of the Link Initialization process and any time the link has been broken.

3.3.6 Read Status Completion Message (Debugging)

Tachyon contains a debugging feature that may be used during hardware or software bringup.

When the host driver sets the Status Request bit in the Tachyon Control register to one, Tachyon writes a block of data (Inbound Read Status Frame Structure), which contains information about Tachyon's most important internal states at that moment, to the next available buffer in the SFSBQ. Tachyon then generates a read_status completion message which points to that information.

3.3.7 Completion Messages with Interrupts

Tachyon generates an interrupt by asserting the TSI Interrupt signal (INT_L) low after updating the IMQ Producer Index. Refer to "7.7.10 Interrupt Signal" on page 243.

Tachyon generates an interrupt with every completion message (except the inbound_ooo_completion message and the inbound_busied_completion message) unless the host programs it not to generate an interrupt and the operation is successful. If the operation is not successful, Tachyon generates an interrupt even if the host programmed it not to generate an interrupt.

It is the host's responsibility to process all the IMQ entries up to the location where its copy of Tachyon's IMQ Producer Index points. If the host does not read all the entries up to its copy of Tachyon's IMQ Producer Index at the time that the interrupt was received, then the host may not see the completion message that caused Tachyon to send the interrupt.

The host informs Tachyon that it has serviced the completion messages by writing to the IMQ Consumer Index register. Tachyon continues to generate an interrupt each time the host updates the IMQ Consumer Index, until the IMQ Consumer Index equals the IMQ Producer Index, i.e. until the host has serviced all the outstanding completion messages.

Tachyon can help the host avoid unnecessary interrupts for both inbound and outbound completions The host is responsible for recognizing non-interrupting completion messages that are followed by a message that requires an interrupt. As an example, while processing an inbound OOO sequence, Tachyon may have generated several completion messages before receiving the entire sequence. At this point, Tachyon may process an outbound completion message that requires an interrupt. The host receives the interrupt and notices that several inbound completion messages are stored before the outbound completion message, but the final inbound completion message is not in the IMQ yet. The host must save the information for the inbound completion messages while waiting for the final inbound completion message.

3.3.8 Interrupt Avoidance Techniques

Tachyon provides three techniques to reduce the number of interrupts sent to the host.

The first technique uses the following bits:

1. Completion Message Disable bit (Word 2, bit 26 in the ODB and the HPDB or Word 1, bit 10 in the Outbound SEST Entry)
2. Completion Message Interrupt Disable bit (Word 2, bit 25 in the ODB and the HPDB or Word 1, bit 9 in the Outbound SEST Entry)

The host sets the Completion Message Disable bit to completely disable the generation of a completion message. This feature is recommended for servicing SCSI requests when the host is the initiator. Tachyon manages the Command and Data phases of the transfer without notifying the host until the Status sequence arrives from the device. At this time, the host can manage the entire exchange with only one interrupt. Tachyon ignores this feature if an error occurs during the operation. In this case, Tachyon generates a completion message and interrupt that indicates the error.

The host sets the Completion Message Interrupt Disable bit to indicate that Tachyon should not generate an interrupt for the completion message. The completion message is posted to the IMQ but the interrupt is not generated. It is assumed that a subsequent transmission requires an interrupt, and the host may prefer to see all of the completion messages at one time. The host may use this feature when it posts several sequences at a time and only wants to know when all of them have completed. For example, when a SCSI target device sends a FCP_XFER_RDY followed by the data sequence, the host may not want to receive an interrupt until both are received.

If the host never posts a command for Tachyon to generate an interrupt, Tachyon does not generate an interrupt on its own. There are no timeout mechanisms for interrupts.

A second technique to reduce interrupts is explained in the following example. Suppose Tachyon posts a completion message to the host and generates an interrupt. If Tachyon generates more messages to the host before the host has had a chance to process the first interrupt and acknowledge it by writing to the Inbound Message Channel's consumer index, interrupts for these added messages do not have to be generated. This means that only one interrupt is pending to the host at any time. It also means that the host can tell if more messages need to be processed without taking the Interrupt Service Routine (ISR) startup penalty. The host compares the IMQ producer index in host memory to the host's copy of the consumer index before exiting its ISR.

A third technique is used during the processing of inbound OOO sequences. As Tachyon processes each in order fragment of an OOO sequence, it generates a completion message, but delays the interrupt until the entire sequence has been processed.

3.4 Acknowledgements (ACKs)

Two acknowledgment models are supported by Tachyon to transfer sequences: the ACK_1 model and the ACK_0 model. The ACK_1 model requires that one acknowledgment frame is returned for each corresponding frame sent. The ACK_0 model requires that one acknowledgment frame is returned when the entire sequence has been transferred.

3.4.1 Acknowledgement of Transmitted Frames

When transmitting frames, Tachyon supports both the ACK_1 and ACK_0 models. Tachyon uses the ACK_0 bit in the ODB to determine which ACK model to use for proper flow control. The host determines whether or not to set the ACK_0 bit from the login parameters exchanged with the node with which it is communicating.

If the host directs Tachyon to use the ACK_1 model, it expects to receive ACKs for each frame transmitted to provide End-to-End flow control. If the host directs Tachyon to use the ACK_0 model, Tachyon sends the entire sequence without End-to-End flow control, and then waits for the single ACK to return.

If the ACK does not match the current outbound sequence, Tachyon sends the frame to the host as an inbound unknown frame with no effect on the OSM.

3.4.2 Acknowledgement of Received Frames

Tachyon relieves the host of many data communication management tasks. One of the largest of these tasks is the generation of ACKs for Class 1 and Class 2 communication. Class 3 frames, by definition, are not acknowledged.

ACKs for Class 1 frames have priority over ACKs for Class 2 frames in the ACK FIFO. Therefore, regardless of the order in which the ACKs are queued, Tachyon transmits ACKs for Class 1 frames before it transmits the ACKs for Class 2 frames.

Tachyon uses the F_CTL field ACK_Form assistance bits defined by FC-PH-2.

The sequence initiator sets the F_CTL bits appropriately to tell Tachyon which ACK model to use when sending ACKs for frames received. These F_CTL bits specify the ACK_1, ACK_0, or ACK_N models.

The following table shows the mapping of the F_CTL bits to the ACK model requested.

F_CTL bit 13	F_CTL bit 12	Description
0	0	No assistance
0	1	ACK_1 model
1	0	ACK_N model
1	1	ACK_0 model

Table 3.7 F_CTL bits 12 and 13

If the initiator sets ACK_N in the F_CTL field and the ACK Generation Assist Enable bit is set to one, Tachyon defaults to the ACK_0 model; otherwise, Tachyon defaults to the ACK_1 model.

PTI 172472

ACK_1 Model for Received Frames

As a sequence recipient, Tachyon supports the required default ACK_1 model. If the initiator does not indicate a preference, or indicates that ACK_1 is to be used, Tachyon responds to each valid frame received with an ACK_1 frame. Tachyon generates the ACK from a copy of the received Fibre Channel header.

To generate an ACK, Tachyon modifies the following inbound FC Header information.

1. In the F_CTL field of the header:
   a. The Exchange Context and Sequence Context bits (23, 22) are inverted.
   b. The Sequence Retransmission bit (9) is cleared to zero.
   c. The Abort Sequence Condition bits (5, 4) are cleared to zero, unless an error is detected.
   d. For Class 1, the Unidirectional bit (8) is set to one.
2. The D_ID and S_ID fields in the header are swapped.
3. The R_CTL field is changed to indicate an ACK frame.

ACK_0 Model for Received Frames

As a sequence recipient, Tachyon also supports the ACK_0 model. To use ACK_0, the ACK Generation Assist Enable bit in the Tachyon Configuration register is set to one. The sequence initiator specifies, by setting the ACK_Form bits in the F_CTL field of the frame header, that ACK_0 is to be used as described in "FC-PH-2 Section 18.5".

Tachyon follows the ACK_0 rules defined in "FC-PH, Section 20.3.2.2 (ACK_0)" except for rule "a)", which states "If ACK_0 is supported by both Sequence Initiator and Sequence Recipient, a single ACK_0 per Sequence shall be used to indicate successful Sequence delivery or to set Abort Sequence bits. An additional ACK_0 shall be used within a Sequence to perform Sequence Interlock." Tachyon, as a sequence recipient, does not know that X_ID Interlock is requested (because Tachyon does not store Login parameters). Therefore, Tachyon does not send an ACK_0 frame for the first frame of the first sequence of an exchange.

ACK_N Model for Received Frames

Tachyon does not support the ACK_N model.

3.5 Busy Responses, Rejects, and ACK Aborts

3.5.1 BSYs, RJTs, and ACK_ABTs of Transmitted Frames

When transmitting frames, Tachyon manages the Busy responses in several ways.

1. If the BSY is in response to the first frame of the sequence and both the Sequence Interlock (lck) and the Start Class 1 Connection (SOFc1) bits in the ODB are not set to one, then:
   a. BSYs are sent to the host without any effect on the state of the OSM.
2. If the BSY is in response to the first frame of the sequence and the Sequence Interlock (lck) bit is set in the ODB then:
   a. Tachyon re-transmits the frame up to 16 times. If the frame has been busied 16 times, then Tachyon terminates the sequence with a Retries Exceeded error (the X bit is set to one in the outbound_completion message). The OSM freezes.
   b. If the first frame of the sequence is busied and the Retry Disable bit is set in the Tachyon Configuration register, then Tachyon terminates the sequence with a Retries Exceeded error (the X bit is set to one in the outbound_completion message). The OSM freezes.
3. If the BSY is in response to the first frame of the sequence and the Start Class 1 Connection (SOFc1) bit is set in the ODB then:
   a. Tachyon re-transmits the frame up to 16 times. If the frame has been busied 16 times, then Tachyon terminates the sequence with a Retries Exceeded error (the X bit is set to one in the outbound_completion message). The OSM freezes.
   b. If the first frame of the sequence is busied and the Retry Disable (rd) bit is set in the Tachyon Configuration register, then Tachyon terminates the sequence with a Retries Exceeded error (the X bit is set to one in the outbound_completion message). The OSM freezes.
   c. If a non-stacked fabric exists (the st bit in the Tachyon Configuration register is zero) and Tachyon detects an inbound Class 1 connection opening after the frame is sent, it treats this condition as a BSY frame received condition and re-transmits the SOFc1 when the inbound Class 1 connection closes. Tachyon re-transmits the SOFc1 because a non-stacked fabric does not send a BSY if it has forwarded an SOFc1.
4. If the BSY is in response to a frame other than the first frame of the sequence, then:
   a. Tachyon sends BSYs to the host with no effect on the state of the OSM.

If a RJT is received and matches the current outbound sequence, the RJT frame is sent to the host with an inbound_unknown_frame completion message. The outbound sequence is terminated and an outbound_completion message is sent to the host indicating the RJT.

If an ACK_ABT is received and matches the current outbound sequence, the outbound sequence is terminated and an outbound_completion message is sent to the host indicating the ACK_ABT. The ACK_ABT frame is not sent to the host.

If a RJT or ACK_ABT is received that does not match the current outbound sequence, the RJT or ACK_ABT frame is sent to the host with an inbound_unknown_frame completion message. This does not affect the state of the OSM.

PTI 172474

3.5.2 Busy (P_BSY) and Reject (P_RJT) Responses of Received Frames

When receiving frames, Tachyon may need to send a negative response to a received frame. This negative response can either be a busy (P_BSY) or a reject (P_RJT) frame.

Tachyon generates P_BSYs under the following conditions.

1. Sequence Terminated, N_Port resource busy; P_BSY Parameter Field: 0x01030000

   Tachyon generates this P_BSY in response to an SOFc1 frame that is received when directly connected to another N_Port and an SOFc1 was already sent but not acknowledged, or a connectionless multiframe sequence is currently being processed.

2. Sequence Active, N_Port resource busy; P_BSY Parameter Field: 0x02030000

   Tachyon generates this P_BSY in response to a connectionless multiframe sequence data frame that cannot be processed at the current time, because a Class 1 connection is open or another connectionless multiframe sequence is currently being processed.

Tachyon does not automatically P_RJT frames because it does not store login or exchange information and, therefore, cannot perform login or exchange validation. Tachyon passes frames with questionable characteristics to the host as unknown frames. The host is responsible for inspecting and rejecting, if necessary, these frames.

3.6 Transmit Details

3.6.1 Descriptions for the Transmit Process Blocks

Refer to "Figure 3.1 Tachyon Internal Block Diagram" on page 17. The host-based and Tachyon internal blocks for the transmit process are described briefly below.

Block	Location	Description
OCQ	Host-based	The host uses the Outbound Command Queue (OCQ) to issue commands that instruct Tachyon to transmit sequences with normal priority. Each OCQ entry contains an ODB.
HPCQ	Host-based	The host uses the High Priority Command Queue (HPCQ) to issue commands to Tachyon and bypass OCQ traffic to transmit single frames with high priority. The main purpose of using the HPCQ is to provide error recovery and to send time sensitive frames that must be sent before ED_TOV expires. The host must use the HPCQ when sending Link Control frames, for example, ACKs, P_BSYs, P_RJTs, etc. Each HPCQ entry contains a High Priority Descriptor Block (HPDB).
Outbound Message Channel	Tachyon	The Outbound Message Channel Outbound Command manages the OCQ. It is also known as the Outbound Message Channel.
High Priority Message Channel	Tachyon	The Outbound Message Channel High Priority Command manages the HPCQ. It is also known as the High Priority Message Channel.
Outbound Block Mover	Tachyon	The Outbound Block Mover DMAs outbound data from host memory to the OSM.
OSM	Tachyon	The Outbound Sequence Manager (OSM) manages the Fibre Channel protocol that sends an entire sequence of outbound data, including processing all ACK frames and error handling. The OSM programs the Outbound Message Channels to retrieve a data sequence from host memory and segment it into individual frames for transmission. The OSM fairly arbitrates between processing ODBs from the OCQ and SCSI hardware assists.
Outbound Frame FIFO	Tachyon	The Outbound Frame FIFO buffers data before transmission to prevent underrun. This FIFO is sized to hold one frame. As Tachyon sends the current frame onto the link, the Outbound Frame FIFO is simultaneously filled with the next frame, maximizing outbound performance.
ACK FIFO	Tachyon	The ACK FIFO holds a maximum of eight ACKs until they can be sent out.

Table 3.8 Transmit Process Blocks

PTI 172476

Block	Location	Description
Frame Manager	Tachyon	The Frame Manager is Tachyon's front end. It is responsible for the FC-1 functions of transmitting and receiving Fibre Channel frames. It is capable of generating interrupts to the host when certain link configuration changes occur to which the host must respond. The interrupt process occurs as part of Link Initialization and any time the link has been broken. The Frame Manager responds to a reset condition by initializing all of its registers to their default values.
OS/CRC Generator	Tachyon	The Ordered Set / CRC Generator encapsulates data into FC-1 frames, generates a 32-bit Cyclic Redundancy Check (CRC) and writes it into the frame, and passes the encapsulated data to the 16B/20B Encoder.
16B/20B Encoder	Tachyon	The 16B/20B Encoder converts outbound 16-bit wide data into two 8-bit pieces, each of which is encoded into a 10-bit transmission character using the 8B/10B encoding algorithm.
20B/10B MUX	Tachyon	The 20B/10B Multiplexer is responsible for selecting the proper width, either 10 bits or 20 bits, of the data path for the specific type of Physical Link Module (PLM) being used. The data width is specified in the Parallel ID field of the PLM interface. A 20-bit data width is used with a 100 Mbyte/sec link speed. All other link speeds transmit 10-bit wide data.

Table 3.8 Transmit Process Blocks (Continued)

PTI 172477

3.6.2 Outbound Command Queue Transmit Details

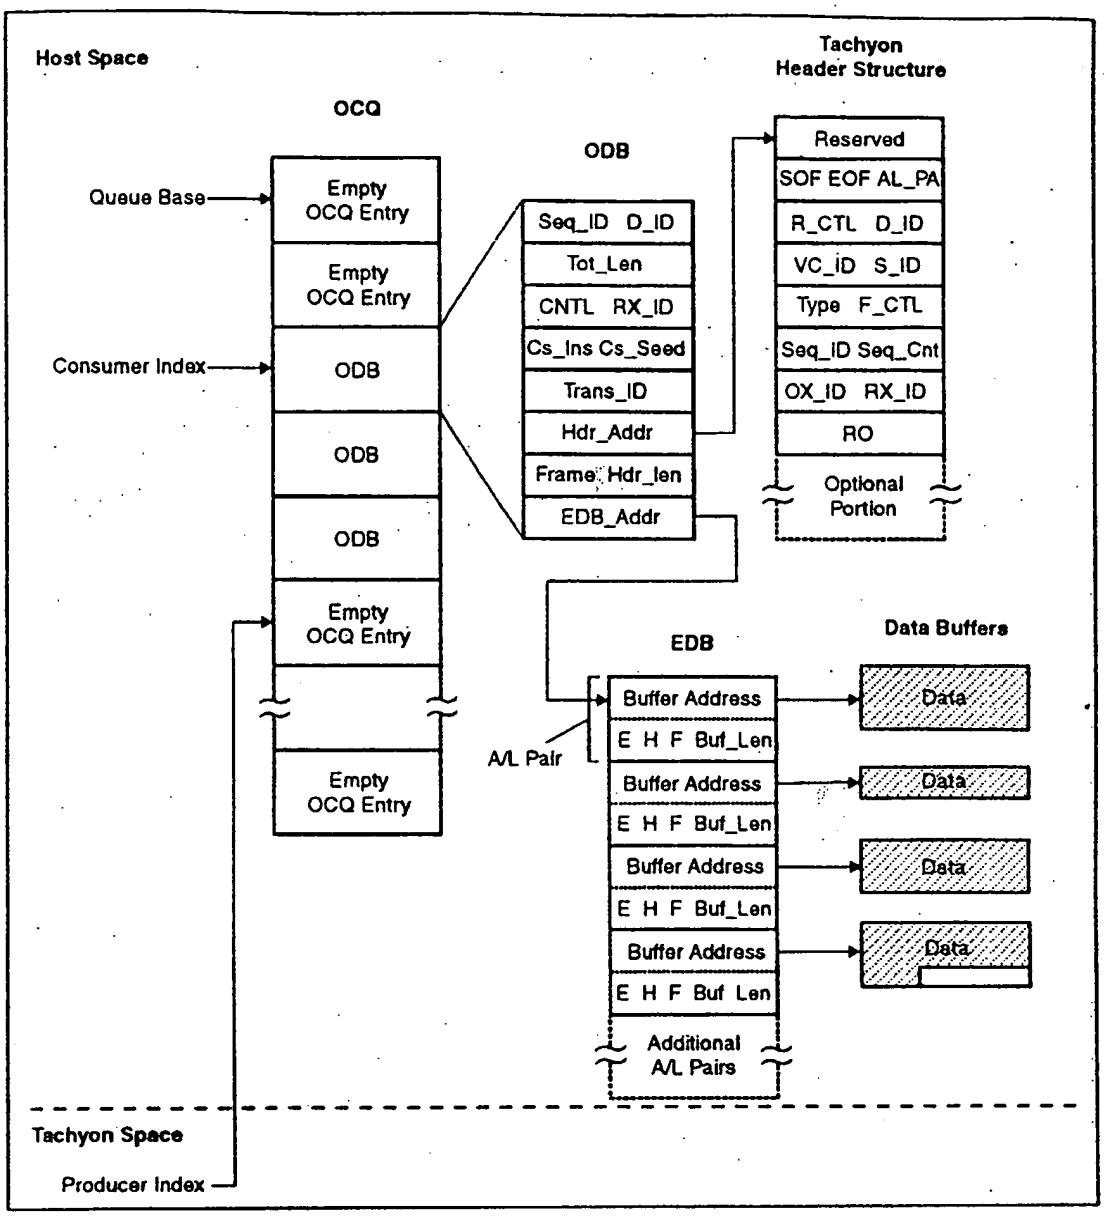

Figure 3.9 Transmit Process via the OCQ

PTI 172478

The host controls outbound sequences (except high priority single frame sequences) with the OCQ. Each outbound sequence is defined by an 8-word entry in the OCQ, called the ODB. The ODB:

1. Defines sequence size, framing, and flow control.
2. Points to the EDB.
3. Points to the Tachyon Header Structure.
4. Specifies Fibre Channel protocol options for the sequence.
5. Controls TCP/UDP checksum assist parameters.
6. Controls outbound_completion messages for the sequence.

The EDB consists of an 2-word long Address/Length (A/L) Pair(s) that defines data buffers in host memory. The first word of the A/L Pair is the Buffer Address that defines the beginning of the buffer. The second word of the A/L Pair contains the Buffer Length that defines the number of bytes in the buffer. The EDB contains as many A/L Pairs as are necessary to completely define the data payload of the sequence.

The Tachyon Header Structure defines all of the Fibre Channel sequence initial header field values not defined in the ODB.

Tachyon processes one single outbound sequence at a time by:

1. Reading the ODB
2. Updating the OCQ Consumer Index
3. Reading the Tachyon Header Structure
4. Reading the EDB and the associated data

The host may send a single frame sequence without using the EDB by defining the entire frame with payload in the Tachyon Header Structure. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47. In this case, the End (E) bit in the ODB must be set to one to indicate that the Tachyon Header Structure contains the entire sequence and an EDB does not exist.

Tachyon waits for all ACKs from the remote node before processing a subsequent sequence.

As Tachyon transmits frames for the outbound sequence, the fabric or the remote node may detect a problem in managing the frame(s). If the problem is non-serious, the fabric or the remote node may respond to the frame with a BSY. The BSY informs Tachyon to either retry sending the frame or stop sending the frame at this time and retry sending it later.

Tachyon provides a feature to allow it to manage the retries for the first frame of a sequence that gets busied. The host enables this feature by setting one of the interlocking bits (SOFc1 or Lck) to one in the CNTL field of the ODB. When the host enables the interlocking feature, Tachyon sends the first frame and waits for an ACK. The ACK indicates that the remote node has the resources to handle a new sequence. If Tachyon receives a BSY instead, Tachyon automatically sends the first frame again. Tachyon can retry 16 times to send the frame. If the Tachyon receives a BSY on the 16th try, Tachyon terminates the sequence by setting the Retries Exceeded bit in the outbound_completion message.

If the remote node returns a BSY for any other frame but the first one when the host enables the sequence interlock, or on the first frame if interlocking is not enabled, Tachyon passes the BSY to the host to be processed. Tachyon does not have the arithmetic hardware necessary to retrace the DMA chains in host memory to determine where the data is located for the busied frame. Tachyon relies on the host to determine if it wants to retry the frame, to calculate where the data is located, and then to resend the frame using the High Priority Channel.

If the fabric or the remote node detects a serious problem with a frame, either the fabric or the remote node may respond to the frame with a RJT. If Tachyon receives a RJT frame in response to a transmitted frame, it passes the RJT frame to the host, aborts transmitting the sequence, sends a completion message to the host, freezes the OSM, and waits for the host to complete its error recovery procedures. The host is responsible for managing the RJT frame.

PTI 172479

Interweaving

Interweaving allows the host software to intermix frames of multiple concurrent sequences. To enable interweaving, the host must set the Continue Sequence bit in the CNTL field of the ODB to one to indicate whether the End Of Sequence bit or EOFt (for Class 3 sequences) should be set in the last frame of the requested data transfer. The host can send the transferred data, or "sequence fragments," with multiple ODBs rather than sending a complete sequence with only one ODB. In this manner, the host can program many frames from different sequences to interweave through Tachyon.

As an example, interweaving may be used in some target applications. When large amounts of data slowly arrive into buffers from several mechanisms, sequence fragments can be forwarded as they arrive instead of each being buffered as an entire sequence before being forwarded.

When sending a sequence with the Continue Sequence bit set to one, Tachyon does not automatically update the following fields in the Tachyon Header Structure.

1. E_C bit of the F_CTL field will not be set on the last frame of the sequence fragment if the E_C bit is set in the ODB.
2. Fill bits of the F_CTL field of the last frame will not assume the values in the ODB.
3. End of Sequence bit in the F_CTL field will not be set to one on the last frame of the sequence fragment.
4. EOFt will not be appended to a Class 3 frame.

When the host sends a sequence with the Continue Sequence bit set to one (i.e., interweaving is used).

1. Fibre Channel time-out rules may be violated, specifically the maximum time between frames, the Error Detect Time-Out Value (ED_TOV).
2. ACK_0 model is not supported. A sequence time-out occurs if both the Continue Sequence bit and the ACK_0 bit are set to one in the ODB.

3.6.3 Outbound Class 1 Sequences

When the host wants Tachyon to send a Class 1 sequence, it builds the ODB and all associated data structures. The host posts the ODB to the OCQ entry referenced by the producer index. The host then updates the producer index by writing the new index value to the producer index register in Tachyon. When Tachyon notices that the producer index has changed and that the outbound channel is available, it DMAs the ODB from the host memory location referenced by the consumer index. The consumer index then increments. Tachyon starts acting upon the request by first comparing the connection state with the ODB SOFc1 bit. If the SOFc1 bit indicates that the host is requesting a connection and one is already established, or if it assumes a connection is already established, Tachyon freezes the OSM and sends an outbound_completion message with an error bit set to the host. The OSM remains frozen until the host completes error recovery and Tachyon resumes normal OSM operation. Tachyon resumes normal OSM operation when the host writes to the Tachyon Control register to unfreeze the OSM.

If the host requests to send a Class 1 sequence and a Class 1 connection does not exist before the sequence is posted, the host must create the Tachyon Header Structure for the sequence with the Initial Start of Frame set to SOFc1 and the SOFc1 bit in the ODB is set to one. Tachyon sends the Tachyon Header Structure as the connect request (SOFc1). The Header Length must not exceed the length of the buffer-to-buffer receive size of the remote node determined during login.

The host driver must determine whether a connection exists or not, and set the ODB SOFc1 bit and the Tachyon Header Structure SOF field with the appropriate values.

If the host requests to send a Class 1 sequence and a Class 1 connection is already established, the host must create the Tachyon Header Structure for the sequence with the Initial Start of Frame set to SOFi1 and the SOFc1 bit in the ODB cleared to zero.

PTI 172480

Independent of the Maximum Frame Length (Frame_Len) value specified in the ODB, the entire amount of data, indicated by the Hdr_Addr and Hdr_Len fields, is sent in the first frame. Therefore, the host must ensure that the Header Length in the ODB is less than the maximum frame size that the responder is able to accept so the responder does not overflow its FIFOs.

As Tachyon reads the Tachyon Header Structure from the host to send the first frame, Tachyon copies the Tachyon Header Structure to internal registers for use in generating Fibre Channel headers for subsequent frames. The D_ID (Fibre Channel address for destination node), RX_ID (Receiver exchange ID to be used for the current sequence), and SEQ_ID (Fibre Channel sequence identifier) fields are present in both the ODB and the Tachyon Header Structure. Tachyon uses the D_ID and RX_ID fields from the ODB and the SEQ_ID from the Tachyon Header Structure in generating the header for each frame. However, for future compatibility, it is recommended that the three fields in the ODB and the Tachyon Header Structure match. Tachyon assumes consistency between the ODB and the Tachyon Header Structure and therefore does not check for mismatches.

Once Tachyon sends the first frame of the sequence, it determines if either the Sequence Interlock bit or the SOFc1 bit is set to one in the CNTL field of the ODB. If either bit is set to one, Tachyon waits until the ACK is received for that frame before continuing. Tachyon then uses the RX_ID value returned in the ACK and substitutes it into the Fibre Channel Header on all subsequent frames of this sequence.

If the Sequnce Interlock bit is cleared to zero, or if the Sequence Interlock bit is set to one and the ACK received, Tachyon continues to DMA data from the host in frame sized blocks and sends them with headers automatically generated from the previously stored header. Tachyon increments the SEQ_CNT and also tracks and inserts the correct Relative Offset field. Tachyon always fills the parameter field of the Fibre Channel header with the current Relative Offset value. If the host does not want the remote node to interpret the parameter field, it must clear bit 3 of the F_CTL field.

As Tachyon sends the frames for the sequence, it also tracks the End-to-End credit (EE_Credit). EE_Credit contained in the CNTL field of the ODB determines the number of frames that Tachyon can send to the remote D_ID without receiving an ACK. Each time Tachyon sends a frame, EE_Credit decrements. Each time Tachyon receives an ACK, EE_Credit is incremented. If EE_Credit goes to zero, frame transmission stops, and an ED_TOV starts. The ED_TOV timer continues until ACKs arrive. If ACKs arrive, transmission resumes. If the ED_TOV timer expires, Tachyon sends an outbound completion message to the IMQ with one of the time-out bits (OT or AT) set to one. Then the OSM freezes.

When it is time to send the last frame of the sequence, Tachyon checks to determine if the E_C bit in the CNTL field in the ODB is set to one. If it is, Tachyon sends the last frame with the E_C bit set to one. This indicates that the connection should be terminated by the remote node. If the E_C bit is not set to one, Tachyon sends the last frame with the E_C bit cleared to zero, which leaves the connection established when the sequence has completed. Tachyon also sets the EOS bit to one on the last frame of the sequence. Only multiframe sequences can be sent when the SOFc1 and E_C bits are both set to one in the CNTL field of the ODB.

When Tachyon receives the last ACK for the sequence, it sends an outbound_completion message to the host. This tells the host that it can de-allocate all memory associated with this ODB and inform any processes waiting on its completion. A part of the outbound completion message is the transaction_id which the host passed down in the ODB. The host may use this to match the completion with the send request.

If a sequence terminates abnormally (an abort or a time-out occurs), Tachyon sets the appropriate error bit in the status field of the outbound_completion message.

The remote node can send a Class 1 frame to Tachyon in an attempt to close the Class 1 connection that Tachyon opened. Tachyon accepts this frame and sends it to the host via the SFSBQ along with an inbound_unknown_frame completion message. The host must recognize that the E_C bit in the F_CTL field of the received frame is set to one so that the host can generate a link reset. Tachyon can then close the outbound Class 1 connection.

PTI 172481

Class 1 Error Recovery

Tachyon does not allow bi-directional data transfers in Class 1 connections. Thus, the only time an N_Port, as the recipient of a Class 1 connection, may send frames is for error recovery. The host sends error recovery frames through the High Priority Chanel.

To ensure that flow control is not violated during Class 1 Error Recovery, the following rules apply:

1. The host should send only one outstanding frame at a time. If more than one frame is sent, Tachyon receives additional ACKs, which may consume inbound buffer resources. If inbound buffer resources are depleted, inbound frames may be discarded.
2. The host must process the ACK for the error recovery frame before sending any additional error recovery frames.

Host Managed Class 1 Connections

The host controls the state of the connection that is established with the remote node. The E_C field within the ODB tells Tachyon whether to retain an established connection for subsequent sequences or to terminate the connection at the end of the current sequence. Because of this, the driver must not mix Class 1 sequences for another D_ID while a connection is established when posting new sequences to Tachyon. The host may post a Class 2 or Class 3 sequence when it has a connection open. Tachyon sends this sequence using the intermix mode of operation.

While a Class 1 connection is held open by the host for multiple sequences, SCSI assisted sequences are blocked from being sent until the connection is closed. Therefore, for SCSI transactions, the Class 1 connection must be closed to resume progress.

3.6.4 Outbound Class 2 Sequences

Class 2 sequences are managed in the same manner as Class 1 sequences, except the host does not need to manage connections or the E_C bit of the ODB.

3.6.5 Outbound Class 3 Sequences

The host manages Class 3 sequences in the same way as Class 2 sequences, except that Class 3 does not use EE_Credit flow control. When Tachyon sends the frames of the Class 3 sequence, they are transmitted as fast as Buffer-to-Buffer Credit (BB_Credit) flow control allows. Also, since Class 3 does not use EE_Credit, ACKs are not received for the frames sent. Therefore, once Tachyon transmits the sequence, the OSM immediately generates the completion message and sends it to the host. Tachyon automatically sends the last frame of a Class 3 sequence with EOFt set to indicate that the sequence is complete.

3.6.6 Transmitting L_Port Open Broadcast Replicate Frames

Refer to "3.9.7 Open Broadcast Replicate Support" on page 86.

PTI 172482

3.6.7 High Priority Command Queue Transmit Details

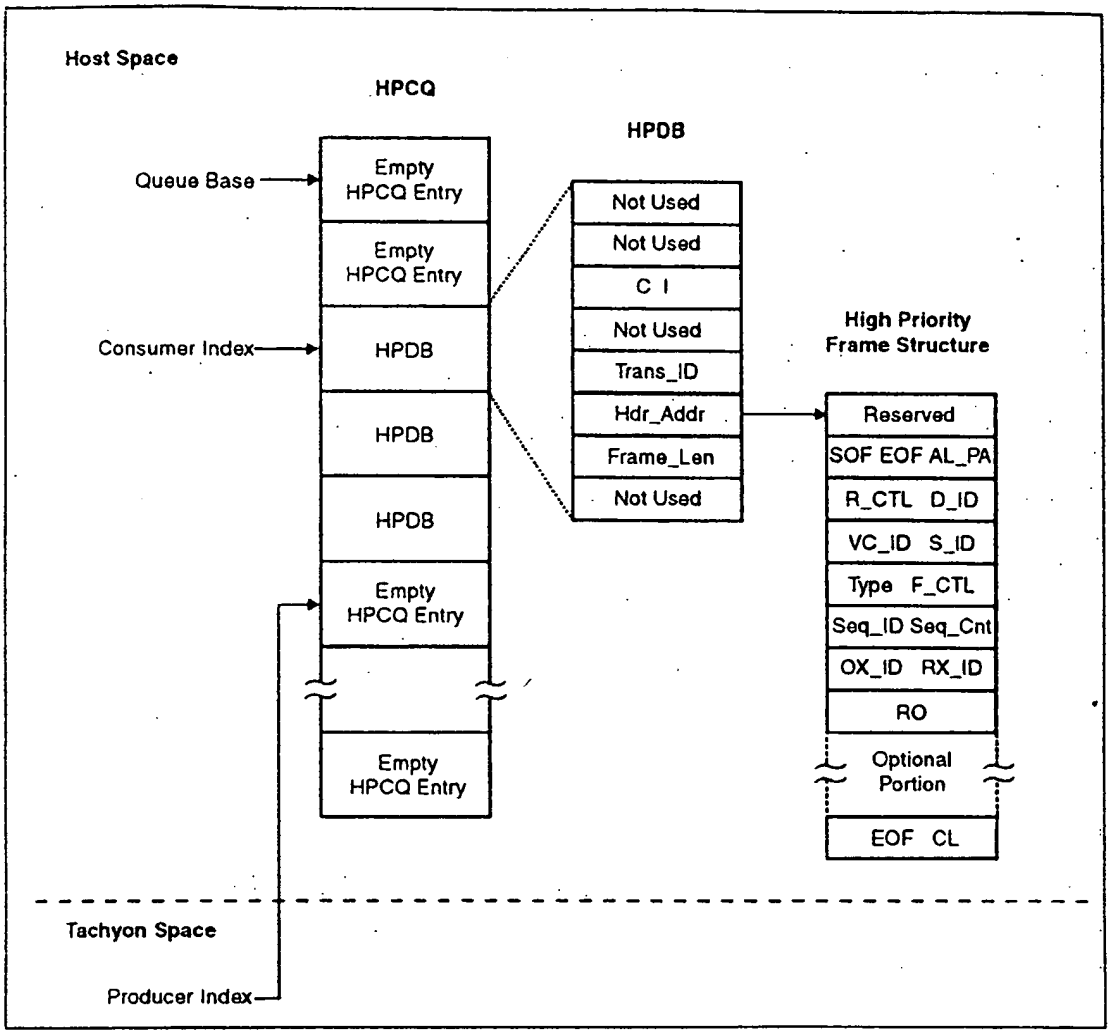

Figure 3.10 Transmit Process via the HPCQ

The host uses the HPCQ to send single frames or single frame sequences. The HPCQ allows the special Fibre Channel error recovery frames to be sent when the Outbound Message Channel is blocked due to an error. If an OCQ sequence is in progress, then Tachyon interweaves frames from the HPCQ with frames from the OCQ. When Tachyon finishes sending the OCQ sequence, then Tachyon sends all remaining HPCQ frames before sending any other OCQ sequence. The host can use the HPCQ to send link control frames, ACKs, and RJTs. The host can also use the HPCQ to re-transmit frames that were busied by the remote node. The host should send Class 1 frames via the HPCQ on an exception basis only, because Tachyon does not check for an open Class 1 connection for a high priority frame.

The host must never send a frame with an EOFdt to close a connection via the HPCQ because Tachyon does not interpret the high priority frame to affect connection status.

PTI 172483

For Class 3, the EOF field in the High Priority Frame Structure must be set to EOFt.

The HPCQ functions much like the OCQ, with a few differences due to the special nature of the high priority frames. The first difference is in the host data structure that describes the data to be sent. As the Transmit Process via the HPCQ figure shows, the entire frame must reside in a single buffer as defined by the High Priority Frame Structure as referenced by the Hdr_addr in the HPDB. Chained HPDBs cannot be sent. Since the High Priority Channel only allows a single frame to be sent, it is the host's responsibility to ensure that the length of each frame is less than the maximum allowed by the remote node or fabric.

Another difference between the HPCQ and the OCQ is that the HPDB contains fewer fields than the ODB counterpart. The fields available in the HPDB are:

1. Trans_ID
2. Hdr_Addr
3. Hdr_Len
4. No_comp (26) and No_int (25) bits of the Control (CNTL) field.

Finally, Tachyon does not wait for ACKs from the remote node before generating a completion message to the host. Once Tachyon DMAs the high priority frame to the Outbound Frame FIFO, Tachyon sends a hi_pri_outbound completion message to the host. If the frame was a Class 1 or 2 sequence, and the remote node generates an ACK for it, the ISM passes the ACK up to the host as an SFS and generates an unknown_frame_completion message. An exception to this rule is if the host uses the HPCQ to re-transmit a busied frame for the current sequence being transmitted. In this case, the OSM recognizes the frame as an ACK for the current sequence and reclaims the credit for it. The host is not notified of the ACK reception.

3.6.8 Small Data Payload Transmit Processes

Three methods are available for transmitting single frame sequences with small data payload; two methods via the OCQ and one method via the HPCQ. Small data payload is defined as "the Tot_Len less than or equal to Frame_len". The maximum Frame_len value for Tachyon is 2048 bytes.

An example of a SFS with small data payload is a FCP Command Sequence (FCP_CMND). Refer to "3.8.11 FCP Write Exchange - Initiator Tachyon" on page 75.

Small Data Payload Transmit Process via the OCQ with the EDB

In the first method, the address in the EDB points to the small data payload. Tachyon may transmit small payloads using this method, but this method is recommended for large payloads (payloads greater than Frame_len value). Refer to "3.6.2 Outbound Command Queue Transmit Details" on page 40.

1. Place the Tachyon Header Structure in buffer "A".
2. Place the data payload in buffer "B".
3. Build an EDB which points to buffer "B".
4. ODB Hdr_Addr points to buffer "A". [Hdr_Len = 32 bytes]
5. ODB EDB_Addr points to EDB.
6. ODB Hdr_len = length of the Tachyon Header Structure.
7. ODB Tot_Len = length of the Tachyon Header Structure's optional portion + the data payload.
8. ODB E bit is cleared to zero.

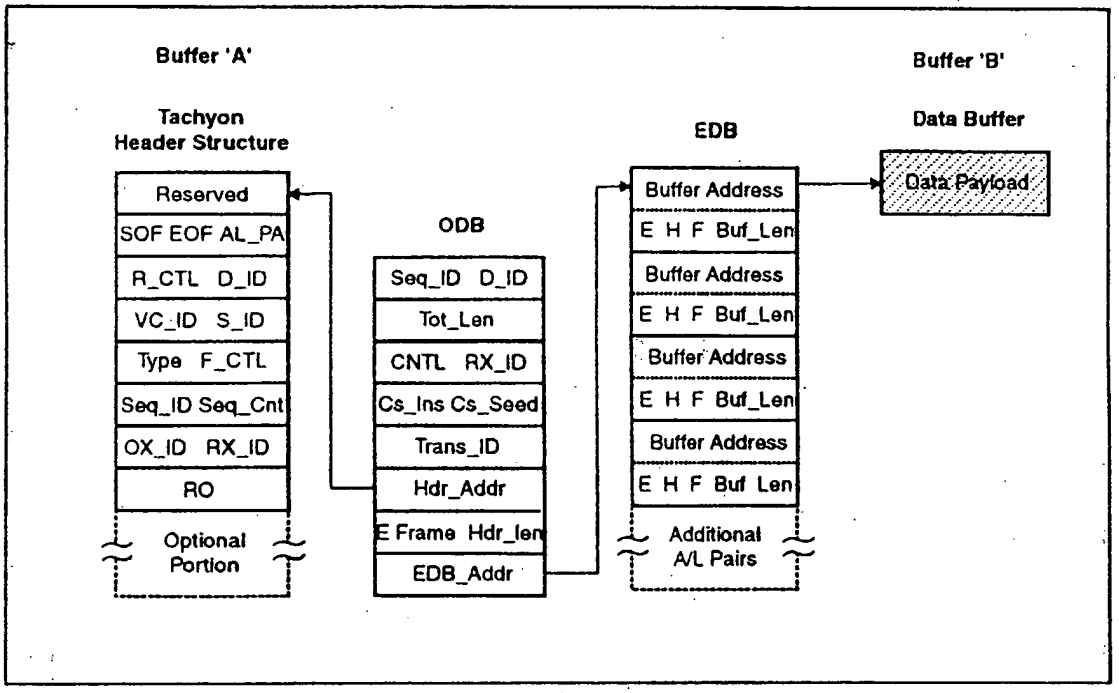

Figure 3.11 Small Data Payload Transmit via the OCQ with the EDB

PTI 172485

Small Data Payload Transmit Process via the Tachyon Header Structure

In the second method, the small data payload is stored in the Optional Portion of the Tachyon Header Stucture. The Optional Portion of the Tachyon Header Structure is available for optional FC Headers (OHDR), as well as for an optional small data payload. This method is recommended for Class 1, Class 2, Class 3 small payloads.

1. Place the Tachyon Header Structure with the small data payload into buffer "C".
2. The ODB Hdr_Addr points to buffer "C".
3. The ODB Hdr_Len is set to the Tachyon Header Structure (which is 32 bytes) + the optional portion (this includes any optional headers and the data payload). The Hdr_len must not exceed Frame_Len + 32 bytes. The data payload must be less than Frame_Len - 32 bytes.
4. The ODB Tot_Len = Length of optional headers (OHDRs) and the data payload.
5. The ODB E bit is set to one.

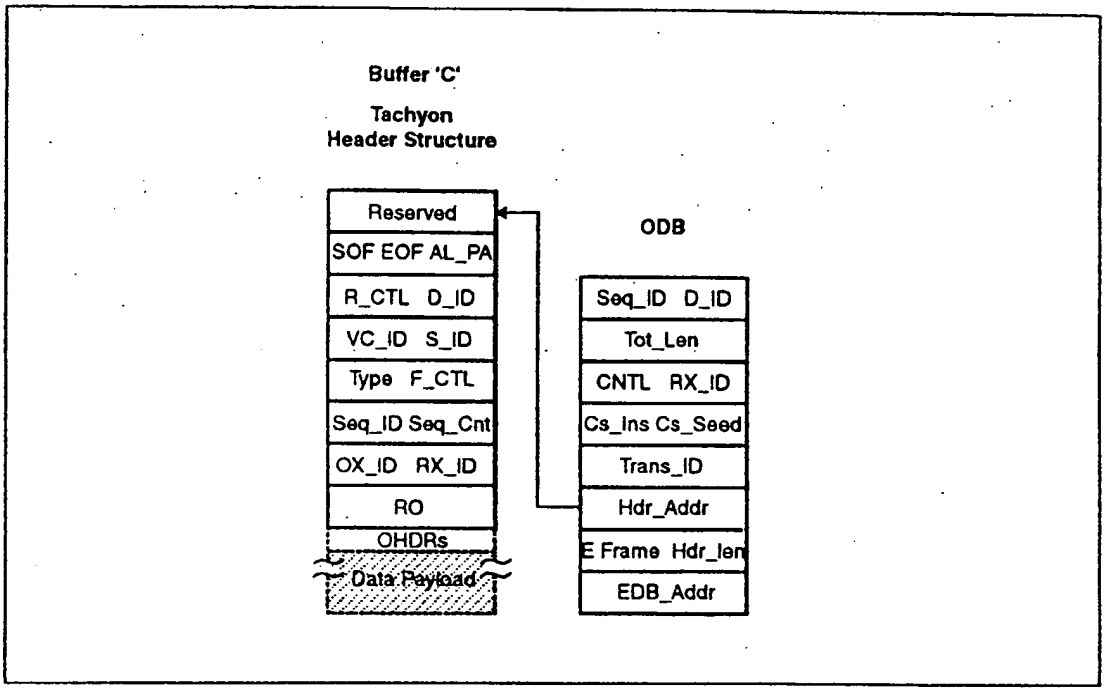

Figure 3.12 Small Data Payload Transmit via the Tachyon Header Structure

PTI 172486

Small Data Payload Transmit Process via the HPCQ

In the third method, the small data payload is stored in the Optional Portion of the High Priority Frame Structure. The Optional Portion of the High Priority Frame Structure is available for optional FC Headers (OHDR), as well as for an optional small data payload. This method is only recommended for Class 3. Class 1 and Class 2 sequences require that the host manually process the ACK frames. This method is not to be used for FC_AL. Refer to "3.6.7 High Priority Command Queue Transmit Details" on page 45.

1. Place the High Priority Frame Structure with the small data payload into buffer "D".
2. The HPDB Hdr_Addr points to buffer "D". The Hdr_Len in the HPDB = 32 bytes + the DATA portion + 4 bytes. The Hdr_len must not exceed Frame_Len + 32 + 4 bytes.
3. The Optional Portion = OHDRs + the data payload.

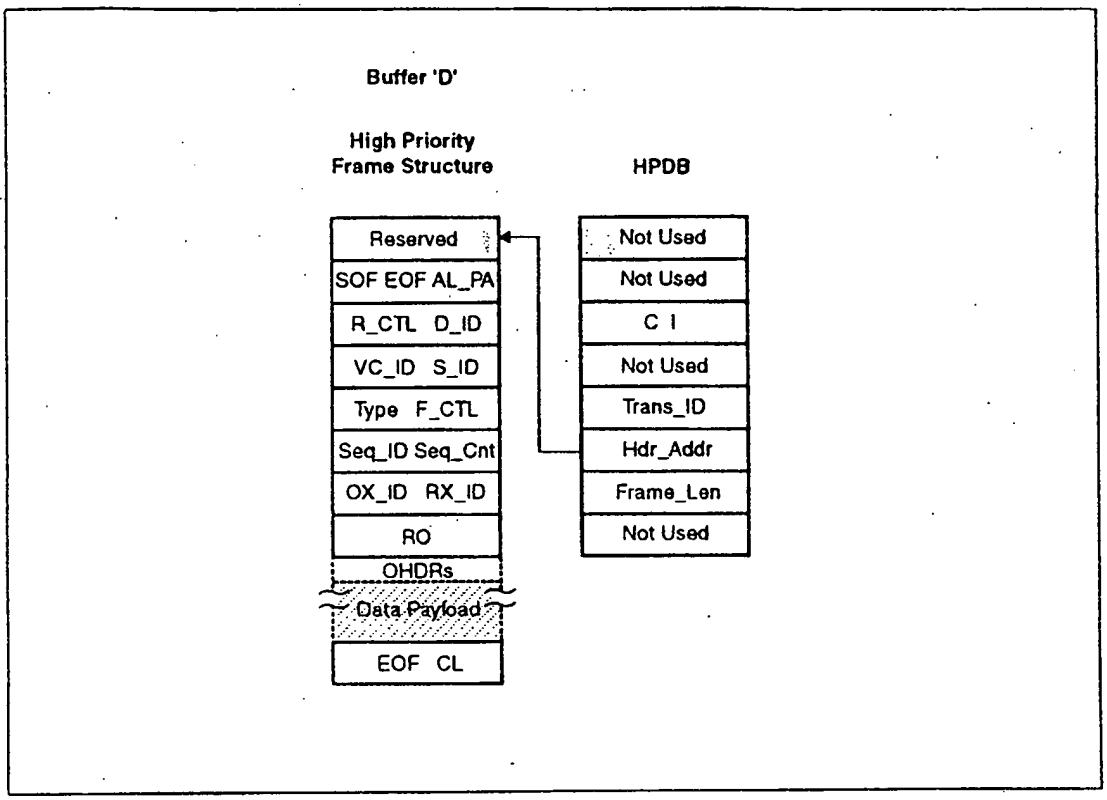

Figure 3.13 Small Data Payload Transmit via the HPCQ

PTI 172487

3.6.9 Recommended Transmit Methods

Depending on the size of data payload, the class of service, and the management of ACKs, the following transmit methods are recommended.

Recommended Transmit Method	Large Payloads	Class 1, Class 2, Class 3, and FC-AL Small Payloads	Class 3 Small Payloads for N_Port
Via the OCQ with the EDB pointing to the data payload	✓		
Via the OCQ with the data payload stored in the Optional Portion of the Tachyon Header Structure		✓	
Via the HPCQ with the data payload stored in the Optional Portion of the High Priority Frame Structure			✓

Table 3.9 Recommended Transmit Methods

PTI 172488

3.7 Non-SCSI Receive Details

3.7.1 Descriptions for the Receive Process Blocks

Refer to "Figure 3.1 Tachyon Internal Block Diagram" on page 17. The Tachyon internal and host-based blocks for the receive process are described briefly below.

Block	Location	Description
Frame Manager	Tachyon	The Frame Manager is Tachyon's front end. It is responsible for the FC-1 functions of transmitting and receiving Fibre Channel frames. It is capable of generating interrupts to the host when certain link configuration changes occur to which the host must respond. The interrupt process occurs as part of Loop Initialization and any time the link has been broken. The Frame Manager responds to a reset condition by initializing all of its registers to their default values.
10B/20B De-MUX	Tachyon	The 10B/20B De-multiplexer is responsible for receiving incoming encoded data, either 10-bits wide or 20-bits wide, from the PLM and packing it into 20 bits for decoding. A 20-bit data width is used with a 100 Mbyte/sec link speed. All other link speeds transmit 10-bit wide data. The data width is specified in the Parallel ID field of the PLM interface.
20B/16B Decoder	Tachyon	The 20B/16B Decoder is responsible for converting 20-bit wide data received from the 10B/20B Demultiplexer from the transmission codes used on the link into 8-bit data bytes.
Elastic Store/ Smoothing	Tachyon	The Elastic Store is responsible for retiming received data from the link clock (RX_CLK) to the internal clock (SCLK). The Elastic Store holds 16-bit wide data.
OS Processor CRC Checker	Tachyon	The Ordered Set Processor is responsible for detecting incoming frame boundaries, verifying the Cyclic Redundancy Check (CRC), and passing the data to the inbound FIFO.
Inbound Data FIFO	Tachyon	The Inbound Data FIFO buffers frames while Tachyon verifies their CRC. It also serves as high-availability, temporary storage to facilitate the Fibre Channel flow control mechanisms. This FIFO is sized to hold a maximum of four 2-Kbyte frames (including headers), though different classes of service may affect the BB_Credit that is used. Refer to "Login Parameters" on page 109.

Table 3.10 Receive Process Blocks

Block	Location	Description
ISM	Tachyon	The Inbound Sequence Manager is responsible for receiving all inbound frames, reassembling multiframe sequences, and establishing sequence completions. Data frames are received by the ISM in single frame sequences, in order multiframe sequences, or Out of Order multiframe sequences. For Class 1 or Class 2, for each data frame received, the ISM generates and sends an ACK or BSY response to theACK FIFO. For each link control frame received, the ISM notifies the OSM of the type of response.
Inbound Data Manager	Tachyon	The Inbound Data Manager maintains the host buffer structures for sequence reassembly of incoming data frames.
Inbound Block Mover	Tachyon	The Inbound Block Mover is responsible for DMAing inbound data into buffers specified by the MFSBQ, SFSBQ, or the SCSI Buffer Manager.
Inbound Message Channel	Tachyon	The Inbound Message Channel is responsible for maintaining the IMQ. This includes supplying the Inbound Data Manager with the address of the next available entry in the IMQ and generating an imq_buf_warn completion message when the number of available entries in the IMQ is down to two.
SFS Buffer Channel	Tachyon	The Inbound Buffer Channel SFS is responsible for managing the SFSBQ. It supplies addresses of empty SFS buffers to the Inbound Data Manager and generates a low buffer warning when the supply of SFS buffers runs low.
MFS Buffer Channel	Tachyon	The Inbound Buffer Channel MFS is responsible for managing the MFSBQ. It supplies addresses of empty MFS buffers to the Inbound Data Manager and generates a low buffer warning when the supply of MFS buffers runs low.
IMQ	Host-based	The Inbound Message Queue notifies the host of inbound and outbound transaction information and status information.

Table 3.10 Receive Process Blocks (Continued)

PTI 172490

Block	Location	Description
SFSBQ	Host-based	The host uses the Single Frame Sequence Buffer Queue (SFSBQ) to inform Tachyon of the location of the buffers that are used to receive single frame sequences. The size of these buffers is defined by the SFS Buffer Length register. Each SFSBQ entry contains eight addresses to buffers packed into a 32-byte structure. The SFSBQ must be configured to pass error frames to the host.
MFSBQ	Host-based	The host uses the Multiframe Sequence Buffer Queue (MFSBQ) to inform Tachyon of the location of buffers that are used to receive and reassemble multiframe sequences. The size of these buffers is defined by the MFS Buffer Length register. Each MFSBQ entry contains eight addresses to buffers packed into a 32-byte structure. If the host never uses the MFSBQ, memory may be saved by setting the MFSBQ length to two entries and ensuring that these entries point to a safe place in memory in case a WRITE occurs.

Table 3.10 Receive Process Blocks (Continued)

The Tachyon architecture differentiates between 1) Single Frame Sequence Reception, 2) Multiframe Sequence, In Order Reception, and 3) Multiframe Sequence, Out of Order Reception. This differentiation is due to Tachyon resources. Tachyon does not need to use resources to receive a single frame sequence. However, for multiframe sequences, Tachyon does need to use resources to manage the reassembly of frames. Due to this resource demand, Tachyon reassembles only one incoming multiframe networking sequence at a time.

Tachyon does not support reception of a Class 2 or Class 3 multiframe sequence if an inbound Class 1 connection is open. However, Tachyon does support the reception of single frame sequences while reassembling a multiframe sequence. This process allows a host to receive short sequences while Tachyon is reassembling a longer incoming sequence.

3.7.2 Single Frame Sequence Reception

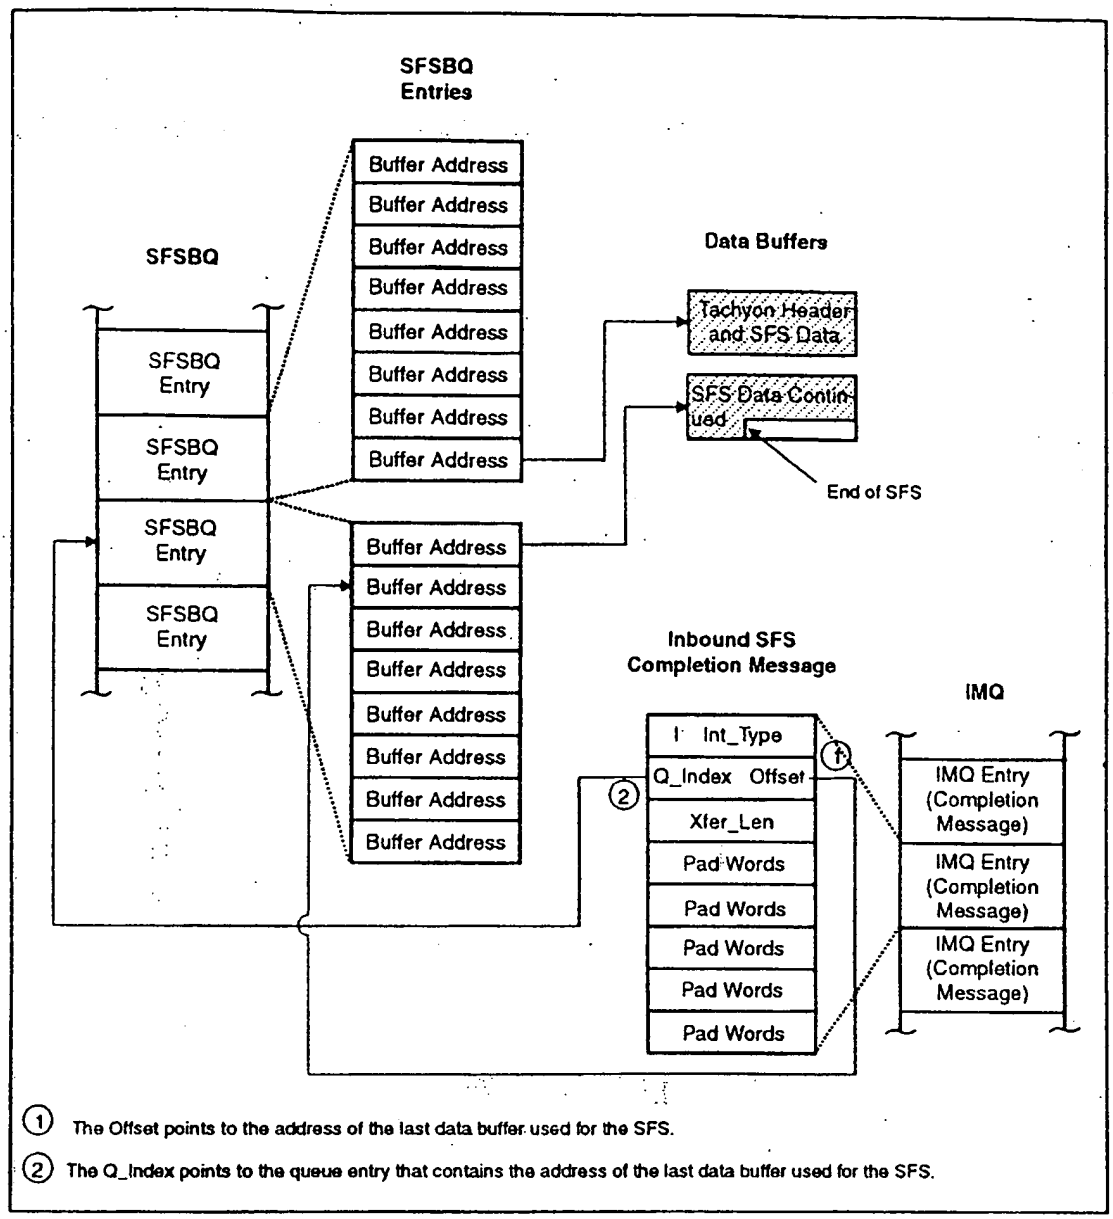

Figure 3.14 SFS Reception

PTI 172493

The host uses the SFSBQ to inform Tachyon of the location of data buffers in host memory. Tachyon uses these data buffers to receive the incoming sequences contained within a single frame. Each 32-byte SFSBQ entry contains eight addresses. The size of these buffers is programmable and is determined by the SFS Buffer Length register.

All buffers must be the same size, and must be a power of 2 bytes in length. The buffers must be aligned on sizeof(buffer) boundaries. The host is the producer and Tachyon is the consumer of the SFSBQ.

As Tachyon receives a single frame sequence, it places the entire SFS (the Tachyon Header Structure followed by the SFS data) in the buffer defined by the address from the SFSBQ. If the SFS is larger than one buffer size, then the remaining data is packed into the next buffers, as required, until all of the SFS is stored.

Next, Tachyon posts an inbound_sfs completion message to the IMQ, generates an interrupt to the host, and increments the SFSBQ Consumer Index register.

Since Tachyon only provides information on the last buffer used (via the Q_Index and the Offset in the inbound_sfs completion message) and the data transfer length, it is the host's responsibility to know the location of the first buffer of the SFS by tracking where Tachyon left off from the previous SFS. All the data buffers used from the first buffer to the last buffer contain the current transfer. The host can either keep track of the index and buffer offset to the first address to a empty buffer, or calculate it by using the transfer length. It is strongly suggested that the host keep track of the first address to an empty buffer, so that it does not have to calculate the buffer in which the first frame starts.

3.7.3 Multiframe Sequence, In Order Reception

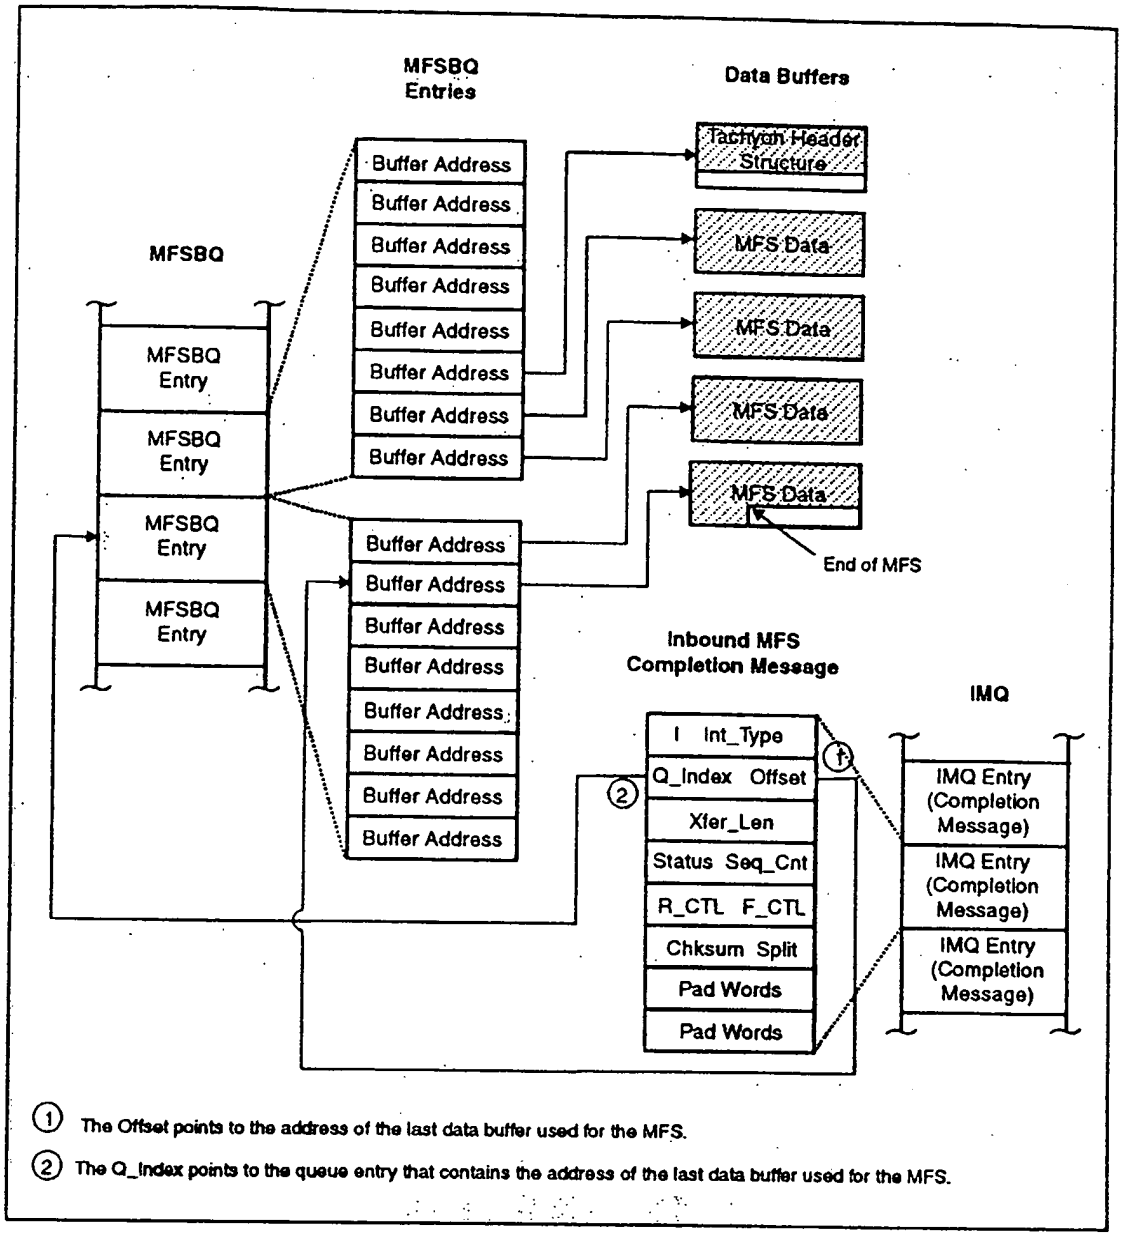

Figure 3.15 MFS In Order Reception

PTI 172495

The host uses the MFSBQ to inform Tachyon of the location of data buffers in host memory. Tachyon uses these data buffers to receive and reassemble incoming sequences that have been split into an arbitrarily large number of frames. Each 32-byte MFSBQ entry contains eight pointers to empty data buffers. The size of these buffers is programmable and determined by the MFS Buffer Length register. All buffers must be the same size, and must be a power of 2 bytes in length. The buffers must be aligned on sizeof(buffer) boundaries. The host is the producer and Tachyon is the consumer of the MFSBQ.

The host should select a buffer length to optimize the memory page usage in the host. This may include making the buffer length equal to the size of a host memory page to allow page remapping after a header/ data split.

When the first frame of a new sequence arrives, Tachyon copies the Tachyon Header Structure into the beginning of the next available MFS buffer. Unlike the SFS receive process, Tachyon packs the data payload of the frame into the *next* MFS buffer following the buffer with the Tachyon Header Structure. As each new frame arrives, Tachyon discards the FC Header and packs the data into the buffers. As the buffer becomes full, the next buffer is used until the entire MFS is stored. Tachyon uses the buffers in the order that they were placed in the MFSBQ. This allows the host to know which buffers have been used. Once all the frames arrive and the sequence is reassembled in memory, Tachyon posts an inbound_mfs_completion message, generates an interrupt to the host to process the entire sequence, and increments the MFSBQ Consumer Index register.

Since Tachyon only provides information on the last buffer used and the data transfer length, it is the host's responsibility to know the location of the first buffer of the transfer by tracking where Tachyon left off from the last transfer. All the data buffers used from the first buffer to the last buffer contain the current transfer. The host can either keep track of the index and buffer offset to the first address to a empty buffer, or calculate it by using the transfer length. It is strongly suggested that the host keep track of the first address to an empty buffer, so that it does not have to calculate the buffer in which the first frame starts.

3.7.4 Multiframe Sequence, Out of Order (OOO) Reception

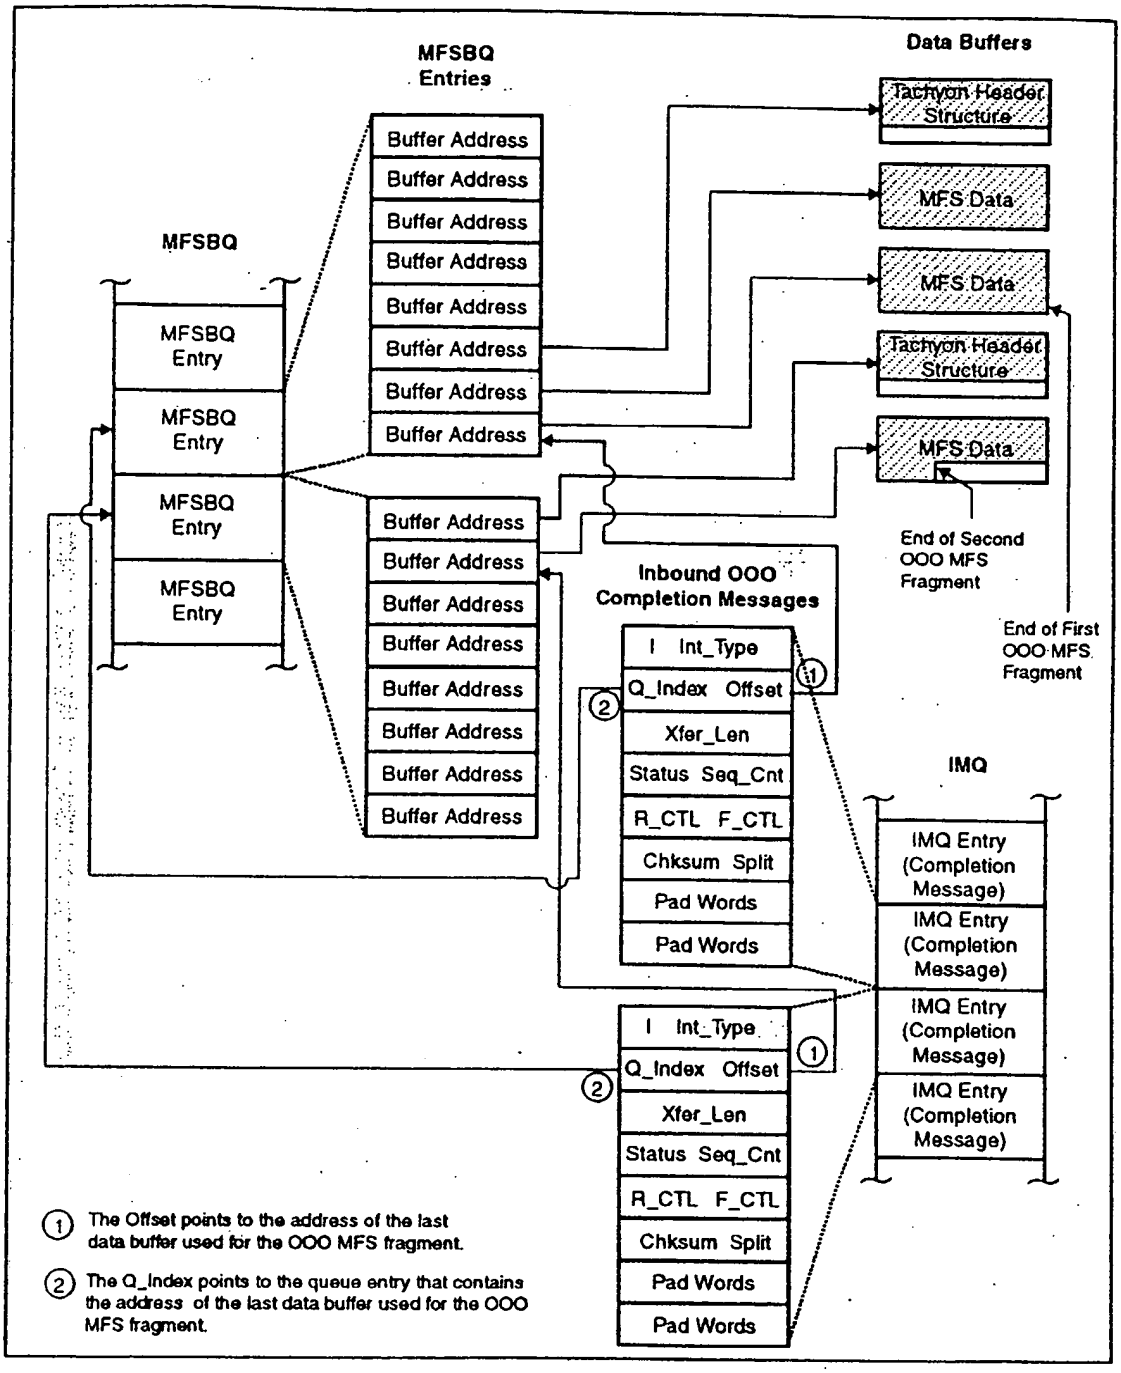

Figure 3.16 MFS Out of Order (OOO) Reception

PTI 172497

Multiframe sequence Out of Order Reception begins the same as In Order Reception. When the first frame of a new sequence arrives, Tachyon copies the Tachyon Header Structure into the beginning of the next available MFS buffer. Tachyon places the data payload of the frame into the following buffer. As each subsequent frame arrives, Tachyon discards the FC Header and packs the data contiguously into host buffers. When the first OOO frame arrives, as defined by its SEQ_CNT field or its Relative Offset (RO) field, Tachyon and the host process this frame differently than the In Order Reception method.

When Tachyon detects an OOO frame, the ISM generates an inbound_ooo_completion message that indicates the in order portion of the MFS and the last MFS buffer used. Tachyon passes the completion message to the IMQ, but does not generate an interrupt until all frames of the sequence are received. Next, the ISM obtains the next available buffer and copies the Tachyon Header Structure of this OOO frame into it. The ISM obtains the next buffer and copies the data payload of this OOO frame into it. At this point, if subsequent frames follow in order from the OOO frame, Tachyon discards the Tachyon Header Structures and packs the data into the host buffers. Tachyon packs this data into each of the buffers on the MFSBQ, obtaining a new buffer when the current buffer is full, until the entire MFS is stored.

If another frame arrives OOO from the previous OOO portion, Tachyon generates a new inbound_ooo_completion message, and the process is repeated. When Tachyon receives the last frame of the sequence, Tachyon passes it to the host and generates an inbound_mfs_completion message. At this time, the ISM generates an interrupt to the host to process the entire sequence. With the information in each of the Tachyon Header Structures that Tachyon passed to the host for each in order portion and the information of the completion messages, the host has enough information to re-order the OOO multiframe sequence.

3.7.5 Multiframe Sequence, Deferred P_BSY Mode

Tachyon reassembles one MFS at a time. When Tachyon receives a new MFS while a MFS is being reassembled, Tachyon busies (P_BSYs) the frame(s) of that new MFS. The remote node that was P_BSY'd can retry the sequence later.

This operation works reasonably well when Tachyon communicates with only one other remote node, e.g. when Tachyon is a client in a client-server configuration. However, a server in a client-server configuration may be receiving MFSs from many clients (remote nodes) at the same time. When Tachyon receives several new MFS frames, other than the current MFS being reassembled, the new frames are P_BSY'd. The remote nodes re-send the frames causing even more interference with the current MFS being reassembled. These excessive retries may cause the remote nodes to stop sending their MFS, which leads to upper layer error recovery overhead.

To manage sequence congestion under light load conditions, Tachyon re-assembles MFSs without help from the host. To manage sequence congestion under heavy load conditions (more than one MFS is being sent to Tachyon at a time), Tachyon uses help from the host and uses the Deferred P_BSY mode. The Deferred P_BSY mode allows Tachyon to be used as a server in a client-server configuration that allows fair access to all remote nodes with reasonable performance.

Deferred P_BSY mode involves the following operation between Tachyon and the host.

1. Tachyon receives the first frame of a new MFS while reassembling an existing MFS. In Deferred P_BSY mode, the host sets the Disable AUTO P_BSY bit in the Tachyon Configuration register so that Tachyon does not automatically send a P_BSY to the remote node. (With Tachyon as a client in a client-server configuration, the Disable AUTO P_BSY bit would be cleared to zero and Tachyon would automatically send a P_BSY.)
2. Tachyon sets the internal Tachyon Deferred P_BSY flag to indicate that it has switched to the Deferred P_BSY mode.

PTI 172498

3. Tachyon sends the new MFS frame to the host via the SFSBQ. Tachyon generates an inbound_busied_frame completion message. (This type of completion message does not generate an interrupt to the host.)
4. The host stores this frame in a Deferred P_BSY queue.
5. When Tachyon receives all the frames for the current MFS it is reassembling, it inspects its internal Deferred P_BSY flag. If Tachyon is not in the Deferred P_BSY mode (see Step 2, above), it generates the final ACK for the MFS, as it normally would.

   If Tachyon is in the Deferred P_BSY mode, Tachyon sets the Deferred ACK bit in the inbound_mfs_completion message, and final ACK generation is deferred to the host. Before the host sends the final ACK to the current remote node, it can send a P_BSY to one of the waiting remote nodes. This allows time for one of the waiting remote nodes to retry its MFS before the current remote node can start a new MFS (this current remote node is not able to stream). In Deferred P_BSY mode, Tachyon attempts to process MFSs in a way that all remote nodes have fair access.

Note	For Deferred P_BSY mode to work properly, the remote node must interlock the first frame of each sequence, i.e., for every sequence, the remote node must only send the first frame of the MFS and wait for an ACK before sending the rest of the sequence. (Refer to the X_ID Interlock bit in the ODB data structure on p. 161.) If the remote node does not support this operation, then the remote node should be limited to a Class 2 EE_Credit of one to force this mode.

6. When the host receives an inbound_mfs_completion message, it needs to inspect its Deferred P_BSY queue. If there is an entry in the queue, the host must create a P_BSY frame for this entry. If the queue is empty, the host should set the Clear Deferred P_BSY bit to one in the Tachyon Control register to return to the normal mode. The host should then send the P_BSY via the HPCQ. The host should only send one P_BSY per inbound_mfs_completion message received.

   The host should time the frames in the Deferred P_BSY Queue and send a P_BSY for each item on the queue before its ED_TOV expires. This forces the remote node to re-transmit the first frame before sequence ED_TOV expires.
7. If Tachyon sets the Deferred ACK bit in the inbound_mfs_completion message, the host needs to create an ACK frame for the last frame of the sequence and send it via the High-Priority Command Queue. It can create this ACK frame from information in the frame header sent to the host for the first MFS frame received and the F_CTL and Seq_Cnt fields in the inbound_mfs_completion message.

The Deferred P_BSY mode is used with Class 2 MFSs only.

3.7.6 Expiration/Security Header Frame Reception

Tachyon does not support frames or sequences that use Expiration/Security Headers. If Tachyon receives a frame that contains an optional Fibre Channel Expiration Header, Tachyon passes the frame to the host as an SFS with an inbound_unknown_frame_completion message, without acknowledging the frame. The host is expected to send a P_RJT frame to terminate the sequence.

3.7.7 Open Broadcast Replicate Reception

Refer to "3.9.7 Open Broadcast Replicate Support" on page 86.

PTI 172499

3.8 SCSI Hardware Assists Details

Tachyon supports SCSI I/O transactions (exchanges) using two methods. The first method uses host-based transaction management. In this method, the host transmits and receives the various FCP sequences using the general transmit and receive processes, as described in the previous sections. By using the host-based transaction management method, Tachyon reassembles only one SCSI unassisted multiframe sequence at a time.

The second method uses Tachyon's SCSI hardware assists. With this method, Tachyon assists the host transaction management through the use of a shared host data structure called the SCSI Exchange State Table (SEST). By using Tachyon's SCSI hardware assists, the host can concurrently reassemble up to 16,384 SCSI assisted sequences.

Tachyon maintains an on-chip cache for up to 16 concurrent inbound transactions. Tachyon uses a Least Recently Used caching algorithm to allow the most active exchanges to complete their transfers with the minimum latency.

3.8.1 Descriptions for the SCSI Hardware Assists Blocks

Refer to "Figure 3.1 Tachyon Internal Block Diagram" on page 17. The Tachyon internal and memory-based blocks for the SCSI hardware assisted transaction process are described briefly below.

Block	Location	Description
SCSI Buffer Manager	Tachyon	The SCSI Buffer Manager is responsible for supplying the Inbound Data Manager with addresses of buffers to be used for inbound SCSI data frames.
SCSI Exchange Manager	Tachyon	In conjunction with the SEST, the SCSI Exchange Manager provides Tachyon with the hardware assists for SCSI I/O transactions. It converts SEST entries to ODB format for the OSM's use.
Register Block	Tachyon	The Register Block contains various status and configuration registers. These registers are used for initialization, fatal error processing, and maintenance of the queues used to transfer data between Tachyon and the host.
SCSI Read/Write Channel	Tachyon	The SCSI Read/Write Channel manages requests from the SCSI Exchange Manager and interfaces to the Tachyon System Interface (TSI) arbiter blocks.
SEST	Host-based	The SCSI Exchange State Table (SEST) contains SCSI Exchange information. Refer to "3.8.2 SCSI Exchange State Table (SEST)" on page 63.

Table 3.11 SCSI Hardware Assists Blocks

3.8.2 SCSI Exchange State Table (SEST)

The SEST is a memory-based data structure that has shared access between Tachyon and the host. The SEST is an array of 32-byte entries. Each used SEST entry corresponds to a current SCSI exchange, or I/O operation.

Each SEST entry contains the following information:

1. Information supplied by the host driver for Tachyon to manage the exchange.
2. Information stored by Tachyon to track the current state of the exchange.

For initiators in SCSI Write transactions, the Outbound SEST Entries contain information indicating where outbound data resides in memory and what parameters to use in the transmission of that data on the Fibre Channel link. For initiators in SCSI Read transactions, the Inbound SEST Entries contain information indicating where inbound data is to be placed in memory.

SEST entries are indexed by an Exchange Identifier (X_ID), either the Originator Exchange Identifier (OX_ID) or the Responder Exchange Identifier (RX_ID). In an initiator application, the OX_ID indexes the SEST. In a target application, the RX_ID indexes the SEST.

The starting address of the SEST is programmable and is defined by the SEST Base register. The length of the SEST is programmable and defined by the SEST Length register. The minimum SEST length is one entry. The maximum SEST length is 16,384 entries. The SEST must be aligned on a sizeof (table) boundary which must be a power of 2 bytes. When calculating the SEST alignment, if the length is not a power of 2, then the length value should be rounded up to the next highest power of 2 value.

As an example, if the SEST has 14 entries, the length is rounded up to 16 entries, and the sizeof (table) equals 16 x 32 or 512 bytes. Therefore, this SEST is aligned on a 512-byte boundary. The starting address for this SEST example must have the nine least significant address bits (bits 8..0) cleared to zero.

When the host determines that the destination is a SCSI target or a SCSI initiator, then the host may, during N_Port login, advertise the number of concurrent sequences equal to its ability to support SCSI sequences, that is, the number of SEST entries should be identical to the number of maximum concurrent sequences, on the condition that only SCSI sequences are allowed. Otherwise, the host should advertise the number of concurrent sequences to be one. The number of concurrent sequences is specified in N_Port Common Service Parameters, refer to "FC-PH, Revision 4.3, Section 23.6.3 and 23.6.8.6".

If SCSI assists are never used, then the SEST does not have to be configured. If the SCSI Enable (se) bit in the Tachyon Configuration register is not set to one, Tachyon never attempts to access the SEST. In this case, host-based transaction management and the general transmit and receive processes are used.

3.8.3 Invalidating and Re-Using Outbound SEST Entries

SCSI Transaction Completes Normally

1. The host receives an inbound_scsi_status completion message.
2. The host can immediately re-use the Outbound SEST Entry. If the host does not immediately re-use this Outbound SEST Entry, then the host should invalidate the entry by clearing the STE Valid (V) bit to zero.

Host Terminates SCSI Transaction When It Is Still Active (e.g. Abort)

To re-use an Outbound SEST Entry that is currently active, e.g., to abort a request, the host must invalidate the entry, terminate the exchange (if necessary), and also verify that the entry is unattached to a linked list of SCSI outbound transactions.

1. To invalidate the Outbound SEST Entry, the host performs the following steps.
   a. Set the SCSI Freeze (sf) bit in the Tachyon Control register (bit 3) to one.
   b. Poll the SCSI Freeze Status (ss) bit in the Tachyon Status register (bit 6) until Tachyon sets it to one. Continue servicing new IMQ entries while polling.
   c. Clear the STE Valid bit (V) in the Outbound SEST Entry to zero.
   d. Clear the SCSI Freeze bit in the Tachyon Control register to zero.
   e. Repeat a. through d. once again, then proceed to f.
   f. At this point, the Outbound SEST Entry is invalidated, however, it cannot be used until it is unattached from a linked list of SCSI outbound transactions. Refer to the Data Phase description of the "FCP Write Exchange - Initiator Tachyon" on p. 75.

WARNING	Failure to invalidate Outbound SEST Entries as described in Step 1. a. through Step 1. f. may cause data in the SEST to be corrupted.

Invalidating an Outbound SEST Entry does not stop Tachyon from sending the data associated with this SEST entry if Tachyon started sending the data before the SEST entry was invalidated.

2. To verify that the Outbound SEST Entry is not attached to a linked list, the host performs the following steps.
   a. Determine if the entry has ever been attached to a linked list.
      1) At initialization of the Outbound SEST Entry, clear the Completion Message Disable bit to zero.
      2) At initialization of the Outbound SEST Entry, write to the Total Sequence Length field with a value that is known to be different than the value that Tachyon will overwrite to this field (i.e. the value of the length returned in the FCP_XFER_RDY frame from the target device).
      3) After the Outbound SEST Entry is invalidated, read its Total Sequence Length field. If the Total Sequence Length field contains the value that was written at initialization in substep 2 (above), then the Outbound SEST Entry was never attached to a linked list and it can be re-used. If the Total Sequence Length field contains a modified value, i.e., if it contains the the value of the length returned in the FCP_XFER_RDY frame from the target device, then this Outbound SEST Entry had been and may still be atttached to a linked list.

PTI 172502

b. Determine if the entry is still attached to a linked list.

1) If the host has seen an outbound completion message for this entry, then the entry is not attached to the linked list anymore. At this point, the entry can be re-used.
2) If the host has not seen an outbound_completion message for this entry, then it must wait for Tachyon to process the entry in the linked list. After Tachyon processes the entry, it posts an outbound completion message with the Host Programming Error (HE) bit set to one. When the host receives this completion message, then the entry can be re-used.

WARNING	Re-using an entry that is still attached to a linked list may cause data in the SEST to be corrupted.

3.8.4 Invalidating and Re-Using Inbound SEST Entries

SCSI Transaction Completes Normally

1. When the SCSI transaction completes normally, Tachyon invalidates and flushes the Inbound SEST Entry.
2. Tachyon sends an inbound_scsi_status completion message. When the host receives this completion message, then the entry can be re-used.

Host Terminates SCSI Transaction When It Is Still Active (e.g. Abort)

To re-use an Inbound SEST Entry that is currently active, e.g., to abort a request, the host must invalidate the entry by performing the following steps.

1. Clear the STE Valid bit (V) in the Inbound SEST Entry to zero.
2. In the Tachyon Flush SEST Cache Entry register, write the SEST index of the Inbound SEST Entry that is to be flushed and set the Update (up) bit to one.
3. Poll the Tachyon Flush SEST Cache Entry register to verify that Tachyon has cleared the Update bit to zero. Continue servicing new IMQ entries while polling.
4. Verify that the STE Valid bit is cleared. Repeat Step 1. through Step 4. until the Update bit is cleared in this manner.

WARNING	Failure to invalidate Inbound SEST Entries as described in Step 1. through Step 4. may cause data corruption.

5. Perform ABTS protocol. When this is complete, the entry can be re-used.

PTI 172503

3.8.5 SCSI Exception Processing in Tachyon

Tachyon manages the entire Fibre Channel Protocol for SCSI, with the following exceptions.

1. If any error occurs during an outbound sequence, Tachyon stops processing the SCSI exchange, sends the host an outbound_completion message with the appropriate error bit set to one, and waits for the host to recover from the error.
2. If an error occurs during an inbound sequence, Tachyon continues processing the SCSI exchange and sends the host a bad_scsi_frame completion message. The possible error conditions include receiving an ABTS frame, receiving a frame with an invalid OX_ID, or receiving a frame with the OX_ID out of range (i.e., an OX_ID not in the range specified by the SEST length).

If the target cannot handle the entire requested Write transaction, it sends multiple FCP_XFER_RDYs to accomplish the SCSI exchange. In this case, the initiator Tachyon processes the FCP_XFER_RDY with a DATA_RO equal to zero, but passes all other FCP_XFER_RDYs (in which DATA_RO is non-zero) to the host for processing.

3.8.6 FCP Read Exchange Example

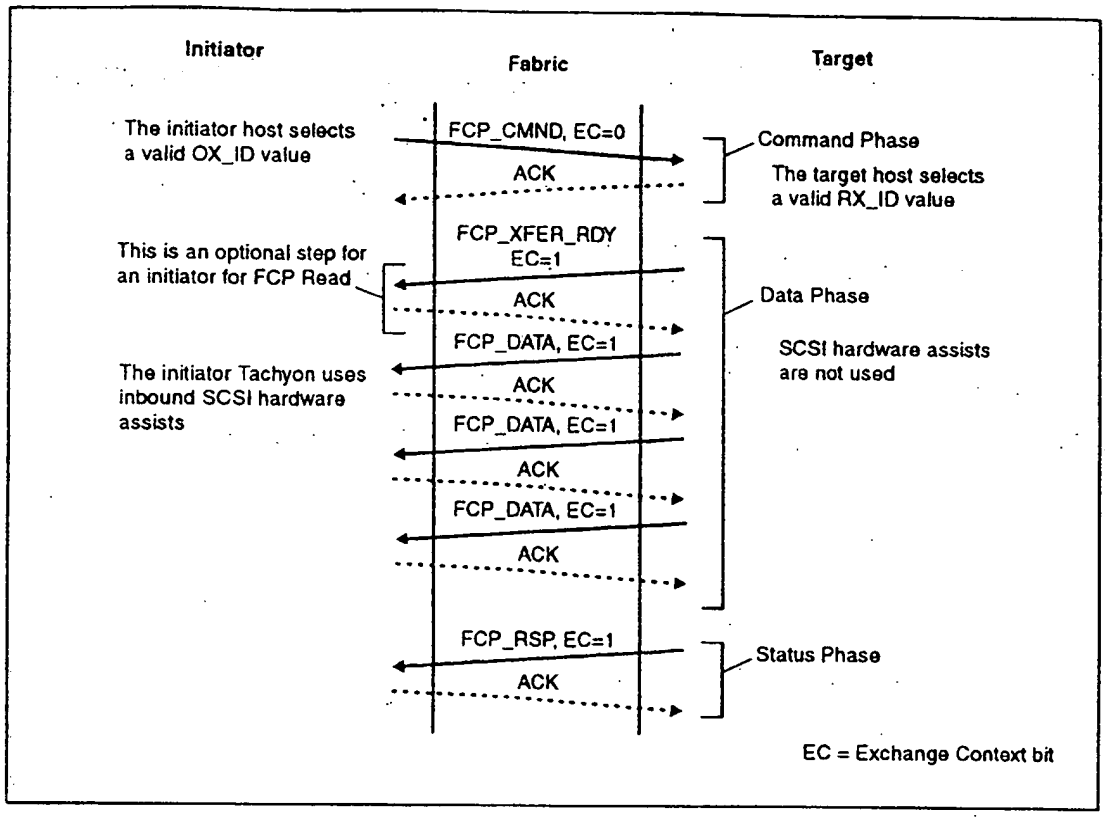

Figure 3.17 FCP Read Exchange Example

3.8.7 FCP Read Exchange, OOO Reassembly - Initiator Tachyon

OOO Reassembly mode enables Tachyon to receive and reassemble OOO sequences for inbound exchanges.

To use OOO Reassembly mode:

- The OOO Reassembly Disable (od) bit of the Tachyon Configuration register must be cleared to zero at initialization.
- SCSI Buffer Length must be a power of 2 bytes.

To receive inbound SCSI data, the initiator host selects a valid OX_ID value which points to an unused location in the SEST. Refer to " Choosing X_ID Values" on page 160. The OX_ID value identifies this particular exchange. Using the OX_ID value, the initiator host builds an Inbound SEST Entry. The Inbound SEST Entry includes the STE Valid bit, the Buffer Offset, and the address of the SDB. The SDB defines host buffers where the Read data is going to be received. For OOO Reassembly, the SDB contains up to 64 pointers to fixed length buffers and the SDB must be aligned on a 256-byte boundary.

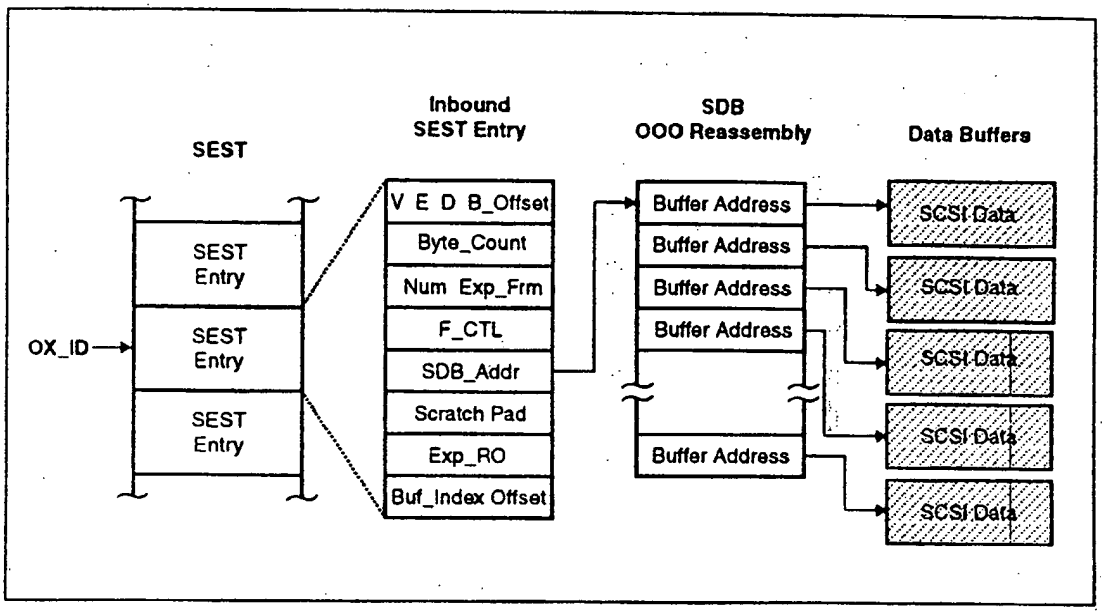

Figure 3.18 FCP Read Exchange - Initiator Host Data Structures

The initiator host clears the Exchange Context bit in the F_CTL field of the Fibre Channel header to zero to indicate that it is the originator of the exchange.

Command Phase

Once the host creates the Inbound SEST Entry, the host creates an FCP_CMND for an FCP Read exchange. The initiator Tachyon sends the FCP_CMND with a valid OX_ID to the target. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

PTI 172506

Data Phase

The initiator Tachyon may receive an FCP_XFER_RDY from the target. This is an optional step for an initiator for an FCP Read because the data frames contain all the information needed to process them.

When Tachyon receives the optional FCP_XFER_RDY from the target, it ACKs the frame (if appropriate) and discards the FCP_XFER_RDY. As each data frame is received, the SCSI Exchange Manager uses the OX_ID to access the appropriate Inbound SEST Entry for the address of the SDB. The SDB and the RO of the data frame determine where data is to be placed in host memory. Tachyon does not examine the TYPE field in the FC Header.

Tachyon maintains an internal cache of 16 Inbound SEST Entries. If the SEST information associated with the received frame is not in cache, then Tachyon writes the "least recently used" cache entry back to the host SEST. Tachyon then fetches into cache the Inbound SEST Entry associated with the received frame. Then it DMAs the Read data to host memory.

The initiator Tachyon automatically handles both single and multiple data phases for inbound data transfers.

Status Phase

When the data phase is complete, the initiator Tachyon receives an FCP_RSP from the target. The FCP_RSP is a Fibre Channel information unit that contains the status information that the SCSI exchange has completed. The initiator Tachyon passes the FCP_RSP to the host via the SFS channel. The initiator Tachyon sends an inbound_scsi_status_completion message to the inititator host. This informs the initiator host that the exchange is completed. The initiator host clears the STE Valid bit to zero which indicates that the Inbound SEST Entry is no longer valid.

The FCP_RSP received by Tachyon as an initiator must be a single frame sequence. If the FCP_RSP is a multiframe sequence (e.g., a large FCP_RSP is sent through a Fabric that only supports small frame sizes), the frames are processed as bad SCSI frames. It is a violation of the FCSI SCSI Profile to send the FCP_RSP as a multiframe sequence.

The Byte Count in the Inbound SEST Entry is updated with the number of bytes received by Tachyon. The initiator host should check this field to detect any lost data.

3.8.8 FCP Read Exchange, In Order Reassembly - Initiator Tachyon

For In Order Reassembly mode for the initiator Tachyon, sequences within the data phase of the same exchange must arrive in order.

To use In Order Reassembly mode:

- The OOO Reassembly Disable bit of the Tachyon Configuration register must be set to one at initialization.
- Tachyon must be operating in an in-order-topology, i.e., point-to-point, a private Arbitrated Loop, a fabric which supports guaranteed in order delivery (negotiated at FLOGIn), or a public Arbitrated Loop using a fabric that supports in order delivery.
- The SCSI Buffer Length must be a multiple of 4 bytes. (For OOO Reassembly, the SCSI Buffer Length is a power of 2.)
- The initiator host must clear the Buffer Offset and Buffer Index fields of the Inbound SEST Entry to zero.

For the initiator host to receive inbound SCSI data, it selects a valid OX_ID value that points to an unused location in the SEST. Refer to " Choosing X_ID Values" on page 160. The OX_ID value identifies this particular exchange. Using the OX_ID value, the initiator host builds an Inbound SEST Entry. For In Order Reassembly, the buffer addresses must be on a word-aligned boundary and the SDB must be aligned on a power of 2 boundary equal to or greater than the SDB length. The SDB length is the number of entries multiplied by four bytes.

The initiator host clears the Exchange Context bit in the F_CTL field of the Fibre Channel header to zero since it is the originator of the exchange.

Command Phase

Once the host creates the Inbound SEST Entry, the host creates an FCP_CMND for an FCP Read exchange. The initiator Tachyon sends the FCP_CMND with a valid OX_ID to the target. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

Data Phase

The initiator Tachyon may receive an FCP_XFER_RDY from the target. This is an optional step for an initiator for an FCP Read because the data frames contain all the information needed to process them.

When Tachyon receives the optional FCP_XFER_RDY from the target, it ACKs the sequence (if appropriate) and discards the FCP_XFER_RDY.

As each data frame is received, the initiator Tachyon determines if the frame is in order for the exchange. If it is not in order, Tachyon passes the frame as a bad SCSI frame to the host via the SFS Buffer Queue, invalidates the SCSI table entry by clearing the V bit in the Inbound SEST Entry to zero, and sends a bad_scsi_frame completion message to the host indicating that an OOO SCSI frame was received erroneously. Refer to "3.3.4 SCSI Completion Messages" on page 31. If the host has set the Bad SCSI Auto ACK bit in the Tachyon Configuration register to one, Tachyon generates an ACK (for Class 1 and Class 2) for the bad SCSI frames and the host is responsible for the error recovery. If the host has cleared the Bad SCSI Auto ACK bit to zero, the host is responsible for the ACK as well as for the error recovery. Tachyon manages all subsequent data frames for that exchange as bad SCSI frames and passes these frames to the host with a bad_scsi_frame completion message.

If the data frame is in order for the exchange, initiator Tachyon checks the OX_ID of the data frame to locate the appropriate Inbound SEST Entry. Tachyon does not examine the TYPE field in the FC Header.

PTI 172508

Tachyon maintains an internal cache of 16 Inbound SEST Entries. If the SEST information associated with the received frame is not in cache, then Tachyon writes the "least recently used" cache entry back to the host SEST. Tachyon fetches into cache the Inbound SEST Entry associated with the received frame. Tachyon then DMAs the Read data to host memory.

An initiator Tachyon automatically handles both single and multiple data phases for inbound I/Os.

Status Phase

When the final data phase is complete, the initiator Tachyon normally receives an FCP _RSP from the target. The initiator Tachyon passes the FCP_RSP to the host via the SFS channel. The initiator Tachyon sends an inbound_scsi_status_completion message to the inititator host. This informs the initiator host that the exchange is completed. The initiator host must clear the STE Valid bit to zero which indicates that the Inbound SEST Entry is no longer valid.

The FCP_RSP received by Tachyon as an initiator must be a single frame sequence. If the FCP_RSP is a multiframe sequence (e.g., a large FCP_RSP is sent through a Fabric that only supports small frame sizes), the frames are processed as bad SCSI frames. It is a violation of the FCSI SCSI Profile to send the FCP_RSP as a multiframe sequence.

The Byte Count in the Inbound SEST Entry is updated with the number of bytes received by Tachyon. The initiator host should check this field to detect any lost data.

3.8.9 FCP Read Exchange - Target Tachyon

For FCP Read exchanges for the target Tachyon, SCSI hardware assists are not used.

The target host must set the Exchange Context bit in the F_CTL field of the FC Header to one to indicate that it is the responder of the exchange.

Command Phase

The target Tachyon receives an FCP_CMND for an FCP Read from the initiator. The target Tachyon sends this FCP_CMND to the target host via the SFS channel. If the SCSI Command Auto ACK bit of the Tachyon Configuration register is set to one, the target Tachyon immeditately returns an ACK (for Class 1 and Class 2) to the initiator Tachyon.

If the SCSI Command Auto ACK bit is cleared to zero, then the target host is responsible for sending the ACK. The target host selects a valid, unused RX_ID value. Refer to " Choosing X_ID Values" on page 160. Even though this value points to an unused entry in the SEST, Tachyon does not use the SEST for this operation (SCSI harware assists are not used for FCP Read for the target). The RX_ID is placed into the header of the ACK and sent via the HPCQ.

Data Phase

The target host builds an ODB, which contains the EDB address. The target host builds an EDB that defines where the Read data is located in the target host memory.

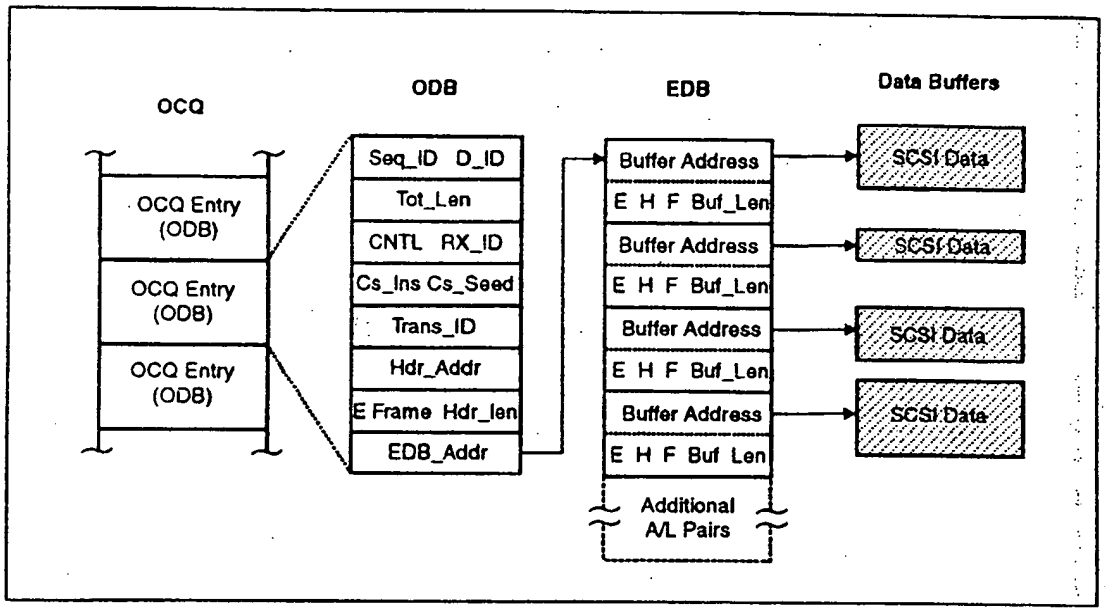

Figure 3.19 FCP Read Exchange - Target Host Data Structures

PTI 172510

The target host may send an FCP_XFER_RDY to the initiator host to indicate that it is ready to send the requested data. The target Tachyon sends the FCP_XFER_RDY(s) with the appropriate RX_ID value to the initiator Tachyon. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

If a problem occurs while sending the FCP_XFER_RDY for the Read data, the target Tachyon notifies the target host via an outbound_completion message. At that point, Tachyon's outbound channel freezes. The target host should perform the appropriate error recovery and unfreeze the outbound channel by writing to the Tachyon Control register.

Using the OCQ, the target Tachyon then sends the appropriate SCSI Read data to the initiator Tachyon.

Status Phase

The target Tachyon sends an FCP_RSP to the initiator: Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

PTI 172511

3.8.10 FCP Write Exchange Example

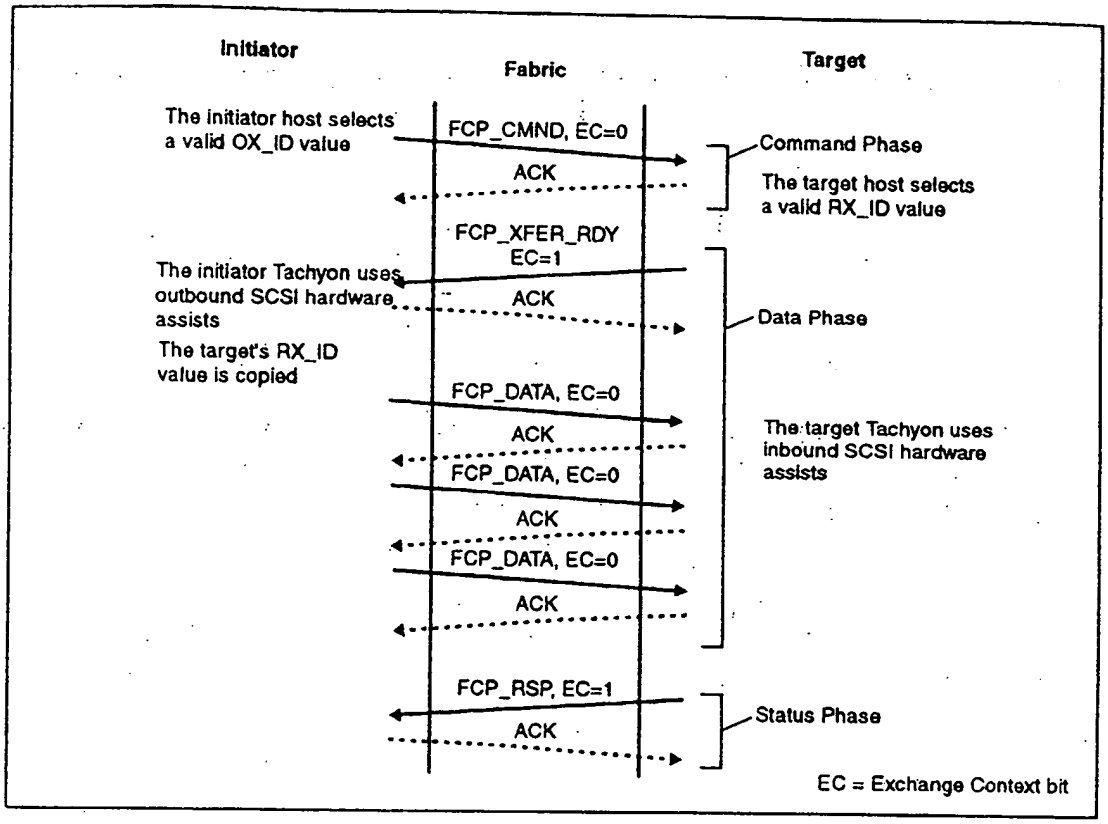

Figure 3.20 FCP Write Exchange Example

PTI 172512

3.8.11 FCP Write Exchange - Initiator Tachyon

For the initiator host to perform an outbound data transfer, it selects a valid OX_ID value. Refer to " Choosing X_ID Values" on page 160. Using the OX_ID value, the initiator host builds an Outbound SEST Entry. The Outbound SEST Entry includes information about the frame size, the class to send the sequence with, and writeable fields that the initiator Tachyon uses to manage the SCSI transfer. The Outbound SEST Entry also contains a pointer to the EDB.

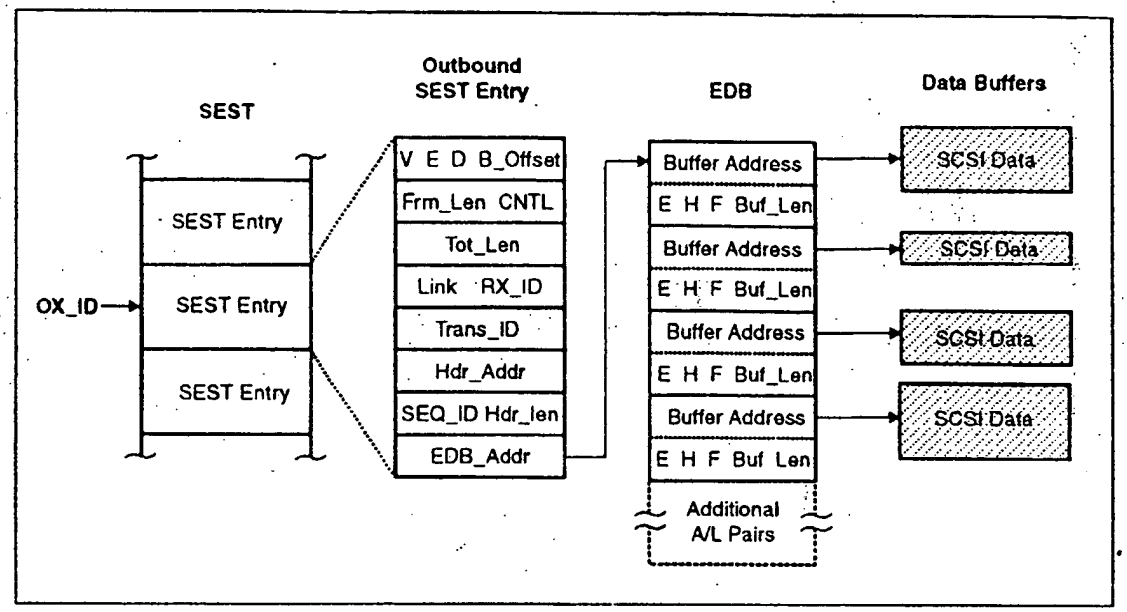

Figure 3.21 FCP Write Exchange - Initiator Host Data Structures

The initiator host clears the Exchange Context bit in the F_CTL field of the FC Header to zero since it is the originator of the exchange. Also, the Fill bits in the F_CTL field are set appropriately for non-word multiple outbound SEST transfers.

Command Phase

Once the host creates the Outbound SEST Entry, the host creates the FCP_CMND for an FCP Write exchange. The initiator Tachyon sends the FCP_CMND with a valid OX_ID to the target. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

Data Phase

When the initiator Tachyon receives the FCP_XFER_RDY from the target, it checks the value of the DATA_RO field. If the DATA_RO is zero, then this FCP_XFER_RDY indicates a single data phase, or the first of a multiple data phase transfer. The initiator Tachyon uses its SCSI hardware assists and manages the data phase for this FCP_XFER_RDY, independent of the host. Tachyon uses the information in the SEST and the FCP_XFER_RDY to build an ODB to be sent through the OSM. If the OSM is busy, the SCSI Exchange State Machine adds the request to a linked list of outbound transactions waiting for transmission. As the OSM becomes available to process a SCSI transfer, the SCSI Exchange State Manager de-queues a waiting transaction and passes it to the OSM. The OSM transmits the Write data to the target Tachyon.

PTI 172513

The data is sent (using the appropriate RX_ID of the received FCP_XFER_RDY) to the address of the target destination that sent the FCP_XFER_RDY (using the S_ID of the received FCP_XFER_RDY).

If the DATA_RO is non-zero, then this FCP_XFER_RDY is part of a multiple data phase transfer. The initiator Tachyon passes this FCP_XFER_RDY as an SFS to the initiator host along with a bad_scsi_frame completion message. The initiator host is responsible for managing the data phase for this multiple data phase transfer by using the general sequence moving services. The OX_ID field in the FCP_XFER_RDY is an index into the SEST and identifies the appropriate Outbound SEST Entry which points to the EDB where the Write data is located. Using the information in the Outbound SEST Entry, the initiator host builds an ODB. The initiator Tachyon uses the ODB to transmit the Write data to the target.

Status Phase

When the data phase completes, the initiator Tachyon receives an FCP_RSP from the target. The initiator Tachyon passes the FCP_RSP to the host via the SFS channel. The initiator Tachyon sends an inbound_scsi_status_completion message to the inititator host. This informs the initiator host that the exchange completed. The OX_ID field in the FCP_RSP is an index into the SEST and identifies the appropriate Outbound SEST Entry. The initiator host must clear the STE Valid bit to zero which indicates that the Outbound SEST Entry in no longer valid.

The initiator host determines a successful transfer by the FCP_RESID field in the FCP_RSP. The FCP_RESID field indicates the number of bytes not transferred successfully. For a successful transfer, the value of the FCP_RESID field is zero. Refer to "FCP, Section 7.4.2".

3.8.12 FCP Write, OOO Reassembly - Target Tachyon

OOO Reassembly mode enables Tachyon to receive and reassemble OOO sequences for inbound exchanges.

To use OOO Reassembly mode:

- The OOO Reassembly Disable bit of the Tachyon Configuration register must be cleared to zero at initialization.
- SCSI Buffer Length must be a power of 2 bytes.

The target host must set the Exchange Context bit in the F_CTL field of the Fibre Channel header to one to indicate that it is the responder of the exchange.

Command Phase

The target Tachyon receives an FCP_CMND for a FCP Write from the initiator. If the SCSI Command Auto ACK bit of the Tachyon Configuration register is set to one, the target Tachyon immediately returns an ACK to the initiator Tachyon, for Class 1 and Class 2.

If the SCSI Command Auto ACK bit is cleared to zero, then the target host is responsible for sending the ACK. The target host selects a valid RX_ID value. Refer to " Choosing X_ID Values" on page 160. The RX_ID is placed into the ACK header and sent via the HPCQ.

Data Phase

If it has not already done so, the target host selects a valid RX_ID value. Using the RX_ID value, the target host builds the Inbound SEST Entry which points to the SDB. The SDB defines host buffers where the Write data is going to be received. For OOO Reassembly, the SDB contains up to 64 pointers to fixed length buffers and the SDB must be aligned on a 256-byte boundary.

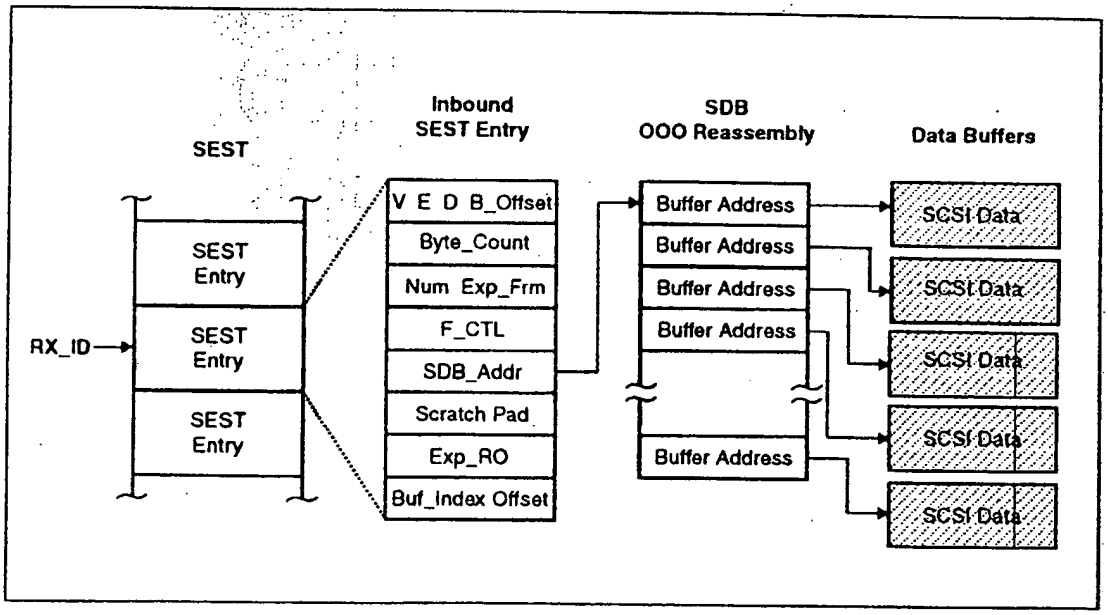

Figure 3.22 FCP Write Exchange - Target Host Data Structures

PTI 172515

If the target host has enough buffers and is ready to receive all of the data from the initiator, the target Tachyon sends the FCP_XFER_RDY with the DATA_RO field cleared to zero. This is a single data phase transfer.

If the target host does not have enough buffers to receive the entire transfer or it does not want all of the data immediately, the target Tachyon sends an FCP_XFER_RDY with BURST_LEN set to less than the total amount of data. This FCP_XFER_RDY indicates a multiple data phase transfer. The DATA_RO field is set to the offset of the data block requested.

The target Tachyon sends the FCP_XFER_RDY(s) with the appropriate RX_ID value to the initiator. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

As each frame is received, the target Tachyon checks the RX_ID of the data frame to locate the appropriate Inbound SEST Entry. Tachyon does not examine the TYPE field in the FC Header.

When the last data frame is received, the target Tachyon sends an inbound_scsi_data_completion message to the target host. The inbound_scsi_data_completion message informs the target host that all frames for the SCSI sequence have been received.

Status Phase

The target host then sends an FCP _RSP to the initiator using one of the small data payload transmit methods. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

3.8.13 FCP Write Exchange, In Order Reassembly - Target Tachyon

For In Order Reassembly mode, frames within a sequence must arrive in order. For the target Tachyon for In Order Reassembly, sequences within the data phase of the same exchange may arrive out of order.

To use In Order Reassembly mode:

- The OOO Reassembly Disable bit of the Tachyon Configuration register must be set to one at initialization.
- Tachyon must be operating in an in order topology, i.e., point-to-point, a closed arbitrated loop, a fabric which supports guaranteed in order delivery (negotiated at FLOGIn), or an open arbitrated loop using a fabric that supports in order delivery.
- The SCSI Buffer Length must be a multiple of 4 bytes.
- The target host must clear the Buffer Offset and Buffer Index fields of the Inbound SEST Entry to zero.

The target host must set the Exchange Context bit in the F_CTL field of the FC Header to one to indicate that it is the responder of the exchange.

Command Phase

The target Tachyon receives an FCP_CMND for a FCP Write from the initiator. If the SCSI Command Auto ACK bit of the Tachyon Configuration register is set to one, the target Tachyon immeditately returns an ACK to the initiator, for Class 1 and Class 2.

If the SCSI Command Auto ACK bit is cleared to zero, then the target host is responsible for sending the ACK. The target host selects a valid RX_ID value. The RX_ID is placed into the ACK header and sent via the HPCQ.

Data Phase

If it has not already done so, the target host selects a valid RX_ID value. Using the RX_ID value, the target host builds the Inbound SEST Entry that points to the SDB. The SDB contains as many buffer addresses as necessary to receive the data. For In Order Reassembly, the buffer addresses must be on a word-aligned boundary, the SDB must be aligned on a power of 2 boundary equal to or greater than the SDB length. The SDB length is the number of entries multiplied by 4 bytes.

If the target host has enough buffers and is ready to receive all of the data from the initiator, then the DATA_RO field in the FCP_XFER_RDY is cleared to zero and the BURST_LEN field is set to the entire amount of data for the transfer. This indicates a single data phase transfer.

Before the FCP_XFER_RDY is sent to the initiator, the target host writes this DATA_RO value in the Exp_RO field of its Inbound SEST Entry. The target Tachyon uses this value to verify that the receiving frames are in order. The target Tachyon then sends the FCP_XFER_RDY to the initiator.

If the target host does not have enough buffers to receive the entire transfer or it does not want all of the data immediately, the target Tachyon sends an FCP_XFER_RDY with BURST_LEN set to less than the total amount of data. This FCP_XFER_RDY indicates a multiple data phase transfer. The DATA_RO field is set to the offset of the data block requested.

The target Tachyon sends the FCP_XFER_RDY(s) with the appropriate RX_ID value to the initiator. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

PTI 172517

As each data frame is received, target Tachyon uses the RO value to determine if the frame is in order for the sequence. If it is not in order, Tachyon passes the frame as a bad SCSI frame to the host via the SFSBQ, invalidates the SCSI table entry by clearing the Inbound SEST Entry to zero, and sends a bad_scsi_frame completion message to the host indicating that an OOO SCSI frame was received erroneously. Refer to "3.3.4 SCSI Completion Messages" on page 31. If the host has set the Bad SCSI Auto ACK bit in the Tachyon Configuration register to one, Tachyon generates an ACK (for Class 1 and Class 2) for the bad SCSI frame and the host is responsible for the error recovery. If the host has cleared the Bad SCSI Auto ACK bit to zero, the host is responsible for the ACK as well as for the error recovery. Tachyon manages all subsequent data frames for that exchange as bad SCSI frames and passes these frames to the host with a bad_scsi_completion message.

If the data frame is in order for the sequence, the target Tachyon checks the RX_ID of the data frame to locate the appropriate Inbound SEST Entry. The Inbound SEST Entry helps Tachyon to determine exactly where within the buffers the data is to be placed. Tachyon does not examine the TYPE field in the Fibre Channel Header.

When the last data frame is received, the target Tachyon sends an inbound_scsi_data_completion message to the target host. The inbound_scsi_data_completion message informs the target host that all frames for the SCSI sequence have been received.

Status Phase

The target host then sends an FCP_RSP to the initiator using one of the small data payload transmit methods. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

PTI 172518

3.9 Loop Details

3.9.1 FC-AL Specification Deviations

Tachyon does not support the following:

1. LIRP and LILP, the Loop Initialization Sequences that carry a 128-byte AL_PA position map. Byte 2 of the LISA Loop Initialization Identifier is cleared to 0x00 to bypass these sequences.
2. The transmission or reception (in any loop state) of N_Port Primitive Sequences; NOS, OLS, LR, and LRR. If Tachyon receives a NOS or OLS after the link is established, it sets the NOS/OLS Received bit in the Frame Manager Status register to one. Then Tachyon sends a frame_mgr_interrupt completion message to the host if Tachyon is in the OFFLINE, OLD_PORT, or INITIALIZE state.
3. The reception of the LPE and LPB Primitive Sequences to Tachyon. Tachyon supports re-transmission of these Primitive Sequences to other devices. To program Tachyon to transmit these Primitive Sequences, refer to "3.9.11 HOST CONTROL State" on page 88.
4. Open Primitive Signal - selective replicate (OPNyr).
5. Maintaining a minimum of six words for an interframe gap. Tachyon reduces the interframe gap to a minimum of two words if its elasticity buffer is approaching an error condition which would force re-initialization of the loop.

Tachyon may exceed FC-AL Specifications for 266 and 531 MBaud data rates. The node delay through Tachyon may be five and a half words before the delete mechanism in the elasticity buffer attempts to remove one word. If a one word delay is included for the GLM, then this condition exceeds the six word specification in FC-AL. This condition does not occur for a 1062 MBaud data rate.

3.9.2 Power-On

At power-on or after a software reset, Tachyon asserts the Lock to Reference signal (a PLM Interface signal) and holds its receiver in reset to 10 msec. During this time, Tachyon sources Idles from its transmitter. After 10 msec, Tachyon de-asserts the Lock to Reference signal. The Frame Manager is initialized to the Offline state.

3.9.3 Online

If Tachyon has acquired synchronization, it forwards valid Ordered Sets that it receives to the transmitter. If Tachyon has not acquired synchronization or receives an invalid Ordered Set, Tachyon sources Idles from the transmitter.

Before the host writes the Initialize command (instructs Tachyon to go online) to Tachyon, the host should read the Frame Manager Status register to determine if a failure condition (e.g., Laser Fault, Out of Synchronization, or Loss of Signal) exists. After the host writes the Initialize command, Tachyon waits for any failure conditions to clear before attempting to initialize. Any failure conditions which exist prior to the attempt to initialize are not reported. It is possible for a failure to occur between the time the Frame Manager Status register is read and the write required to program Tachyon to initialize.

3.9.4 Initialization

To perform Loop Initialization, the host sets the appropriate address bits that correspond to the Loop Initialization phase in which Tachyon wants to acquire an address. Address bits fa, aq, ha, and sa in the Frame Manager Configuration register correspond to the Loop Initialization phases LIFA, LIPA, LIHA, LISA, respectively.

Loop Initialization Phase	Description	Frame Manager Configuration Register	
		Field Label	Bit
LIFA	Loop Initialization Fabric Acquired	fa	10
LIPA	Loop Initialization Previously Acquired	aq	9
LIHA	Loop Initialization Hard (Preferred) Address	ha	8
LISA	Loop Initialization Soft Address	sa	7

Table 3.12 Loop Initialization Bits

The host should set only one of these bits: aq, ha, or sa bits. For more information about the fa bit, refer to "3.9.15 Fabric Operation" on page 90. If the host sets the aq or ha bit, it must also provide a valid AL_PA in the Frame Manager Configuration register since Tachyon attempts to acquire the AL_PA in this register. For valid AL_PA values, refer to the "FC-AL, Version 4.5, Section 10.4.2, Table 15, p.53". If the aq or ha bit is set and the AL_PA in the Frame Manager Configuration register is not available in the LIPA or LIHA phase, respectively, or the host sets the sa bit, Tachyon attempts to acquire an address in the LISA phase. In this case, Tachyon selects the first available bit in the AL_PA bit map which corresponds to the lowest AL_PA available. When Tachyon acquires an AL_PA, it sets the aq bit and clears the ha and sa bits if they were set.

Tachyon generates a frame_mgr_interrupt completion message to the host. At that point, the host can determine the status of the initialization by reading the Frame Manager Status register. If loop initialization succeeds, the lfsm and pfsm fields are set to MONITOR and OFFLINE, respectively. (If N_Port initialization succeeds, the lfsm and pfsm fields are set to OLD_PORT and ACTIVE, respectively.) Also, the Link Up (lup) bit is set. The host should verify that the Non-Participating bit is cleared to zero, indicating that Tachyon was successful in acquiring an AL_PA.

The host must also verify the AL_PA acquired during the initialization process. If the host set either the aq or ha bit and the AL_PA acquired during initialization was not the AL_PA programmed by the host, the host may use the AL_PA that was acquired anyway, may try again to acquire the desired AL_PA, or it may enter the non-participating mode. If the host decides to enter the non-participating mode, the host must force an initialization sequence (Refer to "FC-AL, Version 4.5, Section 10.1, p. 50").

Internal Loopback Mode

- When Tachyon is initialized in internal loopback mode, the Fabric Login Required bit (in the Frame Manager Status register) may be set to one after initialization completes.
- An internal loopback test failure may occur if the Receive Byte Clock (RBC) violates GLM RBC specifications. Use of a loopback hood guarantees that the RBC meets GLM specifications. It is not necessary to have a GLM present to perform an internal loopback test. If RBC is being violated, remove the GLM before performing the test.

PTI 172520

3.9.5 Loop Situations After Initialization

LIP Not Received

If a Loop Initialization Primitive (LIP) is not received within two AL_TIMEs after starting initialization, Tachyon attempts to initialize as an N_Port (Refer to "FC-AL, Version 4.5, Seciton 8.4.3.21, p. 35"). When this occurs, the loop state (lsfm) in the Frame Manager Status register indicates OLD_PORT.

Initializing After Initialization

If the host attempts to initialize the loop after it has already been initialized, a link down condition occurs. Tachyon sends a frame_mgr_interrupt completion message and immediately freezes the OSM so no additional frames are sent. When the OSM error is cleared, the OSM sends the next queued entry unless the OCQ was reset.

After the host resets the OCQ and re-queues sequences, the OSM can send the sequences. One reason that the host may decide to reset the OCQ is that once the loop goes down, the device the host was trying to access may have acquired a new AL_PA, or the device may have gone offline and another device may have aquired its AL_PA. The host would then send data to a non-existent device, or to the wrong device. The PDISC extended link service command can be used to determine if a device has changed without changing the operating parameters of the device.

Drive is Hot-Plugged

If the loop is initialized and a drive on the loop is hot-plugged, Tachyon performs the following:

1. Tachyon sends a frame_mgr_interrupt completion message
2. If an outbound sequence was in progress, Tachyon sends an outbound_completion message with the
3. Link Down (LD) bit set to one. If a networking (non-SCSI) inbound sequence was in progress, Tachyon sends an inbound_mfs completion message with the Link Down (LD) bit set to one. If an active inbound SCSI sequnce was in progress, Tachyon does not send a completion message indicating the Link Down condition.
4. Tachyon sets the Link Down (ldwn) bit in the Frame Manager Status register to one.
5. If the drive performs a Loop Initialization (LIP) to acquire an AL_PA, Tachyon manages the link (re)initialization (LIP, Loop initialization frames, etc.) transparently. Tachyon then interrupts the host to indicate that the link is up.

3.9.6 Opening a Device

This description assumes a Class 3 exchange between Tachyon (local node) and a remote node.

Tachyon uses full-duplex for OPNs except for open broadcast. When loop access is available, Tachyon arbitrates for ownership of the loop. If the host sets the Unfair Access bit in the Tachyon Header Structure, Tachyon can arbitrate regardless of loop access availability.

After Tachyon wins arbitration, it establishes a loop circuit by sending an OPNyx Primitive Signal. For OPNyx, 'y' is equal to the AL_PA value in the Tachyon Header Structure (referring to the remote node) and 'x' is equal to the AL_PA value in the Frame Manager Configuration register (referring to the local node).

PTI 172521

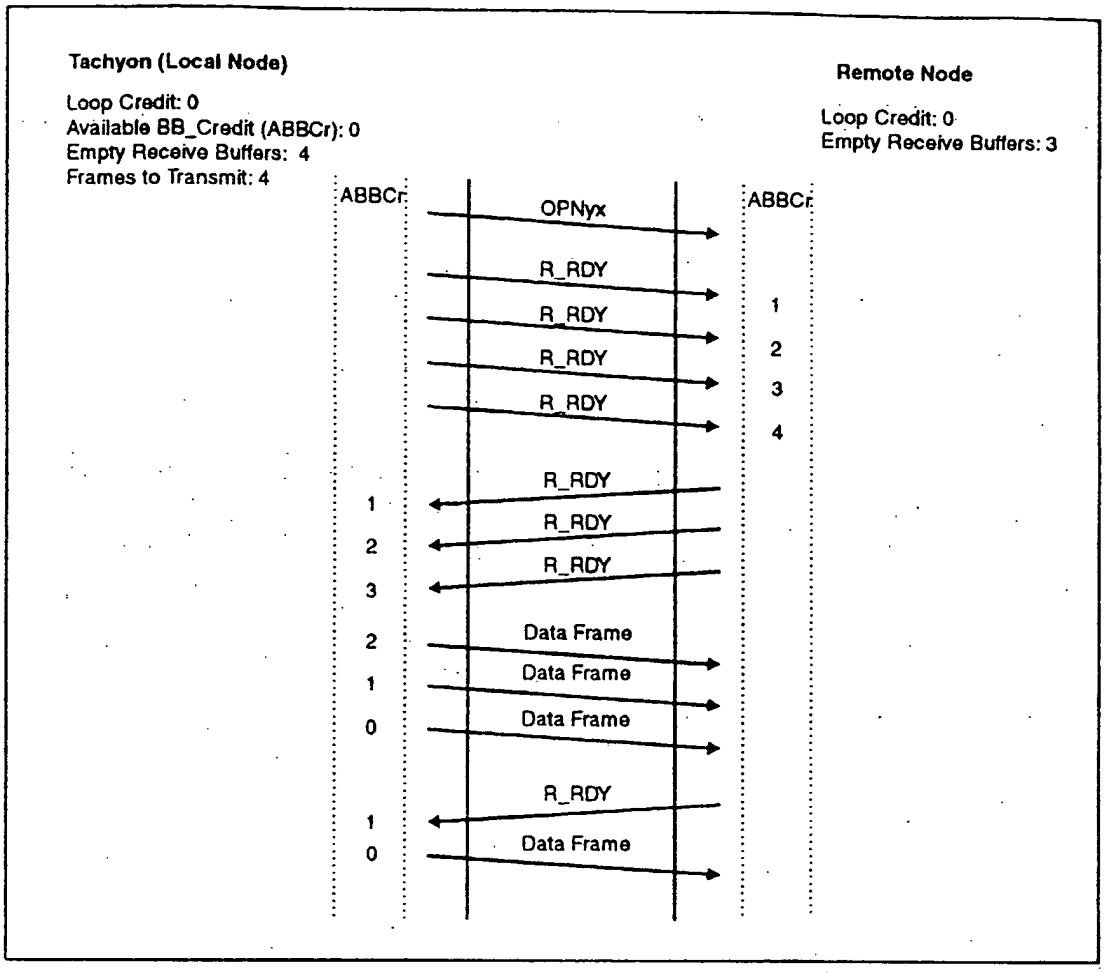

Figure 3.23 Remote Node Advertises Zero BB_Credit at Login

All Tachyon-based devices advertise a Loop Credit (indicated in the Tachyon Header Structure) of zero. Loop Credit is Tachyon's Login BB_Credit. Because Tachyon's advertised Loop Credit is zero, Tachyon sends one R_RDY for each of its empty receive buffers to enable the remote node to send data frames to Tachyon. Tachyon has up to four empty receive buffers and may send up to four R_RDYs to the remote node.

Refer to "Figure 3.23 Remote Node Advertises Zero BB_Credit at Login" on page 84. Tachyon has four empty receive buffers and sends four R_RDYs before sending data frames to the remote node. If the remote node advertises a Loop Credit of zero at Login, then it must send R_RDYs to Tachyon before it receives data frames from Tachyon.

PTI 172522

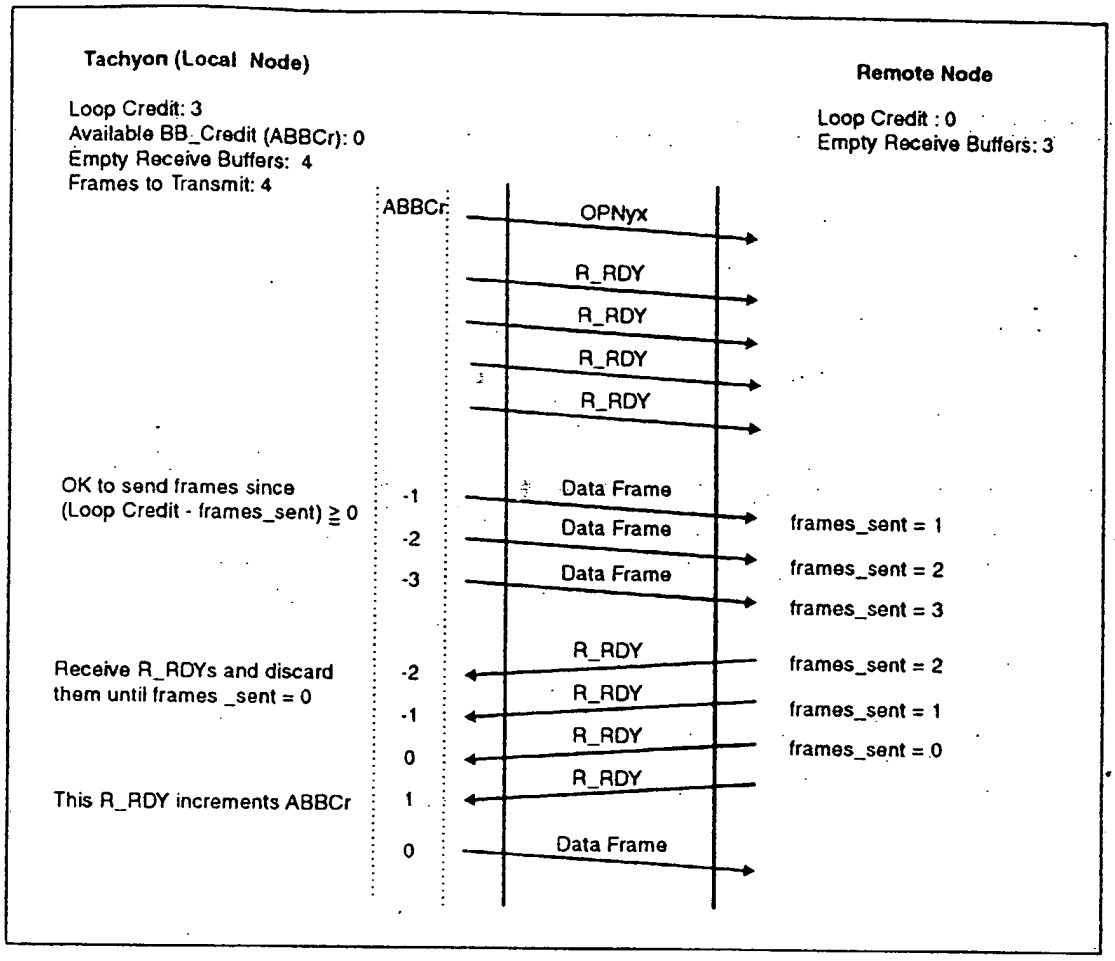

Figure 3.24 Remote Node Advertises Non-Zero BB_Credit at Login

If the remote node advertises a Loop Credit greater than zero, then it guarantees that it has empty receive buffers equal to its Login BB_Credit. Tachyon can send data frames to the remote node equal to the remote node's Login BB_Credit, without waiting for an R_RDY from that remote node. After Tachyon sends the data frames, it waits for the remote node to return that number of R_RDYs. Tachyon discards the Login BB_Credit worth of R_RDYs from the remote node, and does not transmit more than Login BB_Credit frames until an R_RDY, in addition to those discarded R_RDYs, is received from the remote node.

Refer to "Figure 3.24 Remote Node Advertises Non-Zero BB_Credit at Login" on page 85. When the loop circuit is established and the remote node advertises a Loop Credit of three, Tachyon can immediately send three data frames to the remote node. To send additional frames, Tachyon first needs to receive three R_RDYs from the remote node, which correspond to the three frames that Tachyon sent. These frames are discarded by Tachyon. Then the remote node sends another R_RDY so that Tachyon can send an additional frame.

PTI 172523

After Tachyon sends the last frame of the last sequence, it transmits a CLS Primitive Signal and transitions to the XMITTED CLOSE state. Optionally, if the CL bit in the Tachyon Header Structure is set and no more data is immediately available to send to the remote node, Tachyon transmits a CLS and transitions to the XMITTED CLOSE state after each frame.

If the host attempts to open a device and the OPN becomes corrupted, it is possible for the host to receive its own data. Tachyon sends this data to its host as either a SFS or MFS. The host should always check the D_ID to verify the frame is meant for itself. If it discovers that the D_ID is meant for another device, it should check the S_ID to verify it did not the transmit the frame. If it verifies that it sent the frame, then the OPN may have been corrupted. The host should re-send the sequence.

3.9.7 Open Broadcast Replicate Support

Open Broadcast Replicate Transmission

For transmitting L_Port open broadcast replicate frames to all devices using the Open Primitive Signal - broadcast replicate (OPNfr), the following procedure is used:

1. Set the AL_PA in the Tachyon Header Structure to 0xFF.
2. Use SOFi3.
3. Set the LCr field in the Tachyon Header Structure greater than or equal to one.
4. Send single frame sequences. Though not recommended, the host may send multiframe sequences. If the host sends a multiframe sequence, the LCr must have enough credit to send the entire MFS. Credit is not returned from the remote receiving the MFS replicate. Therefore, if LCr is set to less than the number of frames that Tachyon wants to send, a BB_Credit Time-out error occurs.

Open Broadcast Replicate Reception

The Frame Manager reads the AL_PA of the incoming frame. If the AL_PAs is set to 0xff, then the Frame Manager accepts the open broadcast frame and forwards it to the Inbound Data FIFO. Also, the Frame Manager forwards-on the open broadcast frame on the link.

3.9.8 Open Selective Replicate (Multicast)

Tachyon does not support Open Selective Replicate (OPNyr) since this requires that the AL_PA is set to a value other than 0xFF and requires multiple OPNyr(s) per sequence.

3.9.9 TRANSFER (XFER) State

When Tachyon's ACK FIFO is full, Tachyon arbitrates unfairly and uses the TRANSFER (XFER) state. This is the only case that Tachyon uses the XFER state. The XFER state allows Tachyon to perform multiple OPN/CLS sequences in one ARB tendancy.

3.9.10 OPENED State

In the MONITORING state, when Tachyon receives an OPN Primitive Signal addressed to its AL_PA, it transitions to the OPENED state. If Tachyon's Inbound Frame FIFO is full, Tachyon immediately sends a CLS and transitions to the XMITTED_CLOSE state without sending R_RDYs to the L_PORT that is attempting to open Tachyon.

If the Inbound Data FIFO is not full when Tachyon enters the OPENED state, Tachyon sends a number of R_RDYs equal to the number of empty inbound receive buffers available. If the remote node opens Tachyon half-duplex, Tachyon receives frames and only transmits Link Control frames. If the remote node opens Tachyon full-duplex, Tachyon receives frames and transmits both Link Control and data frames. Tachyon transmits data frames only if data frames are available (at the head of the Outbound Frame FIFO) for the port that opened it. Tachyon assumes that the L_Port that established the Loop circuit will initiate the CLS.

When Tachyon receives a CLS, Tachyon transitions to the RECEIVED_CLOSE state. In the RECEIVED_CLOSE state, Tachyon sends all of the frames that are destined for the other L_Port. Tachyon sends frames until BB_Credit is exhausted. When BB_Credit is exhausted or there are no more frames to send, Tachyon forwards the CLS and then transitions to the MONITORING state.

3.9.11 HOST CONTROL State

The host uses the Host Control command in the Frame Manager Status register to support transmission of a Directed Reset, LPE, and LPB.

Directed Reset Generation

1. The host writes the Frame Manager Primitive register with a LIP having the correct values for characters three and four. In the Frame Manager Control register, the host sets the Primitive Sequence bit and the Send Prim_Reg bit.
2. The Frame Manager detects the Host Control command in the INITIALIZING, OPEN_INIT, or MONITORING state. If the Frame Manager was in the MONITORING state, the Frame Manager arbitrates and gains ownership of the loop regardless of fairness prior to transitioning to the HOST_CONTROL state. In the INITIALIZING or OPEN_INIT state, the Frame Manager transitions directly to the HOST_CONTROL state. In either case, the Frame Manager clears the Host Control command on entry, sends at least 12 occurrences of the primitive in the Frame Manager Primitive register, clears the Send Prim_Reg bit, and sends a frame_mgr_interrupt completion message to the host.
3. The Frame Manager continues to send the primitive until it detects the LIP in the Frame Manager Primitive register. When the Frame Manager detects the LIP, it sends a frame_mgr_interrupt completion message to the host.
4. The host writes the Initialize command to the Frame Manager. The Frame Manager exits the HOST_CONTROL state. If the Frame Manager receives a LIP other than the one transmitted in the Frame Manager Primitive register, the Frame Manager sends a frame_mgr_interrupt completion message to the host and sets the LIPf bit in the Frame Manager Status register. A LIPf is a LIP indicating a loop failure; LIP(F8, AL_PS) or LIP(F8,F7). When the Frame Manager receives a LIP other than the one transmitted in the Frame Manager Primitive register, it is the host's responsibility to determine how long to transmit the directed reset.

LPE/LPB Generation

1. The host programs the Frame Manager Primitive register with the LPE/LPB to send. In the Frame Manager Control register, the host sets the Primitive Sequence bit, sets the Send Prim_Reg bit, and programs the Host Control command.
2. The Frame Manager detects the Host Control command in the MONITORING, INITIALIZING, or OPEN_INIT state. If the Frame Manager was in the MONITORING state, the Frame Manager transitions to the ARBITRATING state and wins ownership of the loop regardless of fairness. The Frame Manager enters the HOST_CONTROL state, clears the Host Control command, transmits at least 12 occurrences of the primitive in the Frame Manager Primitive register, clears the Send Prim_Reg bit and sends a frame_mgr_interrupt completion message to the host.
3. The Frame Manager continues to send the primitive until it receives the LPE/LPB back. The Frame Manager then sends a frame_mgr_interrupt completion message to the host. If the Frame Manager receives a LIP, the Frame Manager generates an additional frame_mgr_interrupt completion message and sets the LIPf bit in the Frame Manager Status register. In this way, Tachyon notifies the host that it has received a LIP, but Tachyon, itself, does not respond to the LIP.
4. The host then writes the Exit Host Control command to the Frame Manager Control register provided it has not received a frame_mgr_interrupt completion message due to a LIP. If the host received a frame_mgr_interrupt completion message due to a LIP, the host writes the Initialize command. The Frame Manager detects the Exit Host Control or Initialize command and exits the HOST_CONTROL state.

PTI 172526

3.9.12 Programming the Frame Manager OFFLINE

If Tachyon is on a loop and the host programs the Offline command in the Frame Manager Control register, the loop may go into an unknown state because the host cannot control the exact time that Tachyon actually goes OFFLINE.

To take Tachyon OFFLINE (non-participating) on a loop, perform the following:

1. In the Frame Manager Configuration register, clear the fa, pa, ha, and sa bits to zero.
2. In the Frame Manager Control register, program the Initialize command.
3. In the Frame Manager Control register, program the Offline command.

When none of the fa, pa, ha, and sa bits in the Frame Manager Configuration register are set, initialization behaves normally, but when it completes, Tachyon is in the non-participating state. The non-participating bit in the Frame Manager Status register is set. In addition, if no errors occur during initialization and if loop initialization succeeds, the lfsm and pfsm fields in the Frame Manager Status register are set to MONITOR and OFFLINE, respectively. If no errors occur during initialization and if N_Port initialization succeeds, the lfsm and pfsm fields are set to OLD_PORT and ACTIVE, respectively.

If the host has initialized on the loop, it should not program the Offline command to Tachyon. Instead, if it does not wish to participate in the loop, it should relinquish its AL_PA by re-initializing in non-participating mode. After non-participating mode is established, the Offline command may be programmed. The only exception is in internal and external loopback mode. The host should only use these modes immediately after a reset, or after a laser fault error occurs.

3.9.13 Programming the Frame Manager OFFLINE if in LOOP_FAIL State

If the Frame Manager is configured for loop and is in the LOOP FAIL state, writing the Offline command to the Frame Manager Control register does not force the Frame Manager offline.

Workaround to Force Frame Manager Offline

Read the Frame Manager Status register. If the Loop State Machine (lsfm) is in the LOOP_FAIL state (0xd0), then perform the following steps to force the Frame Manager into the OFFLINE state:

1. In the Frame Manager Control register, program the Host Control command to force the Frame Manager into the HOST_CONTROL state.
2. In the Frame Manager Control register, program the Offline command to force the Frame Manager into the OFFLINE state.

3.9.14 Determining the Link State of Tachyon

Because both the Link Up (lup) and Link Down (ldwn) bits in the Frame Manager Status register can be set at the same time, the recommended method for determining the link's state is to test the Loop State Machine (lfsm) and Port State Machine (pfsm) bits in the Frame Manager Status register as follows.

1. If bit 7 of the Frame Manager Status register is cleared to zero, then the link is up.
2. If the four low order bits (bits 3 .. 0) are set to 0xF, then Point-to-Point is active.
3. If bit 7 in "1." above is set to one or the low order bits in "2." above are not set to 0xF, then the link is down.

PTI 172527

3.9.15 Fabric Operation

Port ID Values

During normal Loop operation, the host must determine if the destination N_Port is on the local loop or on the other side of an FL_Port. To determine whether the destination N_Port is on the local loop, the host compares the upper 16 bits of the destination's N_Port ID with the upper 16 bits of its own N_Port ID. If they match, then the destination N_Port is on the local loop and the host will set the AL_PA value in the Tachyon Header Structure to equal the destination's N_Port ID. If the N_Port IDs do not match, then the destination N_Port is not on the local loop, so the host clears the AL_PA value to zero.

Fabric Login

When the host performs Fabric Login (FLOGI) it must first set the Don't Close Loop bit in the Frame Manager Control register to one and then queue a FLOGI sequence to the OCQ. Setting the Don't Close Loop bit causes the Frame Manager not to close the loop after sending the FLOGI sequence. When the host receives the FLOGI Response (ACC) sequence, it must check the assigned Port Identifier to determine if the Fabric has changed the AL_PA value.

When performing Fabric Login, the Fabric may assign the port a different AL_PA value than the one it acquired during Loop Initialization. If this occurs, two ports could have the same AL_PA value. To prevent duplicate AL_PA values, this port must Login with the Fabric without terminating the Loop Circuit. If the AL_PA value is changed, the port must transition to the INTIALIZATION state without terminating the Loop Circuit (i.e., don't give the other port with the same AL_PA a chance to establish a Loop Circuit and send frames).

If the Fabric has not changed the AL_PA value, then the host must do the following:

1. In the Frame Manager Configuration register, set the fa and aq bits to one and clear the ha and sa bits to zero.
2. In the Frame Manager Control register, set the cl bit to one and clear the Don't Close Loop bit to zero.

If the Fabric has changed the AL_PA value, then the host must do the following:

1. In the Frame Manager Configuration register, set the fa and aq bits to one and clear the ha and sa bits to zero.
2. In the Frame Manager Configuration register , write the Fabric assigned AL_PA value to the AL_PA field.
3. In the Frame Manager Control register, clear the Don't Close Loop bit to zero.
4. In the Frame Manager Control register, program the Initialize command. This causes the Frame Manager to transition to the INTIALIZATION state and start the Loop Initialization process.

If the host knows that no Fabric is connected to this loop, and the host wishes to act as the Fabric, it should do the following:

1. In the Frame Manager Configuration register, set the rf, fa, aq, if, and lr bits to one.
2. In the Frame Manager Control register, program the Initialize command.

Tachyon becomes Loop Master and causes all ports to re-acquire their address and login with the Fabric. When Loop Initialization completes, other devices can perform FLOGI with this port.

At the end of INTIALIZATION, the Frame Manager does not reset the Login Required (lr) bit in the Frame Manager Configuration register. The host should reset the lr bit immediately after receiving the frame_mgr_interrupt completion message.

If a Fabric is present and the host programs the Initialize as Fabric (if) bit in the Frame Manager Configuration register, loop operation will be unpredictable.

PTI 172528

3.9.16 Loop Operation While Impersonating a Fabric

Tachyon may want to initialize as a Fabric (impersonate a Fabric) if it wants to build an N_Port capable of responding to server addresses (e.g., Name Server, Time Server, etc.).

If a Fabric exists on the loop, Tachyon should not initialize as a Fabric. Only one Fabric (the Fabric with the higher priority, i.e., the Fabric with the lower World Wide Name) becomes Loop Master during Loop Initialization. If the host programs Tachyon with Initialize as Fabric and Tachyon has a lower World Wide Name, then Tachyon becomes the Loop Master instead of the "real" Fabric. Then, during Loop Initialization, if Tachyon sees that another Fabric is on the loop and assumes that this Fabric is a "real" Fabric, Tachyon gives up being Loop Master. When this occurs, no Loop Master exists and Loop Initialization does not complete.

PTI 172529

3.10 TCP/UDP Hardware Assists

Tachyon provides TCP/UDP hardware assists to enhance networking performance. These assists offload the host from most or all of the following tasks:

1. Checksum calculation for outbound networking data
2. Checksum calculation for inbound networking data
3. Splitting of inbound networking header and data into separate buffers in host memory

3.10.1 Checksumming

Tachyon provides varying degrees of hardware assistance for both TCP and UDP checksumming for both outbound and inbound sequences. For outbound sequences, Tachyon, with assistance from the host, calculates an exact checksum and inserts it at the proper place in the sequence. For inbound sequences, Tachyon calculates a total sum of all of the payload for non-SCSI sequences. To arrive at the true packet checksum, the host must subtract from the total checksum any network headers that should not have been included in the total checksum. Tachyon does not include the Tachyon Header Structure in the checksum calculation.

Tachyon does not support checksumming across IP fragments.

3.10.2 Outbound Checksum Requirements

Tachyon supports both header and trailer checksums. To perform a header checksum, the host must use a Fibre Channel feature called Random Relative Offset (RRO). Not all Fibre Channel nodes may support RRO, and only sequences destined for nodes who do support RRO can use header checksum assists.

If the remote node does not support RRO, it may support trailers. This capability is negotiated at the TCP level. If the node does not support trailers or RRO, the host cannot take advantage of hardware checksum assists and the host must calculate the checksum in software. UDP does not use trailers; therefore, complete reliance on trailers using applications such as NFS have poor performance unless RRO is used.

Tachyon allows outbound header checksum assists only on sequences that are a multiple of four bytes in length. In this mode the host cannot use the Fill bytes in the F_CTL field in the Tachyon Header Structure.

PTI 172530

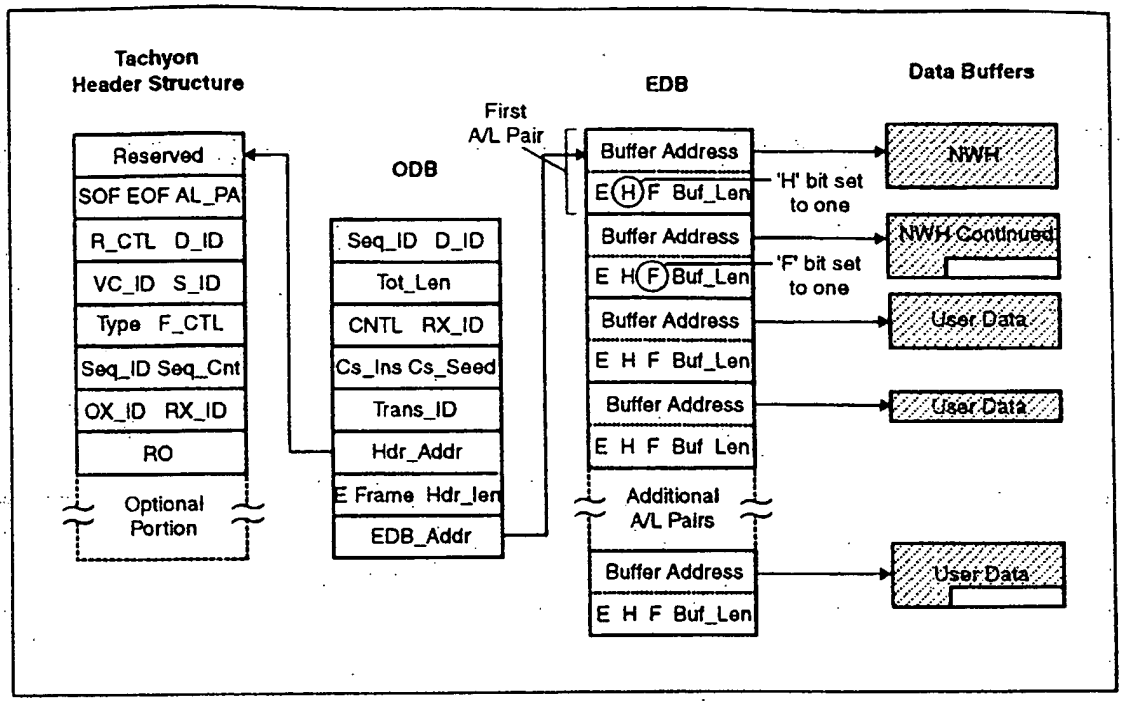

Figure 3.25 Outbound Checksum Requirements

Outbound Header Checksum Requirements

To use header checksum hardware assists, the host sets the Initial Relative Offset (RO) field in the Tachyon Header Structure to point to the first user data buffer. The address of the first A/L pair must point to the beginning of the network header (NWH). If all of the NWH does not fit into one data buffer, the host may use the address of a second A/L pair to point to the remaining portion of the NWH. In this case, the host must set the H bit in the first A/L pair. If all of the NWH fits in one data buffer, then the host only uses one A/L pair and does not set H. In this case, the second A/L pair points to the first data location. The host may not use more than two A/L pairs to point to the NWH.

The host must also compute a seed value that is placed in the ODB for this sequence. This seed value is the checksum of any part of the Network Header that needs to be checksummed. When outbound header checksumming is enabled, the Tot_Len field in the ODB must equal the sum of the optional headers and all of the A/L pairs in the EDB, including the A/L pair with the END bit set. The END bit must be set on the last A/L pair.

Tachyon saves the Network Header A/L pair(s) and processes the data to be sent. Tachyon calculates a running checksum value as it sends the data. Tachyon inserts the calculated checksum into the network header, resets the RO field within the Tachyon Header Structure to zero, and sends this network header out as the last frame.

If the entire sequence fits within one frame, the host should not use TCP/UDP checksumming assists. Tachyon sends the sequence correctly, but performance may suffer since the sequence is broken into two frames.

PTI 172531

Outbound Trailer Checksum Requirements

When the host uses trailer checksumming, it must ensure that header frames are separated from data frames to allow the remote node to split the two. To do this, the host uses the F bit in the A/L pair to instruct Tachyon to send the data in the current A/L pair as a frame by itself and to start a new frame for the next A/L pair. When the host sets the F bit to one in the last A/L pair that points to the NWH, the host forces Tachyon to place the entire NWH in the first frame of the sequence and then to start a new frame for the network data. The remote node can then assume that the first frame is the header and that the data is in the second frame and perform a simple split.

The host should clear the CS_INS field in the ODB to zero even though Tachyon ignores this value during a trailer checksum insertion. Tachyon reads down the sequence from host memory, and replaces the last two bytes of the sequence with the trailer checksum. Tachyon recognizes that if the Tot_Len field is an odd value, the last two bytes are on an odd byte boundary. In this case, Tachyon swaps the bytes of the checksum before the insertion to force proper alignment of the result. The seed value supplied by the host in the CS_Seed field of the ODB should not include the overwritten bytes or the Fill bytes.

Since Tachyon sums the entire sequence (except for the FC header), the host must calculate a seed value that subtracts out the checksum for any network headers that are not to be checksummed. The host places this seed value in the ODB CS_Seed field for Tachyon to add into the running sum.

If the host uses the Relative Offset (RO Present bit set in F_CTL), the RO field in the Tachyon Header Structure must be cleared to zero when using trailer checksumming.

3.10.3 Header/Data Splitting

Tachyon supports two different modes for header/data splitting depending upon whether the network header arrives as the first frame or not. The following discussion assumes that the host has established inbound buffers for Tachyon that are aligned on page boundaries.

In Order Reception

When the sequence reception is in order, Tachyon assumes that the first frame contains the NHW. Tachyon places the first frame into its own buffer and starts the second frame on a new page-aligned buffer. Tachyon does not guarantee that the data begins at the second frame. The host determines this by scanning the received header to determine if the split was correct. If the split was not correct, the host must copy the data to properly align it for delivery to the user.

The Tachyon splitting state machines are disabled when the following conditions occur:

1. Tachyon detects that the sequence is neither an IP packet (TYPE field is 5 = IS8802-2 LLC/SNAP) nor of the type specified in the split_type field of the Tachyon Configuration register.
2. Tachyon detects that the sequence is OOO before the split occurs.

Out Of Order Reception

When the sequence is transmitted using the RRO method, as in header checksumming, the NHW arrives out of order with the network data. When the first frame arrives with the Relative Offset (RO) non-zero, Tachyon assumes the header checksum mode is being used and places the arriving data into the page-aligned buffers. When the header arrives, its RO is not in order. Tachyon changes modes to the OOO mode and generates one completion for the data portion and a separate completion for the header. At this point, Tachyon has split and page-aligned the data. The host must realize that the header completion follows the data and must process it accordingly. The host must check the split and offset fields in the outbound completion message for the data portion of the sequence for information about the out of order split.

PTI 172532

3.11 Network Management Features

Tachyon is designed to operate in the Internet-standard MIB-II environment. It is assumed that the objects defined by MIB-II are to be maintained by the host driver software. However, the driver does not have visibility to a small subset of MIB-II objects in the Interface group, and hence must periodically query Tachyon in order to collect these data. This section describes how some of the Tachyon and Frame Manager registers can be used to track MIB-II Interface object data.

Below are the MIB-II objects in the Interface group for which Tachyon registers may be used to collect network management statistics:

<table>
<tr><th>Object</th><th>Register</th><th>Field</th><th>Description</th></tr>
<tr><td>ifSpeed</td><td>Frame Manager Status register</td><td>PAR_ID</td><td>These two bits can be used to determine the link speeds as follow:
<table><tr><th>Bits</th><th>Link Speed</th></tr><tr><td>01</td><td>266 Mbaud</td></tr><tr><td>10</td><td>531 Mbaud</td></tr><tr><td>11</td><td>1062 Mbaud</td></tr></table></td></tr>
<tr><td>ifOperStatus</td><td>Tachyon Status register</td><td></td><td>Chip error and reset status.</td></tr>
<tr><td></td><td>Frame Manager Status register</td><td>link status, int status, port state</td><td>These fields describe various status of the Frame Manager.</td></tr>
<tr><td>ifInErrors</td><td>Frame Manager Link Error Status #2 register</td><td>rxd_EOFa, gen_EOFa, bad_crc, proto_err</td><td>These values can be added together to determine the total number of inbound packets (frames) discarded due to any error.</td></tr>
</table>

Table 3.13 Network Management MIB II Objects

PTI 172533

4. Initialization and Configuration

This section describes how to estimate the host memory requirements for Tachyon and perform the necessary steps to initialize Tachyon.

Also included in this section is a Tachyon configuration example, testing Tachyon functionality procedure, loopback mode information, and login parameters.

4.1 Estimate Memory Requirements

By adding together the host memory requirements and additional dynamic memory requirements, the total host memory requirements can be estimated. Practical memory usage is between 40 Kbytes and 2 Mbytes, depending on the application. The minimum memory configuration is approximately 33 Kbytes. If SCSI is used, but networking is not, then the minimum memory configuration is approximately 17 Kbytes. Minimum memory configurations are not recommended, however, because memory usage for Tachyon changes dynamically depending on the number of active transactions and the number of buffers they require.

4.1.1 Static Memory Requirements

The minimum amount of host memory required for Tachyon can be estimated by adding the following calculations.

Static memory requirement =

((Number of OCQ entries * 32 bytes) +

(Number of HPCQ entries * 32 bytes) +

(Number of IMQ entries * 32 bytes) +

(Number of MFSBQ entries * 32 bytes) +

(Maximum number of active MFSBQ entries * MFS Buffer Length * 8 bytes) +

(Number of SFSBQ entries * 32 bytes) +

(Maximum number of active SFSBQ entries * SFS buffer length * 8 bytes) +

(Number of SEST entries * 32 bytes) +

(Number of index address registers * 32 bytes))

PTI 172534

4.1.2 Dynamic Memory Requirements

Assumptions about dynamic memory requirements.

1. All OCQ and all HPCQ entries can be used at any time
2. Average data transfer lengths include the header
3. The HPCQ is used only for error recovery
4. All buffers are used for Inbound SEST Entries
5. The number of expected EDB A/L pairs must be multiples of four pairs
6. The SDB OOO Reassembly mode is used for the SEST entries

The amount of additional dynamic memory the host must provide at any time can be estimated by adding the following calculations.

Dynamic memory requirement =

((Maximum number (zero-based) of OCQ entries * average data transfer length in bytes) +

(Maximum number (zero-based) of HPCQ entries * average error recovery transfer length in bytes) +

(Maximum number of expected active inbound SEST entries * 64 bytes * buffer length) +

(Maximum number of expected active outbound SEST entries * average data transfer length in bytes) +

(Maximum number of expected active inbound SEST entries) +

((Maximum number (zero-based) of OCQ entries + Maximum number of expected outbound SEST entries) * number of expected EDB A/L pairs * 32 bytes) +

(Number of additional SFSBQ entries * SFS buffer length) +

(Number of additional MFSBQ entries * MFS buffer length))

PTI 172535

4.2 Initialize Tachyon

The following steps are required to initialize Tachyon. These steps are described in further detail in the subsequent sections.

1. Program the Tachyon Configuration register. If parity is used, this should be the first register programmed to ensure that parity functions properly.
2. Build the host-based and SCSI data structures and inform Tachyon where they are located. These include:
   a. IMQ
   b. OCQ
   c. HPCQ
   d. MFSBQ
   e. SFSBQ
   f. SEST
3. Initialize the Frame Manager.

4.3 Program the Tachyon Configuration Register

Before Tachyon can start processing Fibre Channel requests, Tachyon must be configured for the correct operating modes. The host does this by writing to the Tachyon Configuration register.

Tachyon Configuration register identifies:

1. The optional performance features that are supported by the host.
2. The characteristics of the link to the Fibre Channel that Tachyon is connected.
3. The operational functions within Tachyon to disable for the purpose of debugging.

4.4 Build Data Structures

The length of the host-based data structures must be a power of 2 and must be aligned on sizeof (data structure) boundaries. If the data structures are not built with proper values, or on proper boundaries, or with sufficient space, Tachyon does not function properly.

The host must write two parameters for each of the queues and the SEST to the proper Tachyon queue register. These two parameters are (1) the base address of the queue and (2) the length of the queue, where length is the (zero-basesd) number of queue entries. When the SEST is defined, each entry should be cleared to zero so Tachyon can determine when an invalid entry is being accessed.

Tachyon requires that one entry in each of the queues is always empty. Therefore, all queues must have minimum length of two entries, except for the IMQ. The minimum number of entries in the IMQ is four. The suggested size of the IMQ is at least twice the size of the other queues because Tachyon uses this queue to send completion messages to the host.

If the SFSBQ or MFSBQ length is two entries, memory can be allocated for the first entry. As soon as Tachyon reads the entry, it sends a sfs or mfs_buf_warn completion message to the host. Then, memory is allocated for the next queue entry and the producer index is updated. Tachyon prefetches the pointers. Tachyon continues to use all the buffers in the first queue entry. As soon as Tachyon starts to use the last queue entry's buffers, it issues a buffer warning completion message. At this time, the host should either process the buffers of the queue's first entry or allocate additional memory for the buffers of the queue's first entry, and then write the producer index register. This process is repeated each time Tachyon sends a buffer warning completion message. If this process is used, the host must be able to supply buffers to Tachyon before ED_TOV expires.

If the SFSBQ and MFSBQ length is more than two entries, memory usage can be reduced by not allocating memory for unused entries until the first queue entry is filled. Then the first queue entry can be emptied and its memory allocated for the next unused queue entry. Tachyon's indexes can be updated until the queues empty again, and the procedure repeated. The host must be able to supply buffers to Tachyon before ED_TOV expires.

Tachyon does not generate a buffer warning completion message for an empty condition until the host has written to the inbound buffer queue producer index at least once. This prevents an interrupt from being generated after a reset and before initialization is complete.

For Tachyon to operate properly after reset, the host must allocate buffers to receive data from Tachyon (for SFS, MFS, and SCSI frames). For each inbound buffer queue (SFSBQ and MFSBQ) entry, the host provides multiples of eight empty buffers and writes the eight buffer addresses in the entry. The host then updates the producer index register with the index of the next empty entry (after the last prepared entry).

The two outbound command queues (OCQ and HPCQ) have producer indicies to indicate to Tachyon the location in the queue of a sequence that has been posted for transmission.

The IMQ is assumed to be fully initialized when the host writes the IMQ Producer Index Address register. Once the host writes to this register, Tachyon owns the queue entries defined by these values and can start sending completion messages to the host to process. The host writable register for this queue is the consumer index. The host uses the consumer index to return ownership of processed IMQ entries to Tachyon.

Also, the host must allocate three 4-byte words of host memory to Tachyon for writing host copies of the current values of Tachyon's producer and consumer indexes. The memory addresses of these locations are written to Tachyon's three index address registers (OCQ Consumer Index, HPCQ Consumer Index, and IMQ Producer Index address registers). These registers should be written only after the queues' base, length, and buffer length registers have been configured, because writing a queue's index address register is Tachyon's signal to process the queue's base, length, and index registers.

PTI 172537

The producer index registers for the SFSBQ and the MFSBQ must be the last registers written. As soon as these registers are written, Tachyon processes the first queue entry. At initialization, Tachyon resets this producer index value to zero. The inbound buffer queues (SFSBQ and MFSBQ) do not have host-based indexes because the completion messages indicate which buffers have been consumed. However, if the host needs these values, they are available in Tachyon's readable SFSBQ Consumer Index and MFS Consumer Index registers.

Tachyon uses buffer length registers to indicate the length, in bytes, of SFS, MFS, and SCSI receive buffers. These buffer length registers determine the size of the buffer that gets assigned to each buffer entry. For the SFS and MFS buffer length registers, the buffer size must be a power of two.

The buffer length for SCSI depends upon the SDB Buffer Address Mode bit in the Tachyon Configuration register. If the SDB Buffer Address Mode is cleared, then the buffer size must be a power of two. If this bit is set, the buffer size must be a multiple of four bytes.

Values for the queues and buffer registers must be chosen carefully because most values can be set only at initialization. The only registers that can be modified after initialization are the host-owned OCQ, HPCQ, SFSBQ, MFSBQ Producer Index registers and the IMQ Consumer Index register.

4.5 Initialize the Frame Manager

The Frame Manager responds to a reset by initializing all of its registers to their default values. The Frame Manager state machines are initialized to an offline state and the input data stream is directly connected to the output data stream to act as an electrical bypass of this node.

To start Tachyon operation, the host must configure the Frame Manager. To do this, the host performs the following:

1. Initializes Tachyon registers.
2. Writes the following fields in the Frame Manager Configuration register.
   a. The host sets the AL_PA value to any pre-defined value that it chooses. Valid AL_PA values are defined in the FC_AL Specification. If the host does not care about the value of its loop address or if the host has set the N_Port bit, the host should clear this field to zero.
   b. The BB_Credit field is cleared to zero after reset and should be set to one by the host for N_Port initialization. Tachyon ignores this value for loop or loop initialization and uses a value of zero.
   c. If a loopback test is needed, refer to "4.8 Loopback Mode" on page 108.
   d. If Tachyon is to participate on a connected arbitrated loop, one of the arbitrated loop address options must be selected to indicate whether the Frame Manager should try to get a pre-defined address or just take what it can get during loop initialization. If no pre-defined address is needed, the Soft Address (sa) bit should be set. If the host wants to be non-participating on the loop, all four of the address bits (Soft Address, Hard Address, Previously Acquired, and Fabric Acquired) must be cleared to zero.
   e. If the N_Port bit is set, the arbitrated loop parameters are not used. If the host knows that it is connected to a fabric, it can set the N_Port bit to keep Tachyon from attempting arbitrated loop initialization. This prevents performance degradation. Degradation occurs by constantly arbitrating and opening the loop to transfer data when only two ports are present.
3. Writes the unique address associated with this node to the Frame Manager World Wide Name (Hi and Lo) registers. This value is only used during loop initialization. If the host has set the N_Port bit in the Frame Manager Configuration register, it is not necessary to write these registers.
4. Optionally writes to the Frame Manager RT_TOV/AL_TIME & ED_TOV register. If the default values in this register are not acceptable, they must be initialized to the desired values.

   The Frame Manager uses the Frame Manager RT_TOV/AL_TIME & ED_TOV register for the Receiver Transmitter Time Out Value (RT_TOV) or Arbitrated Loop Time Out Value (AL_TIME). RT_TOV is used for Point-to-Point Initialization and AL_TIME is used for Loop Initialization. RT_TOV should be programmed to 100 milliseconds (ms) and AL_TIME should be programmed to 15 ms.

   If the host does not know its topology before initialization, the Frame Manager RT_TOV/AL_TIME & ED_TOV register must be programmed for 15 ms for Loop Initialization. After the host determines that no loop exists, the RT_TOV/AL_TIME value should be changed to 100 ms.

   Follow these steps to use the Frame Manager RT_TOV/AL_TIME & ED_TOV register:

   a. The host initializes the RT_TOV/AL_TIME register to 15 ms.
   b. The host instructs the Frame Manager to go online. Instruct the Frame Manager to go online by writing the Initialize Command in the Frame Manager Control register (Set bits 2..0 to 110).
   c. The host begins a timer for at least 40 ms.
   d. If the host receives a frame_mgr_interrupt completion message before the 40 ms timer expires and the Frame Manager Status register indicates Link Up, ensure that the Frame Manager has initialized as a loop port. The Frame Manager Status register Loop State Machine bits should not indicate OLD PORT.

PTI 172539

e. If the host receives a frame_mgr_interrupt completion message before the 30 ms timer expires and the Frame Manager Status register indicates Link Up, and the Frame Manager has initialized as a N_Port, re-initialize Frame Manager as described in step f. The Loop State Machine bits in the Frame Manager Status register should indicate OLD PORT.

f. If the 30 ms timer expires, the host should read the Loop State Machine bits in the Frame Manager Status register. If the the Loop State Machine bits indicate OLD PORT state, take Frame Manager offline and change the RT_TOV/AL_TIME value to 100 ms. Also, set the N_Port bit in the Frame Manager Configuration register and take the Frame Manager online again.

g. If the 30 ms timer expires, the host should read the Loop State Machine bits in the Frame Manager Status register. If the Loop State Machine bits indicate INITIALIZE, then no cable is attached and steps 'c' through 'g' should be followed again.

5. Lastly, the host must write to the Frame Manager Control register with the Initialize Command set. This forces Tachyon to start the proper (loop or N_Port) initialization.

When Tachyon completes initialization, it generates a frame_mgr_interrupt completion message. At that point, the host can determine the status of the initialization by reading the Frame Manager Status register. If loop initialization succeeds, the lfsm and pfsm fields are set to MONITOR and OFFLINE, respectively. If N_Port initialization succeeds, the lfsm and pfsm fields are set to OLD_PORT and ACTIVE, respectively. Also, in both cases, the Link Up (lup) bit is set.

If a Tachyon initializes as an L_Port, and another Tachyon initializes as an N_Port, the L_Port Tachyon recognizes the OLSs and sets the NOS/OLS Received bit in its Frame Manager Status register. Setting the NOS/OLS Received bit indicates that a NOS and/or OLS was received after a loop circuit was established. The L_Port Tachyon then generates a frame_mgr_interrupt completion message to its host. The L_Port host can then recognize this condition and, if no other ports exist on the loop, it should re-initialize Frame Manager as an N_Port.

4.5.1 Determine the Link State of Tachyon

Because both the Link Up (lup) and Link Down (ldwn) bits in the Frame Manager Status register can be set at the same time, the recommended method for determining the link's state is to test the Loop State Machine (lfsm) and Port State Machine (pfsm) bits in the Frame Manager Status register as follows.

Loop

If bit 7 of the Frame Manager Status register is cleared to zero, then the link is up. If bit 7 is set to one or the four low order bits (3..0) are not set to 0xF, then the link is down.

N_Port, Point-to-Point

If the four low order bits (bits 3..0) are set to 0xF, then Point-to-Point is active. If the low order bits are not set to 0xF, then the link is down.

4.6 Tachyon Configuration Example

The following table shows a sample mapping of Tachyon registers in host memory and configuration values that could be written to the registers for operation. It also shows how memory is reserved for Tachyon's use. The example does not specify either the maximum or minimum amount of memory that Tachyon could use. Driver writers must work within their memory constraints and adjust the dynamic memory equation to fit their application's needs and requirements.

Assumptions of the Tachyon configuration example:

1. Tachyon resides in memory mapped I/O space
2. Locations 0x0000 through 0x01EF are memory mapped I/O space
3. Real memory begins at address 0x0200
4. The buffer length for the SFS, MFS and SDB is 512 bytes

Location	Value	Description	Comment
0x0000	0x0200	OCQ Base register	
0x0004	0x0007	OCQ Length register, containing the number (zero-based) of OCQ entries	Example shows 8 entries
0x0008	0x0000	OCQ Producer Index register	Initialized to 0x0000
0x000C	0x0700	OCQ Consumer Index Address register	
0x0040	0x0300	HPCQ Base register	
0x0044	0x0007	HPCQ Length register, containing the number (zero-based) of HPCQ entries	Example shows 8 entries
0x0048	0x0000	HPCQ Producer Index register	Initialized to 0x0000
0x004C	0x0720	HPCQ Consumer Index Address register	
0x0080	0x0400	IMQ Base register	
0x0084	0x000F	IMQ Length register, containing the number (zero-based) of IMQ entries	Example shows 16 entries
0x0088	0x0000	IMQ Consumer Index register	Initialized to 0x0000
0x008C	0x0740	IMQ Producer Index Address register	
0x00C0	0x0600	MFSBQ Base register	
0x00C4	0x0007	MFSBQ Length register, containing the number (zero-based) of MFSBQ entries	Example shows 8 entries
0x00C8	0x0004	MFSBQ Producer Index register	Four active entries: 0, 1, 2, and 3
0x00CC	0x0000	MFSBQ Consumer Index register	Initialize to 0x0000
0x00D0	0x01FF	MFS Buffer Length register	Default value is 0x01FF

Table 4.1 Tachyon Configuration Example - Memory Mapped I/O Space

PTI 172541

Location	Value	Description	Comment
0x0100	0x0800	SFSBQ Base register	Aligned on a 512 byte boundary
0x0104	0x000F	SFSBQ Length register, containing the number (zero-based) of SFSBQ entries	Example shows 16 entries
0x0108	0x0004	SFSBQ Producer Index register	Four active entries: 0, 1, 2, and 3
0x010C	0x0000	SFSBQ Consumer Index register	Initialized to 0x0000
0x0110	0x01FF	SFS Buffer Length register	Default value is 0x001F
0x0140	0x0A00	SEST Base register	
0x0144	0x000F	SEST Length register, containing the number (zero-based) of SEST entries	Example shows 16 entries
0x0148	0x01FF	SCSI Buffer Length register	Default value is 0x0003
0x0184 to 0x019C		Tachyon registers	All buffers are 512 bytes, normal mode of operation for the SEST
0x01C0 to 0x01EC		Frame Manager registers	All buffers are 512 bytes, normal mode of operation for the SEST

Table 4.1 Tachyon Configuration Example - Memory Mapped I/O Space (Continued)

PTI 172542

Location	Value	Description	Comment
0x0200 to 0x021F		OCQ entry (ODB) 0	
0x0220 to 0x02FF		OCQ entries (ODBs) 1 through 7	
0x0300 to 0x03FF		HPCQ entries (HPDBs) 0 through 7	
0x0400 to 0x05FF		IMQ entries (completion messages) 0 through 15	
0x0600 to 0x06FF		MFSBQ entries 0 through 7	
0x0700		OCQ Consumer Index register	Tachyon initializes this value to 0x0000
0x0720		HPCQ Consumer Index register	Tachyon initializes this value to 0x0000
0x0740		IMQ Producer Index register	Tachyon initializes this value to 0x0000
0x0800	0x0C00	SFSBQ entry 0, buffer 0	
0x0804	0x0E00	SFSBQ entry 0, buffer 1	
0x0808	0x1400	SFSBQ entry 0, buffer 2	
0x080C	0x1200	SFSBQ entry 0, buffer 3	
0x0810	0x1000	SFSBQ entry 0, buffer 4	
0x0814	0x1600	SFSBQ entry 0, buffer 5	
0x0818	0x1A00	SFSBQ entry 0, buffer 6	
0x081C	0x1800	SFSBQ entry 0, buffer 7	
0x0820 to 0x09FF		SFSBQ entries 1 through 15, each having 8 buffers	
0x0A00 to 0x0BFF	0x0000	SEST	Initialize to 0x0000
0x0C00 to 0x0DFF		SFSBQ entry 0, buffer 0	Base address must be 32-byte aligned
0x0E00 to 0x4BFF		Buffers for active SFSBQ entries	Must be 32-byte aligned
0x4C00 to 0x8C00		Buffers for active MFSBQ entries	Must be 32-byte aligned

Table 4.2 Tachyon Configuration Example - Real Memory

PTI 172543

4.7 Testing Tachyon's Functionality Incrementally

This section describes how the functionality of a hardware design may be checked in incremental blocks after Tachyon initialization.

1. Slave Reads and Writes:
   a. Write a software reset to the Tachyon Control register.
   b. Read the Tachyon Configuration register. It should contain 0x00000002.
   c. Write any value to the Tachyon Configuration register.
   d. Read back and verify the value written to the Tachyon Configuration register (Note: Must keep parity bits set properly).
   e. Repeat steps "c." and "d." (above) for the Frame Manager Configuration register and the Frame Manager World Wide Name (Hi/Lo) register.
   f. Write some value other than 0x00000002 to the Tachyon Configuration register.
   g. Write a software reset to the Tachyon Control register.
   h. Read back the Tachyon Configuration register. It should contain 0x00000002, which is the Parity Even default value.
   i. If it does not, check for hardware problems.
2. Tachyon DMA Reads/Writes and Interrupts:
   a. Write a software reset to the Tachyon Control register.
   b. Set up all the queues for normal operation, and write 0x0001 to the SFSBQ producer index. This causes Tachyon to perform a DMA Read from the SFSBQ, a DMA Write to the IMQ, an update to the IMQ Producer Index, and to generate an interrupt.
   c. Verify that an interrupt occurred.
   d. Verify that the first IMQ entry contains 0x0107, which is the sfs_buf_warning completion message.
   e. Verify that the IMQ producer index in host memory contains 0x0001.
   f. If this does not work,
      1) Verify the chip was reset.
      2) Verify the IMQ base address is on the correct boundary.
      3) Verify the IMQ length. It should be equal to or greater than 0x0004.
      4) Verify the IMQ producer index address is on the correct boundary and is correctly written.
      5) Verify the SFSBQ base address is on the correct boundary.
      6) Verify the SFSBQ length.
      7) Verify the SFS producer index contains 0x0001.
      8) If all of the above are correct, check for hardware problems.
3. Tachyon and Internal/External Loopback:
   a. Write a software reset to the Tachyon Control register.
   b. Set up all the queues for normal operation.
   c. Follow the selftest procedure in the sample C code. Refer to "Tachyon C Code" on page xxii.
   d. If this does not work,
      1) Verify Tachyon was initialized properly and the selftest C code correlates. Refer to "4." (below).
      2) Look for hardware problems with DMA reads and writes.
4. Status Request
   a. Perform a status request on Tachyon by writing 0x00000001 to the Tachyon Control register.
   b. Verify that the last frame sent out via internal loopback is contained in both the inbound and outbound structures.

The procedure described above in 2. also works using the MFSBQ in step 2. b. PTI 172544

4.8 Loopback Mode

To verify correct operation of Tachyon, the host may send loopback sequences. Loopback causes Tachyon to send data to itself, either through an internal loopback path or through the PLM. In internal loopback mode, Tachyon connects the PLM input data stream to the PLM output data stream.

The host uses the Enable Internal Loopback (il) and Enable External Loopback (el) bits in the Frame Manager Configuration register to select the type of loopback. To use loopback mode, Tachyon must be initialized completely and one of the loopback bits must be set. The loopback bits should be set only during initialization testing. If they are set during normal operation, Tachyon will take the loop down.

If the N_Port bit in the Frame Manager Configuration register is not set, then Tachyon initializes on a loop and the AL_PA it acquires for itself must be used or Tachyon enters an indeterminate state. If Tachyon initializes on a loop while in loopback mode, any data sent to an AL_PA other than the AL_PA acquired by Tachyon causes a bad AL_PA error.

Because Tachyon does not support bi-directional connections, only Class 2 and Class 3 sequences may be used in loopback mode.

4.8.1 Loopback Initialization Steps

To place Tachyon into loopback mode and send frames, perform the following steps:

N_Port Loopback Initialization Steps

1. Initialize all Tachyon registers.
2. In the Frame Manager Configuration register, set the N_Port bit and either the Enable Internal Loopback or the Enable External Loopback bit. Also set the BB_Credit to 0x01.
3. Write the Initialize Command (0x06) to the Frame Manager Control register.
4. Tachyon interrupts the host with a Frame Manager interrupt to indicate the link is up.
5. The host sends frames.
6. After all loopback frames have been sent, received, and processed, take Tachyon offline by writing the Offline Command (0x05) to the Frame Manager Control register. Before issuing the Offline command, the N_Port bit in the Frame Manager Configuration register should be cleared to zero.
7. Tachyon interrupts the host to indicate the link went down.
8. Both the Enable Internal Loopback or the Enable External Loopback bits in the Frame Manager Configuration register should be cleared to zero and the BB_Credit set to 0x01 or to the negotiated login value.
9. The port can now be configured to initialize as an N_Port. If the host wants to initialize on the loop, BB_Credit should be cleared to zero, the N_Port bit should be cleared to zero, and one of the loop address bits (aq, ha, sa, fa) should be set. If the fa bit is set, then the aq bit must also be set. If an address bit other than the sa bit is set, then the AL_PA field in the Frame Manager Configuration register must be valid.

The FC-AL Loopback Initialization Steps

1. Initialize all Tachyon registers.
2. In the Frame Manager Configuration register, clear the N_Port bit to zero, set either the il bit or the el bit to one, and set the sa bit.
3. Write the Initialize Command (0x06) to the Frame Manager Control register.
4. Tachyon interrupts the host and indicates the link is up.
5. After all loopback frames have been sent, received, and processed, take Tachyon offline by writing the Offline Command (0x05) to the Frame Manager Control register.
6. Tachyon interrupts the host to indicate the link went down.
7. Both the Enable Internal Loopback or the Enable External Loopback bits in the Frame Manager Configuration register should be cleared to zero and the BB_Credit field reset to default values.
8. The port can now be configured to initialize on the loop or as an N_Port.

PTI 172545

4.9 Login Parameters

Before the host can do any real communication on Fibre Channel, it must login with the Fabric and every other node with which it communicates. Login parameters are necessary for Tachyon to operate with the Fabric, other systems, and other devices.

The following tables indicate the required parameters that the host must use in its Fabric or N_Port Login sequences. Optional and recommended values are indicated in parentheses. All other values are required.

4.9.1 Fabric Login Parameters

	Fabric Login Parameters Sent	Fabric Login Parameters Expected
FC-PH Version	High = Rev 4.3 = 0x09 (optional) Low = Rev 4.2 = 0x08	High = Rev 4.3 = 0x09 (optional) Low = Rev 4.2 = 0x08
BB_Credit	2, if point-to-point/fabric attached 0, if loop	1 or more, if point-to-point, fabric attached 0 or more, if loop
BB Receive Data Field Size	Min: 128 Max: 2048 (recommended)	
Class 1 Service Options	Stacked Connection Requests (optional); Intermix	Stacked Connections (optional); Intermix;
Class 2, 3 Service Options	Sequential Delivery (optional, but recommended)	

Table 4.3 Fabric Login Parameters

4.9.2 N_Port Login Parameters

	N_Port Login Parameters Sent	Destination N_Port Login Parameters Expected
FC-PH Version	High = Rev 4.3 = 0x09 (optional) Low = Rev 4.2 = 0x08	High = Rev 4.3 = 0x09 (optional) Low = Rev 4.2 = 0x08
BB_Credit	2, if point-to-point 0, if loop	1 or more, if point-to-point 0 or more, if loop
Common Features	Random Relative Offset (RO), RO by Information Category; 0x00FF	
Total Concurrent Sequences	255	1 or more
Class 1 Service Options	Intermix	Intermix for point-to-point only

Table 4.4 N_Port Login Parameters

PTI 172546

	N_Port Login Parameters Sent	Destination N_Port Login Parameters Expected
Class 1 Initiator Control	X_ID reassignment supported, ACK_0, and ACK_N capable (all optional)	None
Class 1 Recipient Control	ACK_0 capable (optional); ACK_N incapable; X_ID Interlock not required; Discard Error Policy Only; One Information Category per Sequence	Discard Error Policy
Class 1 Receive Data Field Size	2048	
Class 1 Concurrent Sequences	1	1 or more
Class 1 Credit	3 (maximum)	1 or more
Class 1 Open Sequences per Exchange	1	1 or more
Class 2 Initiator Control	X_ID reassignment supported; ACK_0; ACK_N capable (all optional)	
Class 2 Recipient Control	ACK_0 capable (optional); ACK_N incapable;, X_ID Interlock not required; Discard Error Policy Only; One Information Category per Sequence	Discard Error Policy
Class 2 Receive Data Field Size	2048	
Class 2 Concurrent Sequences	1 (up to 255 if SCSI transactions are supported)	1 or more
Class 2 Credit	4	1 or more
Class 2 Open Sequences per Exchange	1	1 or more
Class 3 Initiator Control	None	
Class 3 Recipient Control	Discard Error Policy Only, One Information Category per Sequence	Discard Error Policy
Class 3 Receive Data Field Size	2048	
Class 3 Concurrent Sequences	1 (up to 255 if SCSI transactions are supported)	1 or more
Class 3 Open Sequences per Exchange	1	1 or more

Table 4.4 N_Port Login Parameters (Continued)

PTI 172547

5. Registers

5.1 Register Overview

Register Name	Page Number	Read/ Write	Address	Reset Value
OCQ Base register	114	W	0x0000	0x00000000
OCQ Length register	114	W	0x0004	0x00000000
OCQ Producer Index register	115	W	0x0008	0x00000000
OCQ Consumer Index Address register (3)	115	W	0x000C	0x00000000
Host's Copy of Tachyon's OCQ Consumer Index	116	R	host memory	
HPCQ Base register	117	W	0x0040	0x00000000
HPCQ Length register	117	W	0x0044	0x00000000
HPCQ Producer Index register	118	W	0x0048	0x00000000
HPCQ Consumer Index Address register (3)	118	W	0x004C	0x00000000
Host's Copy of Tachyon's HPCQ Consumer Index	119	R	host memory	
IMQ Base register	120	W	0x0080	0x00000000
IMQ Length register	120	W	0x0084	0x00000000
IMQ Consumer Index register	121	W	0x0088	0x00000000
IMQ Producer Index Address register (3)	122	W	0x008C	0x00000000
Host's Copy of Tachyon's IMQ Producer Index	122	R	host memory	
MFSBQ Base register	123	W	0x00C0	0x00000000
MFSBQ Length register	123	W	0x00C4	0x00000000
MFSBQ Producer Index register	124	W	0x00C8	0x00000000
MFSBQ Consumer Index register	124	R	0x00CC	0x00000000
MFS Buffer Length register	125	W	0x00D0	0x00000000
SFSBQ Base register	126	W	0x0100	0x00000000
SFSBQ Length register	126	W	0x0104	0x00000000
SFSBQ Producer Index register	127	W	0x0108	0x00000000
SFSBQ Consumer Index register	127	R	0x010C	0x00000000
SFS Buffer Length register	128	W	0x0110	0x00000000
SEST Base register	129	W	0x0140	0x00000000

Table 5.1 Tachyon Memory Map

PTI 172548

Register Name	Page Number	Read/ Write	Address	Reset Value
SEST Length register	129	W	0x0144	0x00000000
SCSI Buffer Length register	130	W	0x0148	0x00000000
Tachyon Configuration register	131	R/W	0x0184	0x00000002
Tachyon Control register	133	W	0x0188	0x00000000
Tachyon Status register	135	R	0x018C	HW dependent
Tachyon Flush SEST Cache Entry register	137	R/W	0x0190	0x00000000
Tachyon EE_Credit Zero Timer register	138	R	0x0194	0x00000000 (2)
Tachyon BB_Credit Zero Timer register	138	R	0x0198	0x00000000 (2)
Tachyon Receive Frame Error Counter register	139	R	0x019C	0x00000000
Frame Manager Configuration register	140	R/W	0x01C0	0x00000000
Frame Manager Control register	142	W	0x01C4	0x00000000
Frame Manager Status register	144	R	0x01C8	0x00000000
Frame Manager RT_TOV/AL_TIME & ED_TOV register	147	W	0x01CC	0x001001F5
Frame Manager Link Error Counters #1 register	148	R	0x01D0	0x00000000
Frame Manager Link Error Counters #2 register	149	R	0x01D4	0x00000000
Frame Manager World Wide Name Hi register	150	R/W	0x01E0	0x00000000
Frame Manager World Wide Name Lo register	150	R/W	0x01E4	0x00000000
Frame Manager Received AL_PA register	151	R	0x01E8	0x00000000
Frame Manager Primitive register	151	W	0x01EC	0x00000000

Table 5.1 Tachyon Memory Map (Continued)

Tachyon Memory Map Notes

1. At reset, Tachyon clears all registers to zero, unless otherwise noted. Reset is defined as power-on, a hardware reset, or a software reset.
2. While Tachyon clears registers to zero at reset, certain registers (such as EE_Credit Zero Timer and BB_Credit Zero Timer registers) begin counting up immediately following a reset. As a result, if the host reads a register immediately after a reset, the value may be non-zero.
3. If the host reads an invalid or Write only (W) register, the transaction completes normally, and Tachyon returns 0x00000000.
4. Tachyon ignores bits 31..9 of the address during register accesses.
5. If the host writes to Read only (R) registers, the transaction completes but Tachyon does not write the data to the register.
6. Some Tachyon registers contain bits that indicate a certain parameter is either active or inactive. For most cases, 1=active and 0=inactive, but there are exceptions. Refer to the detailed description of each register bit for specifics.

PTI 172549

5.2 Length Register Values

The following table summarizes the programmed values for length registers.

Length Register Name	Length (n)	Programmed Length Field Value (n-1)	Minimum Programmed Value	Maximum Programmed Value
OCQ Length register	The Number of OCQ Queue Entries (Must be a power of 2)	n-1	1	255
HPCQ Length register	The Number of HPCQ Queue Entries (Must be a power of 2)	n-1	1	255
IMQ Length register	The Number of IMQ Queue Entries (Must be a power of 2)	n-1	3	255
MFSBQ Length register	The Number of MFSBQ Queue Entries (Must be a power of 2)	n-1	1	255
MFS Buffer Length register	The Number of Bytes of the MFS Receive Buffers (Must be a power of 2)	n-1	511	65,535
SFSBQ Length register	The Number of SFSBQ Queue Entries (Must be a power of 2)	n-1	1	255
SFS Buffer Length register	The Number of Bytes of the SFS Receive Buffers (Must be a power of 2)	n-1	511	4095
SEST Length register	The Number of SEST Entries (Must be a power of 2)	n-1	0	16,383
SCSI Buffer Length register - OOO Reassembly	The Number of Bytes of the SCSI Receive Buffers (Must be a power of 2)	n-1	511	65,535
SCSI Buffer Length register - In Order Reassembly	The Number of Bytes of the SCSI Receive Buffers (Must be a multiple of 4 bytes)	n-1	3	65,535

Table 5.2 Length Registers Information

PTI 172550

5.3 OCQ Registers

5.3.1 OCQ Base Register

Register Address: 0x0000
Reset Value: 0x00000000

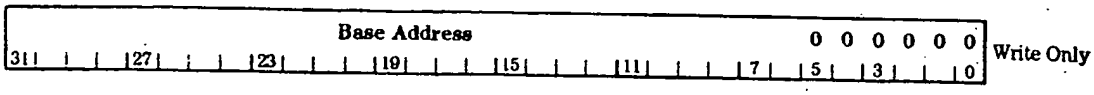

Bit(s)	Field Label	Field Name	Description
31..0	Base Address	Base Address	The physical address in host memory of the start of the OCQ.

Table 5.3 OCQ Base Register

OCQ Base Register Notes

1. The OCQ must be aligned on a sizeof(queue) boundary. For example, if the sizeof (OCQ) equals 512 bytes, then bits 8..0 are cleared to zero. The minimum size of the OCQ is 64 bytes (bits 5..0 are cleared to zero).
2. When the host writes the OCQ Base register, Tachyon clears the corresponding consumer index to zero. Therefore, the host must not write to the OCQ Base register after initialization.

5.3.2 OCQ Length Register

Register Address: 0x0004
Reset Value: 0x00000000

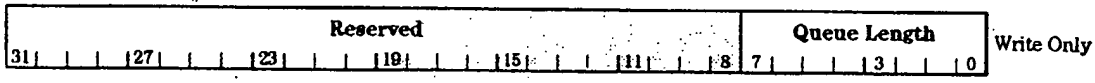

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Length	Queue Length	The number (zero-based) of OCQ entries.

Table 5.4 OCQ Length Register

OCQ Length Register Notes

1. The Queue Length field must be a power of 2. The minimum queue length is 2 entries, corresponding to a programmed value of 1. The maximum queue length is 256 entries, corresponding to a programmed value of 255.
2. When the host writes to the OCQ Length register, Tachyon clears the corresponding consumer index to zero. Therefore, the host must not write to the OCQ Length register after initialization.

PTI 172551

5.3.3 OCQ Producer Index Register

Register Address: 0x0008
Reset Value: 0x00000000

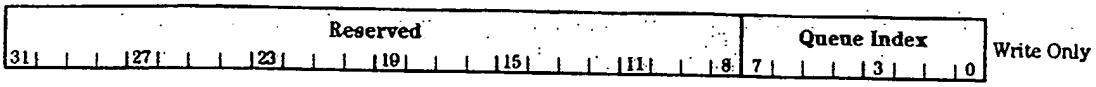

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The producer index value of the next empty OCQ entry.

Table 5.5 OCQ Producer Index Register

OCQ Producer Index Register Note

1. The host uses the OCQ Producer Index register to inform Tachyon that there are new commands to process in the OCQ. After the host uses an OCQ entry, the host writes the index of the next empty OCQ entry to the OCQ Producer Index register. When the host written producer index is different than the internally maintained consumer index, Tachyon processes the new commands.

5.3.4 OCQ Consumer Index Address Register

Register Address: 0x000C
Reset Value: 0x00000000

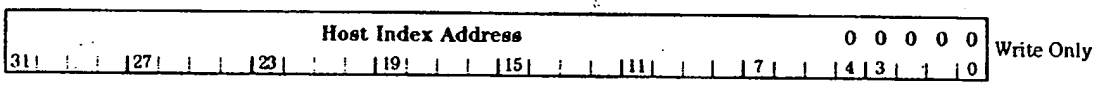

Bit(s)	Field Label	Field Name	Description
31..0	Host Index Address	Host Index Address	The address of the host's copy of the OCQ consumer index.

Table 5.6 OCQ Consumer Index Address Register

OCQ Consumer Index Address Register Notes

1. The OCQ Consumer Index Address register contains the host memory address where Tachyon maintains its consumer index.
2. This index resides in host memory to allow the host fast access to the consumer index.
3. The host index address must be aligned on a 32-byte boundary (bits 4..0 are cleared to zero).

PTI 172552

5.3.5 Host's Copy of the OCQ Consumer Index Register

Register Address: Determined by the host
Reset Value: Determined by the host

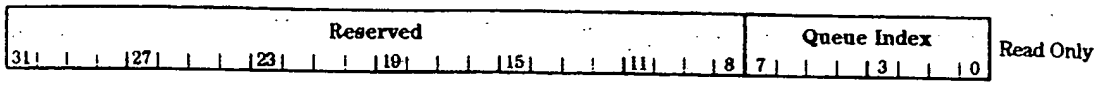

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Value is undefined.
7..0	Queue Index	Queue Index	The consumer index value of the OCQ.

Table 5.7 Host's Copy of the OCQ Consumer Index Register

Host's Copy of the OCQ Consumer Index Register Notes

1. This memory location described by the OCQ Consumer Index Address register contains the host's copy of Tachyon's OCQ consumer index.
2. Tachyon updates this index when it reads an OCQ entry.

PTI 172553

5.4 HPCQ Registers

5.4.1 HPCQ Base Register

Register Address: 0x00040
Reset Value: 0x00000000

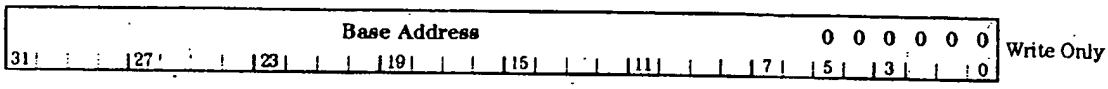

Bit(s)	Field Label	Field Name	Description
31..0	Base Address	Base Address	The physical address in host memory of the start of the HPCQ.

Table 5.8 HPCQ Base Register

HPCQ Base Register Notes

1. The HPCQ must be aligned on a sizeof(queue) boundary. The minimum size of the HPCQ is 64 bytes (bits 5..0 are cleared to zero).
2. When the host writes to the HPCQ Base register, Tachyon clears the corresponding consumer index to zero; therefore, the host must not write to this register after initialization.

5.4.2 HPCQ Length Register

Register Address: 0x0044
Reset Value: 0x00000000

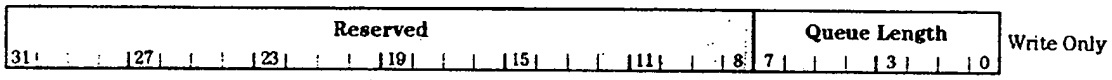

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Length	Queue Length	The number (zero-based) of HPCQ entries.

Table 5.9 HPCQ Length Register

HPCQ Length Register Notes

1. The Queue Length field must be a power of 2. The minimum queue length is 2 entries, corresponding to a programmed value of 1. The maximum queue length is 256 entries, corresponding to a programmed value of 255.
2. When the host writes to the HPCQ Length register, Tachyon clears the corresponding consumer index to zero; therefore, the host must not write to this register after initialization.

PTI 172554

5.4.3 HPCQ Producer Index Register

Register Address: 0x0048
Reset Value: 0x00000000

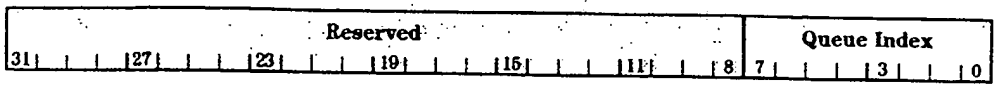

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The producer index value of the next empty HPCQ entry.

Table 5.10 HPCQ Producer Index Register

HPCQ Producer Index Register Note

1. The host uses the HPCQ Producer Index register to indicate to Tachyon that there are new commands to process in the circular queue. After the host uses an HPCQ entry, it writes the index of the next empty entry in the queue to this register. When Tachyon notices that the host written producer index is different from the internally maintained consumer index, Tachyon processes the new commands.

5.4.4 HPCQ Consumer Index Address Register

Register Address: 0x004C
Reset Value: 0x00000000

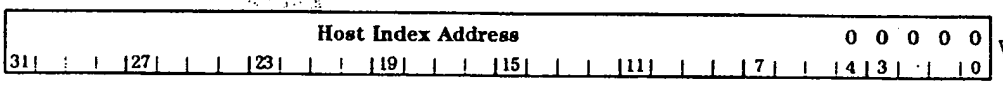

Bit(s)	Field Label	Field Name	Description
31..0	Host Index Address	Host Index Address	The address of the host's copy of the HPCQ consumer index.

Table 5.11 HPCQ Consumer Index Address Register

HPCQ Consumer Index Address Register Notes

1. The HPCQ Consumer Index Address register contains the host memory address where Tachyon maintains its consumer index.
2. This index resides in host memory to allow the host fast access to the consumer index.
3. The host index address must be aligned on a 32-byte boundary (bits 4..0 are cleared to zero).

PTI 172555

5.4.5 Host's Copy of the HPCQ Consumer Index Register

Register Address: Determined by the host
Reset Value: Determined by the host

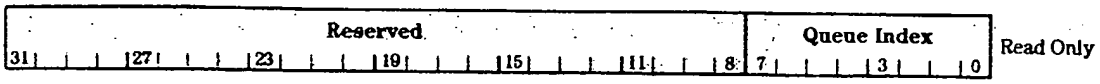

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Value is undefined.
7..0	Queue Index	Queue Index	The consumer index value of the HPCQ.

Table 5.12 Host's Copy of the HPCQ Consumer Index Register

Host's Copy of the HPCQ Consumer Index Register Notes

1. This memory location described by the HPCQ Consumer Index Address register contains the host's copy of Tachyon's HPCQ consumer index.
2. Tachyon updates this index when it reads an HPCQ entry.

PTI 172556

5.5 IMQ Registers

5.5.1 IMQ Base Register

Register Address: 0x0080
Reset Value: 0x00000000

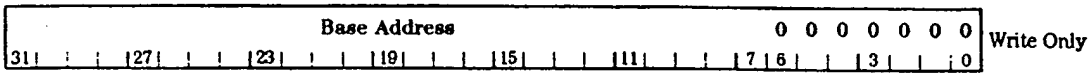

Bit(s)	Field Label	Field Name	Description
31..0	Base Address	Base Address	The physical address in host memory of the start of the IMQ.

Table 5.13 IMQ Base Register

IMQ Base Register Notes

1. The IMQ must be aligned on a sizeof(queue) boundary. The minimum size of the IMQ is 128 bytes (bits 6..0 are cleared to zero).
2. When the host writes to the IMQ Base register, Tachyon clears the corresponding producer index to zero; therefore, the host must not write to this register after initialization.

5.5.2 IMQ Length Register

Register Address: 0x0084
Reset Value: 0x00000000

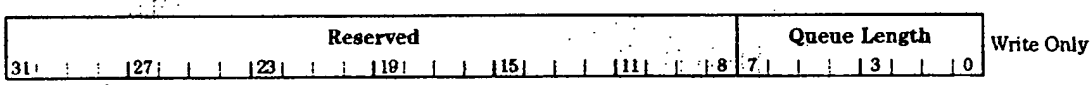

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Length	Queue Length	The number (zero-based) of IMQ entries.

Table 5.14 IMQ Length Register

IMQ Length Register Notes

1. The Queue Length field must be a power of 2. The minimum queue length is 4 entries corresponding to a programmed value of 3. The maximum queue length is 256 entries corresponding to a programmed value of 255.
2. When the host writes to the IMQ Length register, Tachyon clears the corresponding consumer index to zero; therefore, the host must not write to this register after initialization.

PTI 172557

5.5.3 IMQ Consumer Index Register

Register Address: 0x0088
Reset Value: 0x00000000

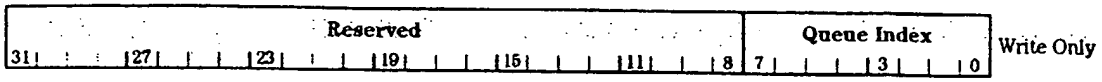

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The consumer index value of the IMQ entry following the last entry processed by the host.

Table 5.15 IMQ Consumer Index Register

IMQ Consumer Index Register Notes

1. This register is used by the host to indicate the completion messages that have been processed by the host and the number of IMQ entries available to Tachyon for posting new completion messages. The host must process completion messages in sequential order and return them to Tachyon in sequential order. As Tachyon generates each interrupt, the host reads the copy of Tachyon's producer pointer in host memory. This value, along with the host's copy of its consumer index, gives the host an indication of how many completion messages Tachyon has posted since the last interrupt. The host processes each completion message and then writes the index of the next entry following the last one it processed.
2. Tachyon uses a pulsed interrupt line that is asserted only once until the host writes back to the IMQ consumer index. This is important in implementations that use latched interrupts. Always enable interrupts before writing the IMQ consumer index. A race condition may occur if the following algorithm is not used.

```
do interrupts
{
     ...
     ...
}
enable interrupts
write IMQ consumer index
```

PTI 172558

5.5.4 IMQ Producer Index Address Register

Register Address: 0x008C
Reset Value: 0x00000000

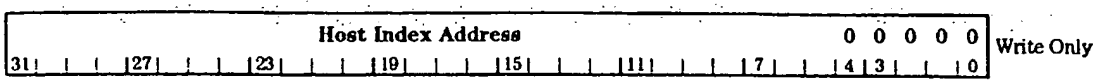

Bit(s)	Field Label	Field Name	Description
31..0	Host Index Address	Host Index Address	The address of the host's copy of the IMQ producer index.

Table 5.16 IMQ Producer Index Address Register

IMQ Producer Index Address Register Notes

1. The Inbound Message Queue Index Address register contains the host memory address where Tachyon maintains its producer index.
2. The host index address must be aligned on a 32-byte boundary (bits 4..0 are cleared to zero).
3. This index resides in host memory to allow the host fast access to the producer index.
4. The host must not access this register after initialization. The host enables the IMQ by writing to the IMQ Producer Index Address register; therefore, the host should initialize this register after the IMQ Base and IMQ Length registers are initialized.

5.5.5 Host's Copy of the IMQ Producer Index

Register Address: Determined by the host
Reset Value: Determined by the host

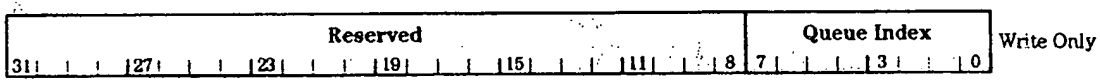

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The address of the host's copy of the producer index value of the IMQ.

Table 5.17 Host's Copy of the IMQ Producer Index

Host's Copy of the IMQ Producer Index Register Note

1. This memory location described by the IMQ Producer Index Address register contains the host's copy of Tachyon's producer index.
2. Tachyon updates this index when it writes an IMQ entry.

PTI 172559

5.6 MFSBQ Registers

5.6.1 MFSBQ Base Register

Register Address: 0x00C0
Reset Value: 0x00000000

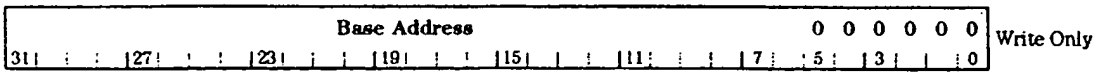

Bit(s)	Field Label	Field Name	Description
31..0	Base Address	Base Address	The physical address in host memory of the start of the MFSBQ.

Table 5.18 MFSBQ Base Register

MFSBQ Base Register Notes

1. The MFSBQ must be aligned on a sizeof(queue) boundary.
2. The minimum size for the MFSBQ is 64 bytes (bits 5..0 are cleared to zero).
3. When the host writes a value to the MFSBQ Base register, Tachyon clears the corresponding consumer index to zero; therefore, the host must not write to this register after initialization.

5.6.2 MFSBQ Length Register

Register Address: 0x00C4
Reset Value: 0x00000000

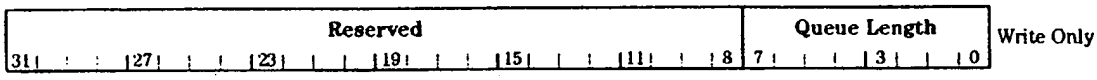

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Length	Queue Length	The number (zero-based) of MFSBQ entries.

Table 5.19 MFSBQ Length Register

MFSBQ Length Register Notes

1. The Queue Length field must be a power of 2. The minimum queue length is 2 entries, corresponding to a programmed value of 1. The maximum queue length is 256 entries, corresponding to a programmed value of 255.
2. When the host writes to the MFSBQ length register, Tachyon clears the corresponding consumer index to zero; therefore, the host must not write to this register after initialization.

PTI 172560

5.6.3 MFSBQ Producer Index Register

Register Address: 0x00C8
Reset Value: 0x00000000

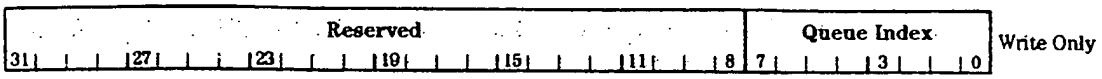

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The producer index value that points to the end of the MFSBQ.

Table 5.20 MFSBQ Producer Index Register

MFSBQ Producer Index Register Notes

1. The host uses the MFSBQ Producer Index Register to pass buffers to Tachyon for use in reassembling incoming multiframe sequences in host memory.
2. The host must initialize all other MFSBQ registers before it writes to this register.

5.6.4 MFSBQ Consumer Index Register

Register Address: 0x00CC
Reset Value: 0x00000000

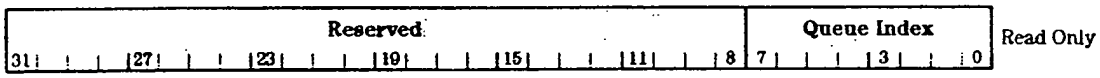

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The MFSBQ entry that Tachyon is currently using to process a MFS.

Table 5.21 MFSBQ Consumer Index Register

MFSBQ Consumer Index Register Note

1. Tachyon updates this register to indicate the MFSBQ entry that it is currently using to process the incoming multiframe sequence.

PTI 172561

5.6.5 MFS Buffer Length Register

Register Address: 0x00D0
Reset Value: 0x00000000

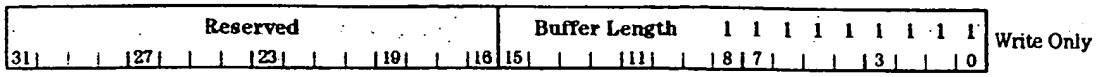

Bit(s)	Field Label	Field Name	Description
31..16	Reserved	Reserved	Initialize to zero.
15..0	Buffer Length	Buffer Length	The (zero-based) length in bytes of the MFS receive buffers.

Table 5.22 MFS Buffer Length Register

MFS Buffer Length Register Notes

1. The Buffer Length field must be a power of 2 bytes.
2. The minimum buffer length is 512 bytes corresponding to a programmed value of 511.
3. The reset value of this register is 511.
4. The maximum buffer length is 65,536 bytes, corresponding to a programmed value of 65,535.
5. When the host writes to this register, Tachyon masks bits 8..0 and always sets them to one.
6. The host should not modify the contents of this register after initialization.

PTI 172562

5.7 SFSBQ Registers

5.7.1 SFSBQ Base Register

Register Address: 0x0100
Reset Value: 0x00000000

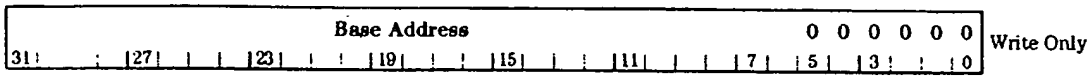

Bit(s)	Field Label	Field Name	Description
31..0	Base Address	Base Address	The physical address in host memory of the start of the SFSBQ.

Table 5.23 SFSBQ Base Register

SFSBQ Base Register Notes

1. The SFSBQ must be aligned on a sizeof(queue) boundary.
2. The minimum size of the SFSBQ is 64 bytes (bits 5..0 are cleared to zero).
3. When the host writes to the SFSBQ Base register, Tachyon clears the corresponding consumer index to zero;therefore, the host must not write to this register after initialization.

5.7.2 SFSBQ Length Register

Register Address: 0x0104
Reset Value: 0x00000000

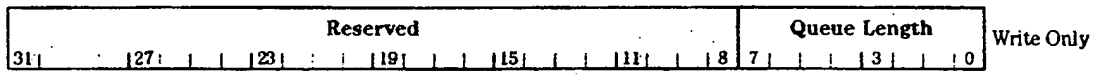

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Length	Queue Length	The number (zero-based) of SFSBQ entries.

Table 5.24 SFSBQ Length Register

SFSBQ Length Register Notes

1. The Queue Length field must be a power of 2. The minimum queue length is 2 entries corresponding to a programmed value of 1. The maximum queue length is 256 entries corresponding to a programmed value of 255.
2. When the host writes to the SFSBQ Length register, Tachyon clears the corresponding consumer index to zero;therefore, the host must not write to this register after initialization.

PTI 172563

5.7.3 SFSBQ Producer Index Register

Register Address: 0x0108
Reset Value: 0x00000000

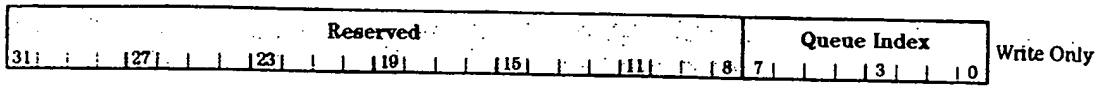

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The producer index value that points to the end of the SFSBQ.

Table 5.25 SFSBQ Producer Index Register

SFSBQ Producer Index Register Notes

1. The host uses the SFSBQ Producer Index register to pass buffers to Tachyon for use in receiving incoming single frame sequences in host memory.
2. The host must initialize all other SFSBQ registers before it writes to this register.

5.7.4 SFSBQ Consumer Index Register

Register Address: 0x010C
Reset Value: 0x00000000

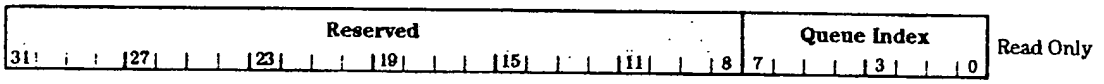

Bit(s)	Field Label	Field Name	Description
31..8	Reserved	Reserved	Initialize to zero.
7..0	Queue Index	Queue Index	The SFSBQ entry Tachyon is currently using to process SFSs.

Table 5.26 SFSBQ Consumer Index Register

SFSBQ Consumer Index Register Note

1. Tachyon updates this register with the index of the SFSBQ entry that it is currently using to process the incoming single frame sequences.

PTI 172564

5.7.5 SFS Buffer Length Register

Register Address: 0x0110
Reset Value: 0x00000000

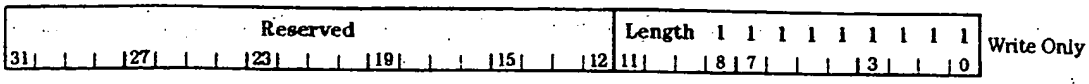

Bit(s)	Field Label	Field Name	Description
31..16	Reserved	Reserved	Initialize to zero.
15..0	Length	Buffer Length	The (zero-based) length in bytes of the SFS receive buffers.

Table 5.27 SFS Buffer Length Register

SFS Buffer Length Register Notes

1. The Length field must be a power of 2 bytes.
2. The minimum buffer length is 512 bytes corresponding to a programmed value of 511.
3. The reset value of this register is 31.
4. The maximum buffer length is 4096 bytes, corresponding to a programmed value of 4095.
5. When the host writes to this register, Tachyon masks bits 8..0 and always sets them to one.
6. The host should not modify the contents of this register after initialization.

PTI 172565

5.8 SEST Registers

5.8.1 SEST Base Register

Register Address: 0x0140
Reset Value: 0x00000000

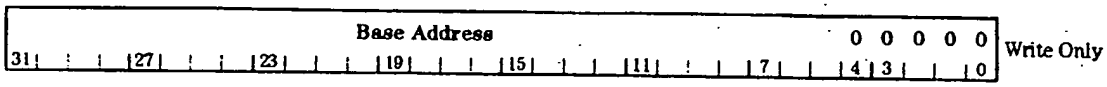

Bit(s)	Field Label	Field Name	Description
31..0	Base Address	Base Address	The physical address in host memory of the start of the SEST.

Table 5.28 SEST Base Register

SEST Base Register Notes

1. The SEST must be aligned on a sizeof(SEST) boundary. For example, if the SEST is 512 bytes, then bits 8..0 of the SEST Base register are cleared to zero.
2. The minimum size for the SEST is 32 bytes (bits 4..0 are cleared to zero).

5.8.2 SEST Length Register

Register Address: 0x0144
Reset Value: 0x00000000

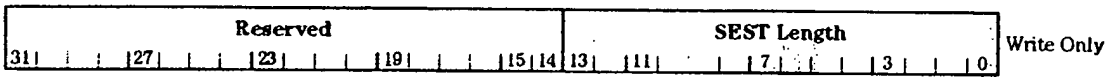

Bit(s)	Field Label	Field Name	Description
31..14	Reserved	Reserved	Initialize to zero.
13..0	SEST Length	SEST Length	The number (zero-based) of SEST entries.

Table 5.29 SEST Length Register

SEST Length Register Note

1. The SEST Length must be a power of 2. The minimum SEST Length is 1 entry, corresponding to a programmed value of 0. The maximum SEST Length is 65,536 entries corresponding to a programmed value of 65,535.

PTI 172566

5.8.3 SCSI Buffer Length Register

Register Address: 0x0148
Reset Value: 0x00000000

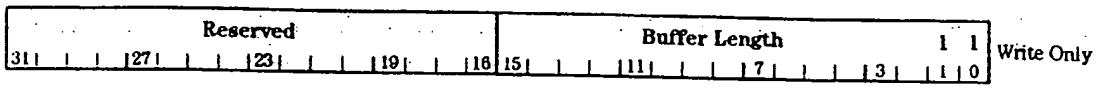

Bit(s)	Field Label	Field Name	Description
31..16	Reserved	Reserved	Initialize to zero.
15..0	Buffer Length	Buffer Length	The (zero-based) length in bytes of the SCSI receive buffers.

Table 5.30 SCSI Buffer Length Register

SCSI Buffer Length Register Notes

1. The host should not modify the contents of this register after initialization.
2. If the host cleared bit 6 in the Tachyon Configuration register to zero (OOO Reassembly) , then the buffer size must be a power of 2, and the minimum value is 512 bytes, corresponding to a programmed value of 511.
3. If the host sets bit 6 in the Tachyon Configuration register to one (In Order Reassembly), then the buffer size must be a multiple of 4 bytes, so that the minimum value is 4 bytes, corresponding to a programmed value of 3.
4. The maximum value is 65,536 bytes, corresponding to a programmed value of 65,535 bytes.
5. When the host writes to this register, it must always set bits 1..0 each to one.

PTI 172567

5.9 Tachyon Registers

5.9.1 Tachyon Configuration Register

Register Address: 0x0184
Reset Value: 0x00000002

ne	se	Reserved	db	Split Type	Res	bs	sc	ws	rs	fn	od	ad	rd	pt	pa	pe	st	Read/Write
31	30	29..25	24	23..16	15..14	13	12	11..10	9..8	7	6	5	4	3	2	1	0	

Bit(s)	Field Label	Field Name	Description
31	ne	TCP/UDP Assists Enable	Set to one to enable TCP/UDP hardware assists.
30	se	SCSI Enable	Set to one to enable SCSI FCP hardware assists.
29..25	Reserved	Reserved	Initialize to zero.
24	db	Disable AUTO P_BSY	When the host sets this bit to one, Tachyon does not automatically P_BSY new MFS frames while a current MFS is being re-assembled. Instead, it passes the new frames to the host.
23..16	Split Type	Split Type	The host-specified Fibre Channel frame type for which to perform header/data splits. Tachyon splits by default for type 5 (LLC/SNAP packets). The host uses this field to specify an additional frame type. If the host clears this field to zero, Tachyon disables splitting for a second type.
15..14	Res	Reserved	Initialize to zero.
13	bs	Bad SCSI Auto ACK	If the host sets this bit to one, Tachyon generates an ACK for bad SCSI frames. If the host clears this bit to zero, the host is responsible for generating the ACK.
12	sc	SCSI Command Auto ACK	If the host sets this bit to one, Tachyon generates an ACK for the SCSI command frame. If the host clears this bit to zero, the host is responsible for generating the ACK.
11..10	ws	Write Stream Size	The maximum number of Write transactions to perform per bus tenancy.
			ws value: 00 → Maximum Stream Count: 1
			ws value: 01 → Maximum Stream Count: 4
			ws value: 10 → Maximum Stream Count: 16
			ws value: 11 → Maximum Stream Count: 64
9..8	rs	Read Stream Size	The maximum number of Read transactions to perform per bus tenancy.

Table 5.31 Tachyon Configuration Register

PTI 172568

Bit(s)	Field Label	Field Name	Description	
			rs value	Maximum Read Stream Count
			00	1
			01	4
			10	16
			11	64
7	fn	ACK Generation Assist Enable	The host sets this bit to one to inform Tachyon, as the recipient of frames, to look at the ACK generation assist bits (bits 12 and 13) of the F_CTL field of the Fibre Channel header. These bits inform Tachyon which ACK model to use when responding to received frames. Refer to "3.4.2 Acknowledgement of Received Frames" on page 34.	
6	od	OOO Reassembly Disable	When the host sets this bit to one, Tachyon does not reassemble OOO SCSI FCP frames. This allows for more freedom in selecting receive buffer sizes. The host should not set this bit if the topology allows OOO delivery.	
5	ad	ACK Disable	When the host sets this bit to one, Tachyon disables automatic ACK generation. Tachyon then passes all received frames to the host as unknown frames via the SFSBQ. Tachyon does not perform any automatic processing of sequences. The host should use this for debugging purposes only.	
4	rd	Retry Disable	When the host sets this bit to one, Tachyon disables retries of the first frame of a sequence. The first BSY that Tachyon receives will terminate the sequence.	
3	pt	Point-to-Point	The host sets this bit to one to indicate that it has detected that Tachyon is directly connected to another N_Port. If this bit is set to one and the OSM tries to establish an outgoing Class 1 connection, then the ISM busies any attempts for inbound Class 1 connections.	
2	pa	Parity Enable	Set to one to enable parity checking within Tachyon.	
1	pe	Parity Even	The host sets this bit to one for even parity and clears it to zero for odd parity. This bit is valid only if the host has parity enabled. The default value for this bit is one. Refer to "7.6.14 Parity" on page 213.	
0	st	Stacked Connect Requests	The host sets this bit to one when it detects that the fabric supports stacked connect requests.	

Table 5.31 Tachyon Configuration Register (Continued)

Tachyon Configuration Register Notes

1. The host uses the Tachyon Configuration register to initialize and configure Tachyon's operating modes.
2. The Stacked Connect Requests (st) and Point-to-Point (pt) bits of the Tachyon Configuration register are the only bits that the host can modify after initial configuration.

5.9.2 Tachyon Control Register

Register Address: 0x0188
Reset Value: 0x00000000

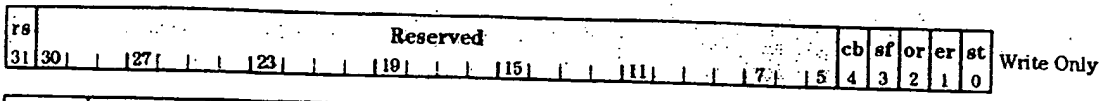

Bit(s)	Field Label	Field Name	Description
31	rs	Software Reset	The host sets this bit to one to initiate a software reset of Tachyon. 1. At reset, Tachyon clears all registers to zero, unless otherwise noted. Reset is defined as power-on, a hardware reset, or a software reset. 2. While Tachyon clears registers to zero at reset, certain registers (such as EE_Credit Zero Timer and BB_Credit Zero Timer registers) begin counting up immediately following a reset. As a result, if the host reads a register immediately after a reset, the value may be non-zero. 3. If the host reads an invalid or Write only (W) register, the transaction completes normally, and Tachyon returns 0x00000000. 4. Tachyon ignores bits 31..9 of the address during register accesses. 5. If the host writes to Read only (R) registers, the transaction completes but Tachyon does not write the data to the register.
30..5	Reserved	Reserved	Initialize to zero.
4	cb	Clear Deferred P_BSY	The host sets this bit to one to clear Tachyon's Deferred P_BSY flag. This allows Tachyon to ACK MFSs normally until the next MFS collision occurs.
3	sf	SCSI Freeze	The host sets this bit to one to stop the processing of XFER_RDYs for outbound exchanges and Outbound SEST Entries. The host clears this bit to zero to resume processing. This bit is cleared to zero for normal operation. This bit affects outbound operations only.
2	or	OCQ Reset	The host sets this bit to one to clear the OCQ parameters. Tachyon resets the OCQ producer and consumer indices to zero. Tachyon may not return an outbound_completion message for the last OCQ entry read if the OCQ is reset.

Table 5.32 Tachyon Control Register

PTI 172570

1	er	Error Release	The host sets this bit to one to unfreeze the OSM from the Error state. Tachyon does not allow the OSM to unfreeze if the Link Down bit in the Frame Manager Status register is set to one.
0	st	Status Request	The host sets this bit to one to request Tachyon to DMA current status to the host.

Table 5.32 Tachyon Control Register (Continued)

Tachyon Control Register Notes

1. The host uses the Tachyon Control register to reset all or portions of Tachyon as well as to request a dump of all internal Tachyon status.
2. The OCQ reset does not necessarily occur the instant that the host writes to the Tachyon Control register. Tachyon performs the reset when the channel is at an idle point to prevent hanging the remainder of Tachyon in the middle of an operation. The Tachyon Status register has a status bit that indicates whether or not the reset procedure is still in progress. Tachyon should always complete the reset procedure in less than ~30 clocks, but Tachyon updates the reset status bit for cases where the host must guarantee the reset is complete before proceeding.

5.9.3 Tachyon Status Register

Register Address: 0x018C
Reset Value: HW dependent

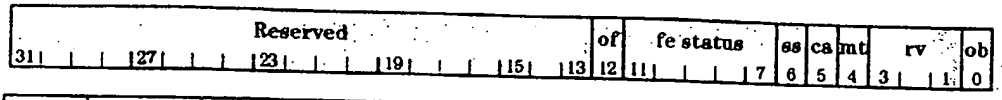

Bit(s)	Field Label	Field Name	Description
31..13	Reserved	Reserved	Value is undefined.
12	of	Outbound FIFO	Tachyon sets this bit to one when the outbound FIFO either is empty or is in the process of being emptied by the Frame Manager. If Tachyon clears this bit to zero, then the FIFO contains a frame that is not being moved. Refer to "10.3.1 Blocked Outbound Frame FIFO error recovery" on page 259.
11..7	fe status	Fatal Error Status	Allows the host to determine what caused ERROR_L to be asserted.

State	Definition
00000	Tachyon has not logged an error
00001	Frame Manager outbound data parity error
00010	Slave write data parity error
00100	Type signals not asserted correctly for slave transactions.
01000	Address parity error
10000	DMA read data parity error

6	ss	SCSI Freeze Status	Tachyon sets this bit to one to indicate that the processing of XFER_RDYs and Outbound SEST Entries in the SEST has stopped. This bit corresponds to bit 3 of the Tachyon Control register.
5	ca	OCQ Reset Status	Tachyon sets this bit to one to indicate that a reset of the OCQ is in progress.
4	mt	Receive FIFO Empty	Tachyon sets this bit to one to indicate that the inbound data FIFO is empty.
3..1	rv	Chip Revision	Tachyon's hardware revision code. Refer to the Tachyon Status register note 2 on page 136.
0	ob	OSM Frozen	Tachyon sets this bit to one to indicate that the OSM is in a Frozen state due to an error. Refer to "10.2 OSM Freeze" on page 258.

Table 5.33 Tachyon Status Register

PTI 172572

Tachyon Status Register Notes

1. The host uses the Tachyon Status register to read error information and operational status of the Tachyon chip.
2. The Chip Revision field reads as follows:

Description	HP Part Number	Bits 3..1	Release Date
Pre-Release 5V Tachyon	HPFC-300K	000	January, 1995
Pilot 3.3V Tachyon	HPFC-500P (HPFC-5000)	001	July, 1995
Pre-Production 3.3V Tachyon Revision 1	HPFC-500K (HPFC-5000)	001	August, 1995
Pre-Production 3.3V Tachyon Revision 2	HPFC-5000	010	February, 1996
Production 3.3V Tachyon Revision 2	HPFC-5000	010	March, 1996

Table 5.34 Tachyon Chip Revisions

No functional differences exist between HPFC-500P and HPFC-500K (Revision 001b). No functional differences are planned between the Pre-Production 3.3V Tachyon Revision 2 and the Production 3.3V Tachyon Revision 2 (Revision 010b).

PTI 172573

5.9.4 Tachyon Flush SEST Cache Entry Register

Register Address: 0x0190
Reset Value: 0x00000000

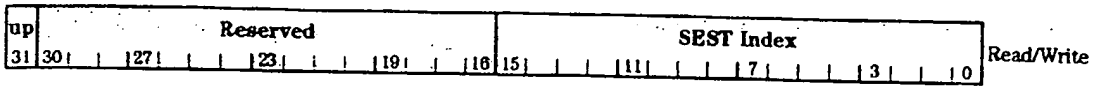

Bit(s)	Field Label	Field Name	Description
31	up	Update	This bit indicates that Tachyon is able to handle an OX_ID flush request from the host. Set to one by the host when it wants to update the SEST entry. Cleared to zero by Tachyon when the cache update is complete.
30..16	Reserved	Reserved	Initialize to zero.
15..0	SEST Index	SEST Index	The index of the SEST entry the host wants updated.

Table 5.35 Tachyon Flush SEST Cache Entry Register

Tachyon Flush SEST Cache Entry Register Notes

1. The host uses this register to force Tachyon to write back any internal status for the indicated Inbound SEST Entry and to invalidate the Inbound SEST Entry. The host performs this operation when it detects an error with the exchange and wants to abnormally terminate it. The host should not use this register to flush Outbound SEST Entries since they are not cached on Tachyon. To terminate an outbound transaction abnormally, the host needs to clear only the Valid bit in the Outbound SEST Entry.
2. If no internal status is associated with the indicated Inbound SEST Entry, then Tachyon does nothing. Therefore, the host should clear the Valid bit in the Inbound SEST Entry before writing the register to prevent race conditions where Tachyon is trying to start processing an incoming transfer at the same time the host is trying to invalidate and flush the Inbound SEST Entry.

5.9.5 Tachyon EE_Credit Zero Timer Register

Register Address: 0x0194
Reset Value: 0x00000000

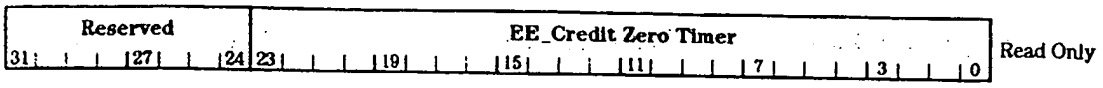

Bit(s)	Field Label	Field Name	Description
31..24	Reserved	Reserved	Value is undefined.
23..0	EE_Credit Zero Timer	EE_Credit Zero Timer	The amount of time Tachyon has waited to transmit, but could not, because it had no EE_Credit.

Table 5.36 Tachyon EE_Credit Zero Timer Register

Tachyon EE_Credit Zero Timer Register Note

1. This register provides a timer that runs whenever the EE_Credit for a sequence goes to zero. This indicates to the host the time that Tachyon is congested due to lack of EE_Credit. The timer is clocked by a 10us pulse which can result in the timer rolling over approximately every two minutes in a heavy congestion environment. If the host uses this timer, the host should read the timer frequently enough to detect rolling over.

5.9.6 Tachyon BB_Credit Zero Timer Register

Register Address: 0x0198
Reset Value: 0x00000000

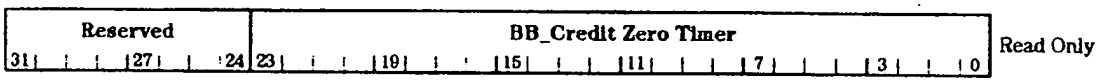

Bit(s)	Field Label	Field Name	Description
31..24	Reserved	Reserved	Value is undefined.
23..0	BB_Credit Zero Timer	BB_Credit Zero Timer	The amount of time Tachyon has waited to transmit, but could not, because it had no BB_Credit.

Table 5.37 Tachyon BB_Credit Zero Timer Register

Tachyon BB_Credit Zero Timer Register Note

1. This register provides a timer that runs whenever the BB_Credit goes to zero. This indicates to the host the time that Tachyon is congested due to lack of credit. The timer is clocked by a 10us pulse which can result in the timer rolling over approximately every two minutes in a heavy congestion environment. If the host uses this timer, the host should read the timer frequently enough to detect rolling over.

PTI 172575

5.9.7 Tachyon Receive Frame Error Counter Register

Register Address: 0x019C
Reset Value: 0x00000000

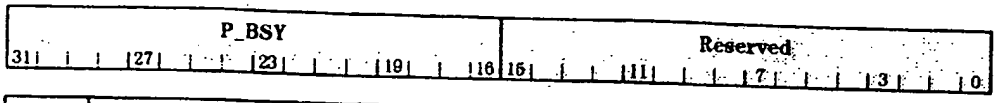

Read Only

Bit(s)	Field Label	Field Name	Description
31..16	P_BSY	P_BSYs Sent	The number of frames received whose response was a P_BSY. Reading this register resets this counter to zero.
15..0	Reserved	Reserved	Value is undefined.

Table 5.38 Tachyon Receive Frame Error Counter Register

Tachyon Receive Frame Error Counter Register Note

1. This register allows the host to monitor the number of busies (P_BSYs) transmitted by Tachyon.

5.10 Frame Manager Registers

5.10.1 Frame Manager Configuration Register

Register Address: 0x01C0
Reset Value: 0x00000000

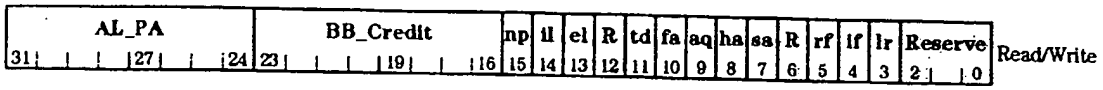

Bit(s)	Field Label	Field Name	Description
31..24	AL_PA	AL_PA	The AL_PA is the address the Frame Manager uses when arbitrating on the loop and the address that it recognizes in an OPN primitive signal. The host sets the AL_PA to an initial value. Tachyon sets this address after acquiring an AL_PA on the loop.
23..16	BB_Credit	BB_Credit	The amount of BB_Credit, as determined from login parameters. Valid only if Tachyon is an N_Port.
15	np	N_Port	Force Tachyon to start initialization as a non-loop N_Port. This bit should be cleared before issuing the Offline command (Frame Manager Control register).
14	il	Enable Internal Loopback	Enable Tachyon's internal loopback path.
13	el	Enable External Loopback	Assert the PLM loopback signal (EWRAP) to loopback data at the PLM.
12	R	Reserved	Initialize to zero.
11	td	Timer Disable	Disable the ED_TOV timer. For debugging purposes only.
10	fa	Loop Initialization Fabric Acquired Address	Indicates whether the AL_PA field was previously assigned by a fabric. If Tachyon does not acquire the AL_PA the host wrote into the AL_PA field, then Tachyon clears this bit to zero.
9	aq	Loop Initialization Previously Acquired Address	Indicates whether the AL_PA field was previously assigned by Loop Initialization. Tachyon sets this bit to one after acquiring an AL_PA on the loop. Refer to Frame Manager Configuration register note 3 on page 141.
8	ha	Loop Initialization Hard (Preferred) Address	Indicates whether the host set the AL_PA field from switches or another fixed address method. Tachyon clears this bit to zero after acquiring an AL_PA on the loop.

Table 5.39 Frame Manager Configuration Register

PTI 172577

Bit(s)	Field Label	Field Name	Description
7	sa	Loop Initialization Soft Address	Indicates no AL_PA was provided, so Tachyon must participate in self-assignment during Loop Initialization. After acquiring an AL_PA on the loop, Tachyon clears this bit to zero.
6	R	Reserved	Initialize to zero.
5	rf	Respond to Fabric Address	Recognize the fabric AL_PA value of zero in addition to the assigned AL_PA. The host should not set this bit to one if a fabric exists on the loop.
4	if	Initialize as Fabric	This bit causes the Frame Manager to use a D_ID and S_ID of zero when generating LISM frames on FC-AL. The host should not set this bit to one if a fabric exists on the loop.
3	lr	Login Required	This bit indicates that Tachyon should assert the Fabric Login bit in LIFA, LIPA, LIHA, and LISA frames. This bit is valid only when the host also sets the Initialize as Fabric (if) bit.
2..0	Reserve	Reserved	Initialize to zero.

Table 5.39 Frame Manager Configuration Register (Continued)

Frame Manager Configuration Notes

1. The host should only modify the Frame Manager Configuration register while the Frame Manager is in the OFFLINE state (except for BB_Credit, refer to Note 2. below). For a Fabric as an L_Port, other Frame Manager Configuration register bits may change. Refer to "3.9.15 Fabric Operation" on page 90.
2. The BB_Credit field (bits 23..16) and the Fabric Acquired Address bit (fa, bit 10) are the only fields/bits which can be modified while the Frame Manager is online.
3. While online as an N_Port, if the BB_Credit field is changed, the host should issue the Link Reset command (in the Frame Manger Control register) to synchronize the BB_Credit with the remote node.
4. If BB_Credit is not initialized to a non-zero value, then Tachyon never issues a BB_Credit warning in N_Port mode. The BB_Credit default is zero.
5. If the host sets the Fabric Acquired Address (fa, bit 10) to one, then the host must also set the Previously Acquired Address (aq, bit 9) to one or an error condition may occur.
6. Setting more than one of the following bits is illegal: Previously Acquired Address (aq, bit 9), Hard Address (ha, bit 8), and Soft Address (sa, bit 7). Tachyon does not check for this situation. Therefore, setting more than one address bit produces unpredictable results. The host must set one of these bits (fa, aq, ha, and sa) for proper loop participation; otherwise, Tachyon cannot acquire an AL_PA.
7. If the host wants non-participating mode on the loop, the host should clear bits 7..10 to zero.

PTI 172578

5.10.2 Frame Manager Control Register

Register Address: 0x01C4
Reset Value: 0x00000000

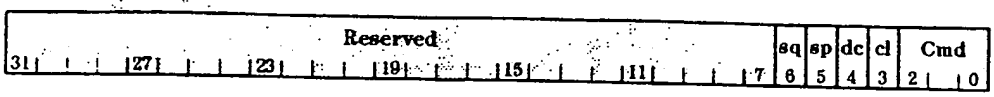

Bit(s)	Field Label	Field Name	Description
31..7	Reserved	Reserved	Initialize to zero.
6	sq	Primitive Sequence	This bit has a meaning only in the Host Control state when a one is written to the Send Primitive Register (sp) bit. One occurrence of the primitive in the Primitive Register field of the Frame Manager Primitive register is sent when this bit contains a one. At least 12 occurrences of the primitive are sent when this bit contains a zero.
5	sp	Send Primitive Register	This bit has a meaning only in the Host Control state. It has no meaning when cleared to zero. When a one is written to this bit, it directs the Loop State Machine to send the primitive in the Primitive Register field of the Frame Manager Primitive register. The number of times the primitive is sent is controlled by the Primitive Sequence (sq) bit. This bit is automatically cleared to zero after the primitive(s) have been sent. Writing a one to this bit while primitives are being sent has no effect.
4	dc	Don't Close Loop Request	This bit forces the Loop State Machine to remain open after sending a sequence. When this bit is set to one, the loop is not closed until the Close Loop Request (cl) bit in the Frame Manager Control register is set to one. The Close Loop Request bit overrides the effect of the Loop Close (cl) bit in the Tachyon Header Structure for the current outbound sequence, preventing CLOSE after any frame in the sequence. When Tachyon closes the loop, it does not automatically clear the Don't Close Loop Request (dc) bit to zero.
3	cl	Close Loop Request	This bit forces the Loop State Machine to send a loop CLS. Tachyon does not transmit the CLS until it has sent all of the frame data from the outbound FIFO. After Tachyon closes the loop, it automatically clears this bit to zero.
2..0	Cmd	State Machine Command	Command Port State Machine operations:

Field Value	Command Name	Description

Table 5.40 Frame Manager Control Register

PTI 172579

Bit(s)	Field Label	Field Name	Description		
			000 001	NOP	No command.
			010	Host Control	Directs the Loop State Machine to transition to the Host Control state.
			011	Exit Host Control	Directs the Loop State Machine to exit the Host Control state.
			100	Link Reset	Forces the Port State Machine to the LR1 state where LR is transmitted.
			101	Offline	Forces the Frame Manager to go offline and stay offline. This is the power up state. Refer to "3.9.12 Programming the Frame Manager OFFLINE" on page 89.
			110	Initialize	Forces the Frame Manager to start initialization. Refer to the Frame Manager Control register note 1 below.
			111	Clear LF	Enables the Port State Machine to transition out of the Link Failure state.

Table 5.40 Frame Manager Control Register (Continued)

Frame Manager Control Register Note

1. If the host programs the Initialize Command when the link is up in N_port mode, Tachyon ignores the command. If the host programs the Initialize Command when the link is up in L_port mode, Tachyon performs Loop Initialization.

PTI 172580

5.10.3 Frame Manager Status Register

Register Address: 0x01C8
Reset Value: 0x00000000

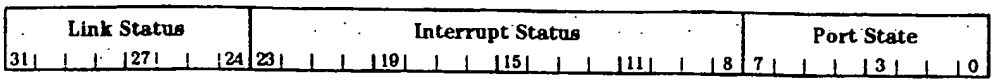

Bit(s)	Field Label	Field Name	Description
	Link Status	Link Status	Link Status comprises the following sub-fields.
31	lp	Loop	A loop circuit has been established.
30	tp	Transmit Parity Error	This bit is set to one when Tachyon detects an internal data parity error. The host clears this error only by performing a hardware reset.
29	np	Non-Participating	The port could not successfully complete Loop Initialization and is not participating in the loop.
28..27	par_id	Parallel ID	These two bits reflect the external PAR_ID bits received by the Frame Manager. Bit 28 is PAR_ID[1] and bit 27 is PAR_ID[0]. Refer to "7.3 Physical Link Module (PLM) Interface" on page 200.
26	lflt	Laser Fault	The PLM has detected a laser fault. The host must reset the PLM with EWRAP.
25	os	Out of Synchronization	The port has lost synchronization with the incoming data stream.
24	ls	Loss of Signal	The port is not receiving a signal from the remote node.
	Interrupt Status	Interrupt Status	Interrupt Status comprises the following sub-fields. These conditions generate a completion message and an interrupt to the host. The host clears these conditions by reading this register.
23..20	Reserved	Reserved	Value is undefined.
19	nos/ols	NOS/OLS Received	Tachyon received a NOS or OLS.
18	lst	Loop State Timeout	The Loop State Machine is in one of the following states: ARB, OPEN, OPENED, XMIT CLS, or RX CLS for longer than ED_TOV.
17	lipf	LIPf	LIPf has been received. A LIPf is a LIP indicating a loop failure, either LIP(F8, AL_PS) or LIP(F8, F7).
16	ba	Bad AL_PA	This bit indicates that an OPN primitive signal was sent to an AL_PA which did not respond.

Table 5.41 Frame Manager Status Register

PTI 172581

Bit(s)	Field Label	Field Name	Description
15	pr_rxd	Primitive Received	The primitive in the Primitive Register field of the Frame Manager Primitive register was received while in the Host Control state.
14	pr_sent	Primitive Transmitted	The primitive in the Primitive Register field of the Frame Manager Primitive register has been transmitted (12 times if the sq bit of the Frame Manager Control register is set to one.)
13	lg	Fabric Login Required	Loop Initialization detected a new fabric. The N_Port must perform fabric login.
12	lf	Link Failure	This bit indicates that one of the following conditions has been detected: Loss of Sync for greater than ED_TOV, Loss of Signal, Laser Fault, or Elastic Store Error.
11	ce	Credit Error	This bit indicates a BB_Credit of zero for ED_TOV error condition.
10	ew	Elastic Store Error	This bit indicates an error condition caused by an Elastic Store over-write or under-read. This indicates a serious timing difference between the receive and transmit clocks.
9	lup	Link Up	This bit indicates that the link has gone into the ACTIVE state or that the link has transitioned from O_I_PROTOCOL to MONITORING on the loop. If both this bit and the ldwn bit are asserted, the loop has reinitialized or a link reset has occurred.
8	ldwn	Link Down	This bit indicates that the link has gone inactive or that Loop Initialization has been reinitialized. The OSM freezes when the link goes down and needs to be unfrozen when the link is back up. Tachyon does not unfreeze the OSM if this bit is set to one.
	Port State	Port State	Port State comprises the following sub-fields.
7..4	lfsm	Loop State Machine	The Fibre Channel Arbitrated Loop State Machine. Refer to the "FC-AL" for state definitions.
3..0	pfsm	Port State Machine	N_Port State Machine. Refer to the FC-PH for state defintions.

State	Definition	State	Definition
0000	MONITORING	1000	INITIALIZING
0001	ARBITRATING	1001	O_I INIT FINISH
0010	ARBITRATION WON	1010	O_I PROTOCOL
0011	OPEN	1011	O_I LIP RECEIVED
0100	OPENED	1100	HOST CONTROL
0101	XMITTED CLOSE	1101	LOOP FAIL
0110	RECEIVED CLOSE	1111	OLD PORT
0111	TRANSFER		

Table 5.41 Frame Manager Status Register (Continued)

PTI 172582

Bit(s)	Field Label	Field Name	Description			
			State	Definition	State	Definition
			0000	OFFLINE	0110	LR2
			0001	OL1	0111	LR3
			0010	OL2	1001	LF1
			0011	OL3	1010	LF2
			0101	LR1	1111	ACTIVE

Table 5.41 Frame Manager Status Register (Continued)

PTI 172583

5.10.4 Frame Manager RT_TOV/AL_TIME & ED_TOV Register

Register Address: 0x01CC
Reset Value: 0x001001F5

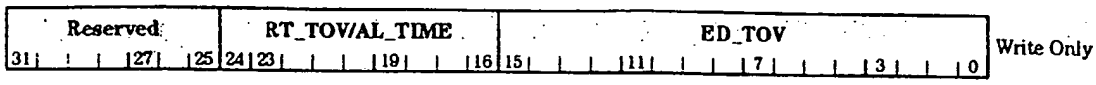

Bit(s)	Field Label	Field Name	Description
31..25	Reserved	Reserved	Initialize to zero.
24..16	RT_TOV/AL_TIME	RT_TOV/AL_TIME	The host uses this register to program both the Receiver Transmitter Time Out Value (RT_TOV) and the Arbitrated Loop Time Out (AL_TIME) value. The host should program this value to 100 milliseconds (ms) for RT_TOV and 15 ms for AL_TIME. The default is 15 ms. If the host does not know the topology prior to initialization, this value must be programmed to 15 ms for Loop Initialization and reprogrammed to 100 ms if it is determined that no loop exists.
15..0	ED_TOV	ED_TOV	The Error Detect Time Out Value (ED_TOV) in milliseconds. The default is 500 ms.

Table 5.42 Frame Manager RT_TOV/AL_TIME and ED_TOV Register

Frame Manager RT_TOV/AL_TIME & ED_TOV Register Note

1. The host should only modify the Frame Manager RT_TOV/AL_TIME & ED_TOV register while the Frame Manager is in the OFFLINE State.

PTI 172584

5.10.5 Frame Manager Link Error Status Counters #1 Register

Register Address: 0x01D0
Reset Value: 0x00000000

Loss of Signal Count				Bad Rx Char Count				Loss of Sync Count				Link Fail Count				Read Only
31			27			24	23			19		16	15		11		8	7			3		0	

Bit(s)	Field Label	Field Name	Description
31..24	**Loss of Signal Count**	Loss of Signal Count	Indicates the number of times the Frame Manager detected a low to high transition on the lnk_unuse signal.
23..16	**Bad Rx Char Count**	Bad Received Character Count	Indicates the number of times the 8B/10B decode detected an invalid 10-bit code. FC-PH denotes this value as "Invalid Transmission Word during frame reception." This field may be non-zero after initialization. After initialization, the host should read this value to determine the correct starting value for this error count.
15..8	**Loss of Sync Count**	Loss of Sync Count	Indicates the number of times the loss of sync is greater than RT_TOV.
7..0	**Link Fail Count**	Link Fail Count	Indicates the number of times the Frame Manager detected a NOS or other initialization protocol failure that caused a transition to the Link Failure state.

Table 5.43 Frame Manager Link Error Status Counters #1 Register

Frame Manager Link Error Status Counters #1 Register Note

1. This register contains four error counters that Tachyon increments each time the associated error occurs. When these counters reach the maximum count (0xff), they rollover to 0x00 and continue counting. The host should poll these error counters every second or so. Since these errors should be infrequent, each counter is only 8 bits in length. The host must maintain longer counters if needed. A read of this register resets the counters to zero.

PTI 172585

5.10.6 Frame Manager Link Error Status Counters #2 Register

Register Address: 0x01D4
Reset Value: 0x00000000

Received EOFa	Generated EOFa	Bad CRC Count	Protocol Error Count	
31 ... 24	23 ... 16	15 ... 8	7 ... 0	Read Only

Bit(s)	Field Label	Field Name	Description
31..24	Received EOFa	Received EOFa	The number of frames containing an EOFa delimiter that Tachyon has received.
23..16	Generated EOFa	Generated EOFa	The number of problem frames that Tachyon has received that caused the Frame Manager to attach an EOFa delimiter. Frames that Tachyon discarded due to internal FIFO overflow are not included in this or any other statistic.
15..8	Bad CRC Count	Bad CRC Count	The number of bad CRC frames that Tachyon has received.
7..0	Protocol Error Count	Protocol Error Count	The number of protocol errors that the Frame Manager has detected.

Table 5.44 Frame Manager Link Error Status Counters #2 Register

Frame Manager Link Error Status Counters #2 Register Note

1. This register contains four error counters that Tachyon increments each time the associated error occurs. When these counters reach the maximum count (0xff), they rollover to 0x00 and continue counting. The host should poll these error counters every second or so. Since these errors should be infrequent, each counter is only 8 bits in length. The host must maintain longer counters if needed. When the host reads this register, Tachyon resets the counters to zero.

PTI 172586

5.10.7 Frame Manager World Wide Name Hi Register

Register Address: 0x01E0
Reset Value: 0x00000000

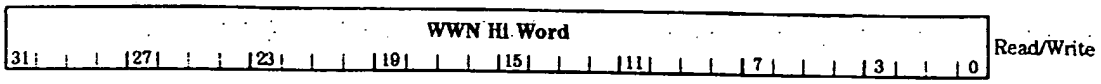

Bit(s)	Field Label	Field Name	Description
31..0	WWN Hi Word	World Wide Name High Word	The most significant four bytes of the WWN.

Table 5.45 Frame Manager World Wide Name Hi Register

Frame Manager World Wide Name Hi Register Note

1. The host uses this register to indicate the high order 4 bytes of the 8-byte unique World Wide Name (WWN) that Tachyon should use during initialization. Note that Tachyon does not specify which type of WWN to use. It simply copies all 8 bytes from the registers to the loop initialization frames. The host should only modify the Frame Manager World Wide Name Hi register while Frame Manager is in the Offline state.

5.10.8 Frame Manager World Wide Name Lo Register

Register Address: 0x01E4
Reset Value: 0x00000000

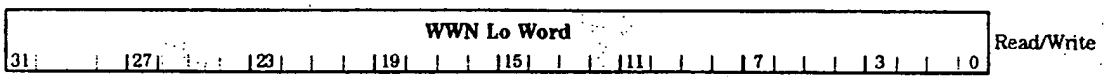

Bit(s)	Field Label	Field Name	Description
31..0	WWN Lo Word	World Wide Name Low Word	The least significant four bytes of the WWN.

Table 5.46 Frame Manager World Wide Name Lo Register

Frame Manager World Wide Name Lo Register Note

1. The host uses this register to indicate the low order 4 bytes of the 8-byte unique World Wide Name (WWN) that Tachyon should use during initialization. Note that Tachyon does not specify which type of WWN to use. It simply copies all 8 bytes from the registers to the loop initialization frames. The host should only modify the Frame Manager World Wide Name Lo register while Frame Manager is in the Offline state.

PTI 172587

5.10.9 Frame Manager Received AL_PA Register

Register Address: 0x01E8
Reset Value: 0x00000000

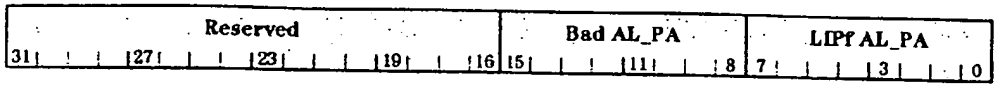

Bit(s)	Field Label	Field Name	Description
31..16	Reserved	Reserved	Value is undefined.
15..8	Bad AL_PA	Bad AL_PA	The AL_PA of an OPN primitive signal that was sent and returned without being accepted by the remote node.
7..0	LIPf AL_PA	LIPf AL_PA	The AL_PA of the most recent LIPf.

Table 5.47 Frame Manager Received AL_PA Register

Frame Manager Received AL_PA Register Note

1. This register contains the 8-bit AL_PA which was received on the last LIPf or as a bad AL_PA. The LIPf AL_PA is only valid when the lipf bit is set in the Frame Manager Status register. The Bad AL_PA is only valid when the ba bit is set in the Frame Manager Status register.

5.10.10 Frame Manager Primitive Register

Register Address: 0x01EC
Reset Value: 0x00000000

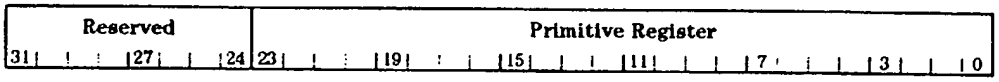

Bit(s)	Field Label	Field Name	Description
31..24	Reserved	Reserved	Initialize to zero.
23..0	Primitive Register	Primitive Register	The lower three bytes of the ordered set to be sent.

Table 5.48 Frame Manager Primitive Register

Frame Manager Primitive Register Note

1. The host writes to this register to tell Tachyon the lower 24 bits for the Primitive it should transmit in the host control state. Refer to "FC-AL" and "FC-PH" specifications for a list of ordered sets.

6. Data Structures

6.1 Boundary Alignment

The following table summarizes the boundary alignment of host data structures.

Host Data Structures	Boundary Alignment
OCQ	sizeof(OCQ)
OCQ Consumer Index Address	32 byte
HPCQ	sizeof(HPCQ)
HPCQ Consumer Index Address	32 byte
IMQ	sizeof(IMQ)
IMQ Producer Index Address	32 byte
SFSBQ	sizeof(SFSBQ)
SFS data buffer	sizeof(buffer)
MFSBQ	sizeof(MFSBQ)
MFS data buffer	sizeof(buffer)
Tachyon Header Structure	4 byte (32 byte is recommended)
High Priority Frame Structure	4 byte (32 byte is recommended)
EDB	32 byte
SEST	sizeof(SEST), which must be a power of 2 bytes
SDB - OOO Reassembly	256 byte
SDB - In Order Reassembly	power of 2 bytes, equal to or greater than the SDB length
SCSI Buffer Length - OOO Reassembly	sizeof(buffer)
SCSI Buffer Length - In Order Reassembly	4 byte

Table 6.1 Host Data Structure Boundary Alignment

PTI 172589

6.2 Circular Queues

6.2.1 OCQ

OCQ Entry (ODB)	Entry 0
OCQ Entry (ODB)	Entry 1
OCQ Entry (ODB)	...
OCQ Entry (ODB)	Entry n-1

Entry	Name	Description
0 .. n-1	OCQ Entry (ODB)	Each entry is a 32-byte ODB (Outbound Descriptor Block).

Table 6.2 OCQ

6.2.2 HPCQ

HPCQ Entry (HPDB)	Entry 0
HPCQ Entry (HPDB)	Entry 1
HPCQ Entry (HPDB)	...
HPCQ Entry (HPDB)	Entry n-1

Entry	Name	Description
0 .. n-1	HPCQ Entry (HPDB)	Each entry is a 32-byte HPDB (High Priority Descriptor Block).

Table 6.3 HPCQ

6.2.3 IMQ

IMQ Entry (Completion Message)	Entry 0
IMQ Entry (Completion Message)	Entry 1
IMQ Entry (Completion Message)	Entry 2
IMQ Entry (Completion Message)	Entry 3
IMQ Entry (Completion Message)	...
IMQ Entry (Completion Message)	Entry n-1

Entry	Name	Description
0 .. n-1	IMQ Entry (Completion Message)	Each entry is a 32-byte Completion Message.

Table 6.4 IMQ

PTI 172590

6.2.4 SFSBQ

SFSBQ Entry	
SFSBQ Entry	Entry 0
SFSBQ Entry	Entry 1
SFSBQ Entry	...
SFSBQ Entry	Entry n-1

Entry	Name	Description
0 .. n-1	SFSBQ Entry	Each entry is 32-bytes long and contains the addresses of eight SFS buffers in host memory.

Table 6.5 SFSBQ

6.2.5 MFSBQ

MFSBQ Entry	
MFSBQ Entry	Entry 0
MFSBQ Entry	Entry 1
MFSBQ Entry	...
MFSBQ Entry	Entry n-1

Entry	Name	Description
0 .. n-1	MFSBQ Entry	Each entry is 32-bytes long and contains the addresses of eight MFS buffers in host memory.

Table 6.6 MFSBQ

PTI 172591

6.3 Host-Based Data Structures

6.3.1 Tachyon Header Structure

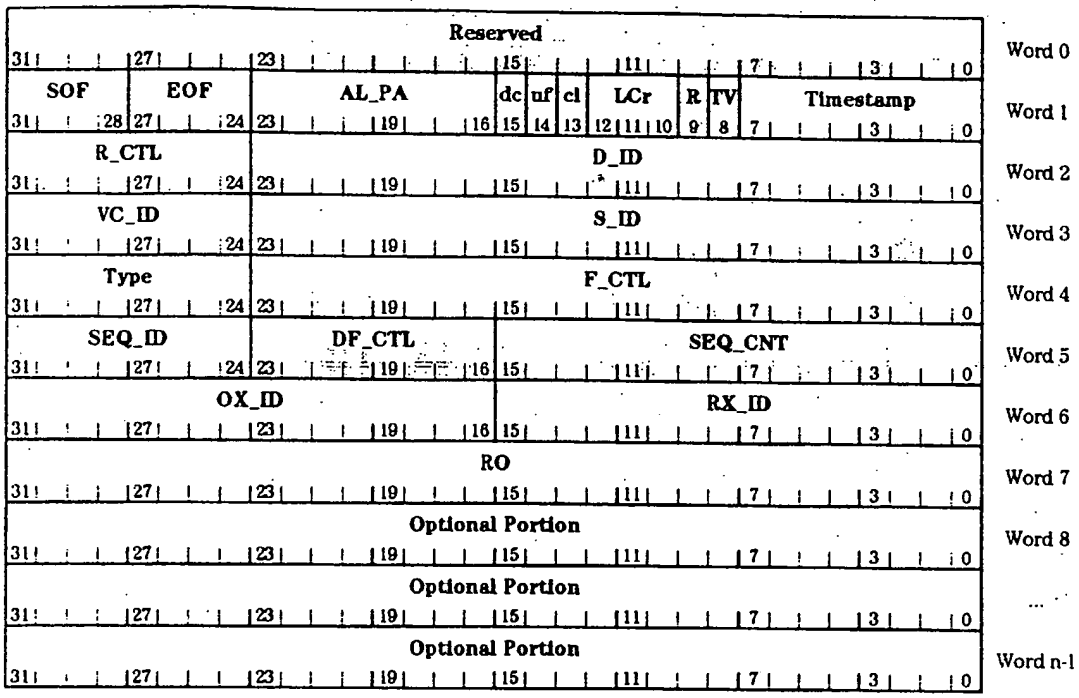

Word	Bit(s)	Field Label	Field Name	Description
0	31..0	Reserved	Reserved	Initialize to zero.
1	31..28	SOF	Initial Start of Frame	The host programs the SOF value for the first frame. Tachyon then modifies the SOF vlaue on all subsequent frames, using the values: SOFc1 = 3 SOFi1 = 5 SOFn1 = 9 SOFi2 = 6 SOFn2 = 10 SOFi3 = 7 SOFn3 = 11 If an invalid SOF is sent, a SOFf is substituted for it.

Table 6.7 Tachyon Header Structure

PTI 172592

Word	Bit(s)	Field Label	Field Name	Description
	27..24	EOF	Initial End of Frame	The host programs the EOF value for the first frame. Tachyon then modifies the EOF value for subsequent frames as necessary, using the values: EOFdt = 1 EOFa = 4 EOFn = 5 EOFt = 6 If an invalid EOF is sent, an EOFa is substituted for it. When sending single frame sequences, the host should program the EOF value to EOFn. Tachyon automatically updates the value to EOFt.
	23..16	AL_PA	FC-AL Physical Address	Address of the destination loop device. This value should be cleared to zero if the remote node is on the far side of a fabric (the D_ID value indicates the fabric address of the remote node). An AL_PA is required only for loop operation. If Tachyon is not on a loop, this field is ignored.
	15	dc	Disable CRC	For debugging purposes or when the host generates the CRC for the frame. Should be limited to High Priority frame use only.
	14	uf	Unfair Access	On a loop, the unfair access rules allow earliest transmission of a sequence of frames. Required only for loop operation. If Tachyon is not on a loop, this bit is ignored.
	13	cl	Loop Close	When this bit is set to one, it indicates that Tachyon closes the loop after each frame is sent, unless frames destined for the same AL_PA are still in the Data or ACK FIFOs. After these frames are sent, Tachyon closes the loop. Required only for loop operation. If Tachyon is not on a loop, this bit is ignored.
	12..10	LCr	Loop Credit	The BB_Credit given to the host to use during login to send a sequence. Required only for loop operation. If Tachyon is not on a loop, this field is ignored.
	9	R	Reserved	Initialize to zero.
	8	TV	Timestamp Valid	This flag indicates whether the Timestamp (bits 7..0) contain valid data.
	7..0	Timestamp	Timestamp	This field contains the timestamp used to determine when a frame is older than ED_TOV.
2	31..24	R_CTL	R_CTL	These five fields are the Fibre Channel Header values to use for all frames of this sequence.
	23..0	D_ID	D_ID	

Table 6.7 Tachyon Header Structure (Continued)

PTI 172593

Word	Bit(s)	Field Label	Field Name	Description
3	31..24	VC_ID	VC_ID	Note that the D_ID value of the Header should match the D_ID value of the Outbound Descriptor Block.
	23..0	S_ID	S_ID	
4	31..24	Type	Type	
	23..0	F_CTL	Initial F_CTL	The Header field used on all frames except the last, which Tachyon modifies to values given in the ODB. For non-word multiple outbound SEST transfers, the Fill bits in the F_CTL field are set automatically.
5	31..24	SEQ_ID	Sequence ID	Used for all frames of this sequence. This value should match the Sequence ID in the ODB.
	23..16	DF_CTL	Initial DF_CTL	Tachyon modifies this field to zero on subsequent frames.
	15..0	SEQ_CNT	Initial Sequence Count	Tachyon send the first frame with this value and then increments the value for each subsequent frame of the sequence. Tachyon passes the final Sequence Count back to the host in a completion message.
6	31..16	OX_ID	OX_ID	The OX_ID for this exchange. Selected by the host initiating the exchange and assigned in the first frame of an exchange. Refer to " Choosing X_ID Values" on page 160.
	15..0	RX_ID	RX_ID	The RX_ID for this exchange. Selected by the host responding to the exchange. The initiating host must set this value to 0xFFFF in the first frame of the exchange. The responder host's selected RX_ID is returned in the ACK frame sent in response to the last data frame of a sequence. Refer to " Choosing X_ID Values" on page 160.
7	31..0	RO	Initial Relative Offset	Tachyon tracks and fills in the proper RO as each frame is received.
8 .. n-1	31..0	Optional Portion	Optional Portion	Various optional Fibre Channel Headers and a small data payload may be placed in the same buffer as the Fibre Channel Header. The optional headers (OHDRs) and small data payload must be less than or equal to Tachyon's Frame_Len value. Tachyon only sends the ODHRs and a small data payload on the first frame of the sequence.

Table 6.7 Tachyon Header Structure (Continued)

PTI 172594

Tachyon Header Structure Notes

1. All normal outbound and inbound frames use the above format for the Tachyon Header Structure. The Tachyon Header Structure is a copy of the Fibre Channel Header which resides in host memory.
2. For a normal outbound sequence, Tachyon reads the 32 bytes of the FC Header and stores them internally for use in creating Tachyon Header Structures for subsequent frames. Tachyon then updates several fields during the course of sending the sequence.
   a. Tachyon clears the DF_CTL field to zero on all frames except the first frame.
   b. Tachyon sets the E_C bit of the F_CTL field to one on the last frame of the sequence if the host set the E_C bit to one in the ODB. This does not occur if the host has set the CONT_SEQ bit in the ODB to one.
   c. Tachyon sets the fill bits of the F_CTL field of the last frame to the values in the ODB. (Tachyon does not do this if the host has set the CONT_SEQ bit in the ODB to one.)
   d. Tachyon sets the EOS bit of the F_CTL field to one on a single frame sequence or on the last frame of the multiframe sequence. Tachyon does not do this if the host has set the CONT_SEQ bit in the ODB to one.
   e. Tachyon sets EOFt on the last frame of a Class 3 sequence. Tachyon does not do this if the host has set the CONT_SEQ bit in the ODB to one.
   f. Tachyon sets SOFxn on all frames of the sequence except the first frame.
3. All fields in the Tachyon Header Structure and the ODB that have the same name must have the same values.

PTI 172595

Choosing X_ID Values

The initiator host of the exchange must assign an OX_ID in the Tachyon Header Structure. As an initiator, Tachyon uses the OX_ID as an index into the SEST for a SCSI transaction. The target host must assign an RX_ID in the Tachyon Header Structure. As a target, Tachyon uses the RX_ID as an index into the SEST for SCSI transactions. The host initiator or target must use the following format when choosing X_ID values.

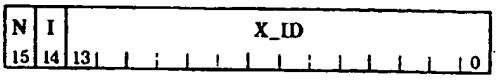

Bit(s)	Field Label	Field Name	Description
15	N	Non-SCSI Bit	The host sets this bit to one if the X_ID being chosen is for a non-SCSI assisted transaction. For a SCSI-assisted transaction, this bit is cleared to zero.
14	I	Inbound Bit	For a SCSI-assisted transaction, this bit is used to verify that the sequence in progress corresponds to the direction programmed in the corresponding SEST entry. For an Inbound SEST Entry, this bit is set to one. For an Outbound SEST entry, this bit is cleared to zero. For a non-SCSI-assisted transaction (bit 15, above, is set to one), this bit is ignored.
13..0	X_ID	X_ID	The exchange ID selected for this exchange. For SCSI, this value must be the index of the SEST Entry for the SCSI transaction. The initiator host must assign an OX_ID. Tachyon uses the OX_ID as an index into the SEST. The target host must assign a RX_ID. Tachyon uses the RX_ID as an index into the SEST.

Table 6.8 Choosing X_ID Values

PTI 172596

6.3.2 Outbound Descriptor Block (ODB)

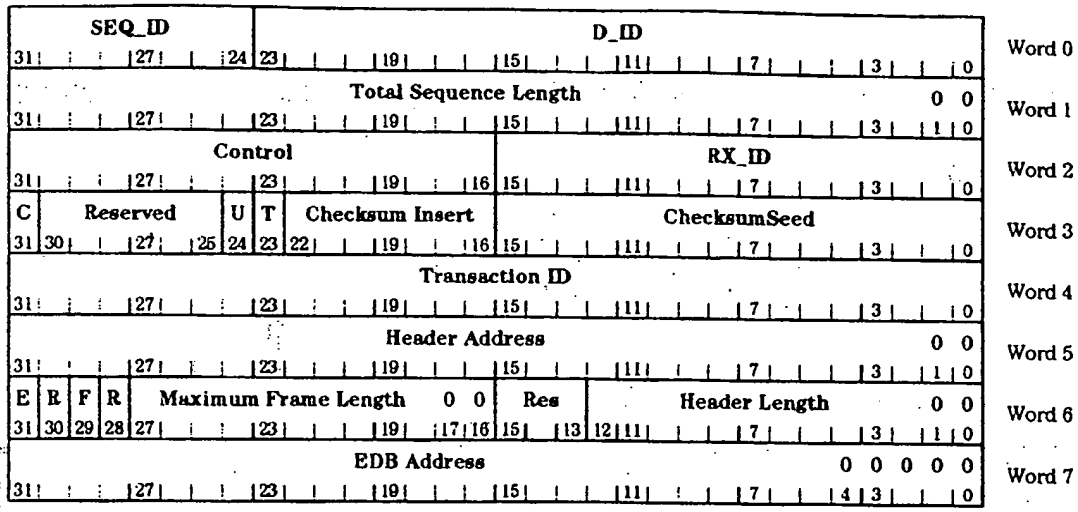

Word	Bit(s)	Field Label	Field Name	Description
0	31..24	SEQ_ID	SEQ_ID	The unique FC sequence identifier for this transfer. This value should match the SEQ_ID value in the Tachyon Header Structure.
	23..0	D_ID	D_ID	The FC address of the destination node.
1	31..0	Total Sequence Length	Total Sequence Length	The length in bytes of the data to send. This length does not include the length of the Tachyon Header Structure, but does include any optional headers. Use the Fill bytes in the Control field of the ODB to determine the Total Sequence Length: Total Sequence Length = Fill bytes + exact length. The Total Sequence Length must be equal to or greater than [Header Length - 32 bytes] and must be aligned on a 4-byte boundary.
2		Control	Control	Various chip control bits for this sequence are contained in this field, as follows.
	31..30	Cl	Class	The FC class of service to use. 1 = Class 1 2 = Class 2 3 = Class 3 0 = Undefined

Table 6.9 Outbound Descriptor Block

PTI 172597

Word	Bit(s)	Field Label	Field Name	Description
	29	Lck	Sequence Interlock	Set to one when Tachyon is to wait for an ACK after transmitting the first frame of a sequence before proceeding with the rest of the sequence. This bit is used to enable retries of the first frame of the sequence and to provide the X_ID Interlock Protocol.
	28	SOFc1	Start Class 1 Connection	Set to one to create a Class 1 connection with the beginning of this sequence. The Tachyon Header Structure, as described in this ODB, is sent as the connect request.
	27	E_C	End Connection	Set to one to terminate the connection at the end of this sequence.
	26	No Comp	Completion Message Disable	Set to one if the host does not want the Outbound Sequence Manager to generate a completion message for this sequence.
	25	No Int	Completion Message Interrupt Disable	Set to one if the host does not want an interrupt when the completion message is posted back to the host.
	24	ACK_0	ACK_0	Use the ACK_0 model for transmitting the sequence. Send all frames without waiting for individual frame ACKs, but wait for a final ACK for the entire sequence. Note that a sequence timeout occurs if the ACK_0 bit and the CONT_SEQ bit are both set.
	23..22	Fill	Fill Bytes	The F_CTL fill byte field set in the last frame of the sequence. The value of this field indicates how many pad bytes (0, 1, 2, or 3) are to be inserted to make the frame length a multiple of four bytes.
	21	Reserved	Reserved	Initialize to zero.
	20	CONT_SEQ	Continue Sequence	Set to one when the end of sequence is not to be indicated in the last frame. A sequence timeout occurs if the CONT_SEQ bit and the ACK_0 bit are both set.
	19..16	EE_Credit	EE_Credit	The EE_Credit, given to the host during login, to use to send the sequence. The maximum EE_Credit is 15.
	15..0	RX_ID	RX_ID	The responder exchange ID to use for this sequence.

Table 6.9 Outbound Descriptor Block (Continued)

PTI 172598

Word	Bit(s)	Field Label	Field Name	Description
3	31	C	Checksum Enable	Set this bit to enable hardware assisted checksumming. Checksum Enable is enabled only if TCP/UDP assists are also enabled, i.e., bit 31 in the Tachyon Configuration register is set to one.
	30..25	Reserved	Reserved	Initialize to zero.
	24	U	UDP Checksum	Set this bit to enable UDP Checksum. UDP Checksum can only be enabled if TCP/UDP assists are also enabled, i.e., bit 31 in the Tachyon Configuration register is set.
	23	T	Trailer Checksum	Enable Trailers where the checksum overwrites the last two bytes of data DMAed from the host. The checksum does not include the value of the last two bytes before overwriting, nor does it include any fill bytes. Trailer Checksum can only be enabled if TCP/UDP assists are also enabled, i.e., bit 31 in the Tachyon Configuration register is set.
	22..16	Checksum Insert	Checksum Insert	The Relative Offset location within the sequence where the checksum value is to be placed. The Relative Offset location must be on a two byte boundary. Checksum Insert is used only if TCP/UDP assists are enabled. Note that this field is ignored if doing trailer checksumming.
	15..0	Checksum Seed	Checksum Seed	The value the checksum accumulator is to start with when calculating sequence checksums. The Checksum Seed is only used if TCP/UDP assists are enabled.
4	31..0	Transaction ID	Transaction ID	This value is used by the driver to match completions to requests. Tachyon returns this value unmodified on completions.
5	31..0	Header Address	Header Address	The address in host memory of the Tachyon Header Structure and Optional Headers, if they exist. The address must be aligned on a 4-byte boundary.
6	31	E	End	Set this bit if the header contains the entire sequence.
	30	R	Reserved	Initialize to zero.
	29	F	Frame Boundary	Set this bit when the host wants an EOF delimiter sent after the data pointed to by the Tachyon Header Structure A/L pair.
	28	R	Reserved	Initialize to zero.

Table 6.9 Outbound Descriptor Block (Continued)

PTI 172599

Word	Bit(s)	Field Label	Field Name	Description
	27..16	Maximum Frame Length	Maximum Frame Length	The maximum payload length per frame. Refer to the login parameters on p. 109. The Receive Data Field Size login paramenter defines the minimum value of this field as 128 bytes. Since Tachyon only supports an outbound frame payload of 2048 bytes, including any optional headers, the maximum value of this field is 2048 bytes. The length must be aligned on a 4-byte boundary.
	15..13	Res	Reserved	Initialize to zero.
	12..0	Header Length	Header Buffer Length	This Header Buffer Length is set to the Tachyon Header Structure (which is 32 bytes) + the optional portion (this includes optional headers and data payload). The Hdr_len must be less than or equal to [32 bytes + Maximum Frame Length] or to 2080 bytes, whichever is less. For SOFc1, the Header Length must not exceed the (Class 2 Login Receive Data field size + 32 bytes) or 2080 bytes, whichever is less. Refer to the login parameters on p. 109. The length must be aligned on a 4-byte boundary.
7	31..0	EDB Address	Extended Descriptor Block Address	The address of the EDB that contains the addresses of the data buffers to be sent. The EDB Address must be 32-byte aligned.

Table 6.9 Outbound Descriptor Block (Continued)

WARNING For Tachyon to operate properly, do not set the Frame Boundary Bit (F, Word 6, bit 29) if the entire sequence is contained in the Header Buffer A/L.

ODB Notes

1. The ODB is a 32-byte structure that defines the Tachyon Header Structure, optional headers, checksum information, pointers to the data, etc.
2. All fields in the Tachyon Header Structure and the ODB that have the same name must have the same values.

PTI 172600

6.3.3 Extended Descriptor Block (EDB)

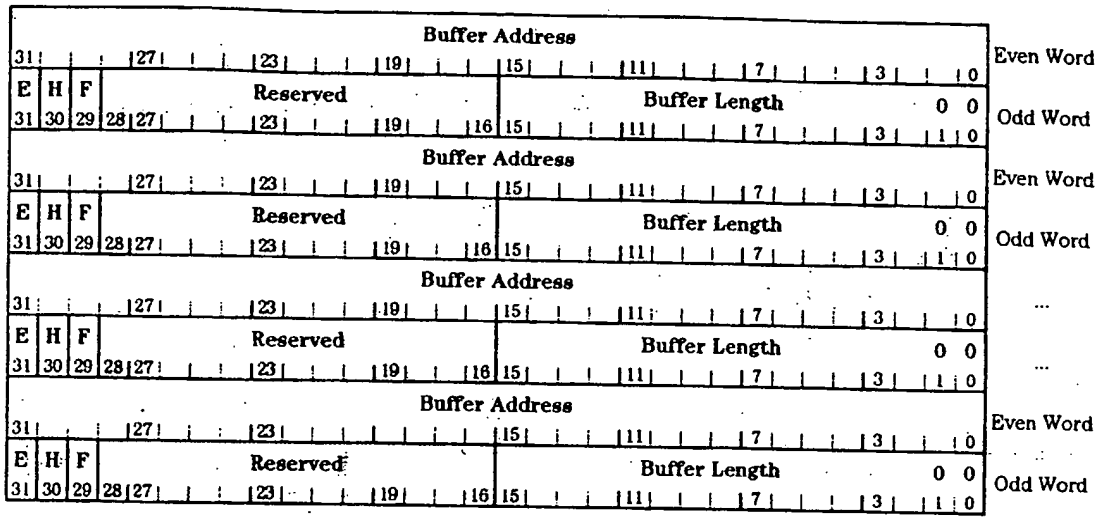

Word	Bit(s)	Field Label	Field Name	Description
Even	31..0	Buffer Address	Buffer Address	The address of a buffer in host memory aligned on a 4-byte boundary.
Odd	31	E	End	This bit is set in the last A/L pair of the EDB.
	30	H	Header	This bit is set in the first A/L pair when a second A/L pair is used to point to additional network header information. The Header Bit should only be set if the Checksum Enable bit in the ODB is set. This bit is not used if the entire network header is pointed to by one A/L pair.
	29	F	Frame Boundary	This bit is set when the host wants an EOF delimter sent after the block pointed to by the current A/L pair.
	28..16	Reserved	Reserved	Initialize to zero.
	15..0	Buffer Length	Buffer Length	The number of valid bytes in the buffer. The Buffer Length must be a multiple of 4 bytes.

Table 6.10 Extended Descriptor Block (EDB)

PTI 172601

EDB Notes

1. The EDB defines the data payload of a FC sequence.
2. The EDB must be 32-byte aligned.
3. Each entry within the EDB consists of an Address/Length pair (8 bytes).
4. Each A/L pair describes a buffer in host memory that contains data that is to be transmitted.
5. A/L pairs in the EDB are always fetched in blocks of 4 pairs (32 bytes).
6. The EDB contains as many A/L pairs as necessary to completely define the sequence.
7. If the number of A/L pairs in the EDB is not a multiple of 4, then the memory locations between the last A/L pair and the next 32 byte boundary is fetched and must be readable.
8. There is no facility for extending the length of the EDB to another area in memory, therefore the complete EDB must reside in a physically contiguous space.
9. If Tachyon has determined that the remaining data to be sent (Tot_Len) is less than or equal to Frame_Len, any subsequent A/L pairs with the Frame Boundary (F) bit set are ignored. This does not apply if outbound header checksumming is enabled.
10. If outbound header checksumming is enabled, the Tot_Len field in the ODB must equal the sum of the optional headers and all of the A/L pairs in the EDB, including the last A/L pair which must have the END bit set.
11. The EDB must be allocated in 32-byte chunks regardless of whether all the entries are used. The buffer length (Buf_Len) may range from 0x0004 to 0xFFFC. Because only word transfers are allowed, the largest practical buffer length is 32 Kbytes (0x8000 Kbytes) for each A/L pair, if the host wishes to keep data aligned on power of 2 byte boundaries. This is not a requirement.

6.3.4 High Priority Frame Structure

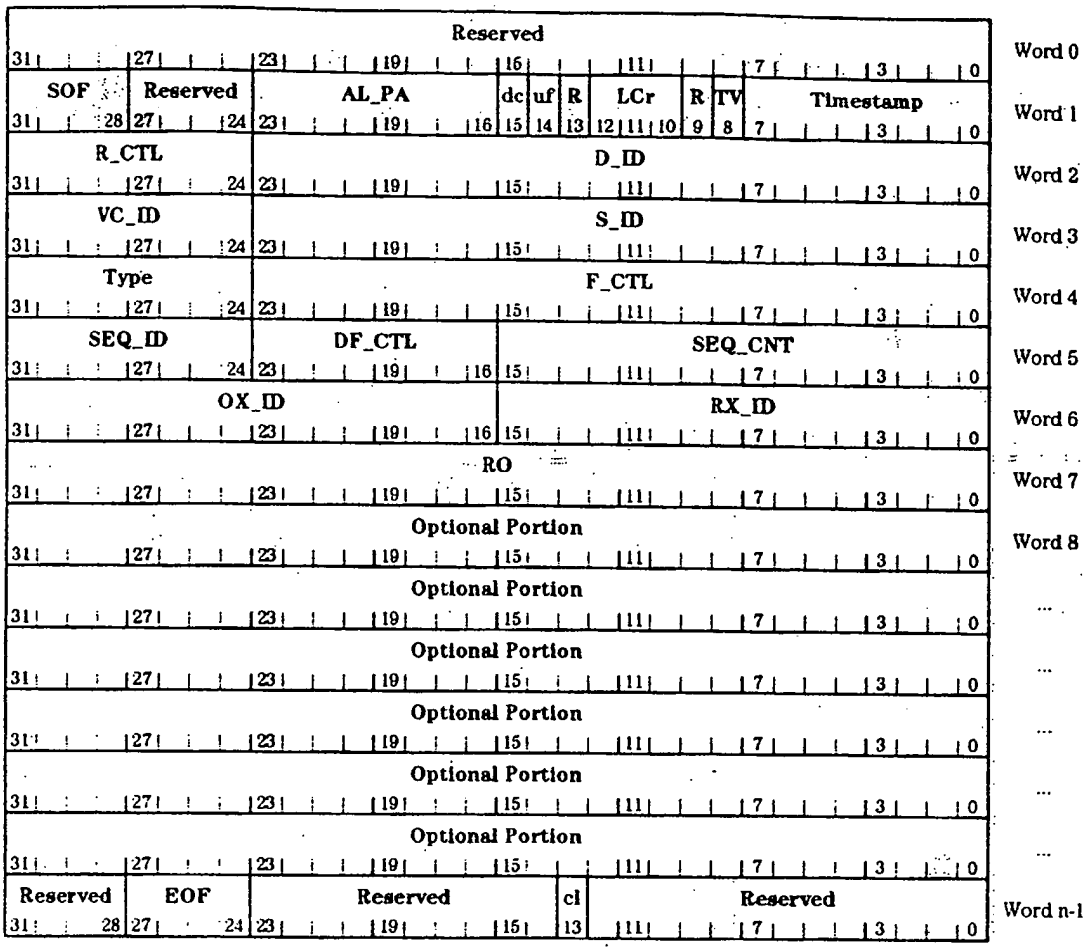

Table 6.11 High Priority Frame Structure

High Priority Frame Structure Notes

1. The descriptions of the fields for the High Priority Frame Structure are identical to the descriptions of fields in the Tachyon Header Structure.
2. Unlike the Tachyon Header Structure for the normal outbound sequence, Tachyon does not overwrite the High Priority Frame Structure with any information from the ODB. Therefore, it is important that the host ensures that all fields in a High Priority Frame Structure are accurate.
3. For data frames that Tachyon does not automatically ACK, the host must generate the ACK. The host must use the TV and timestamp values in the data frames that were received from Tachyon. The host must provide the same timestamp to allow the outbound timestamp verifier to prevent a response from being sent after ED_TOV since the data frame was received.
4. The Frame_Len field in the HPDB includes the 32 bytes of the header, the data (plus fill bytes, if any), and the EOF word. Units are in bytes. The length must be word-aligned (bits 1..0 = 0).
5. For Class 3, the EOF field should be set to EOFt.

PTI 172603

6.3.5 High Priority Descriptor Block (HPDB)

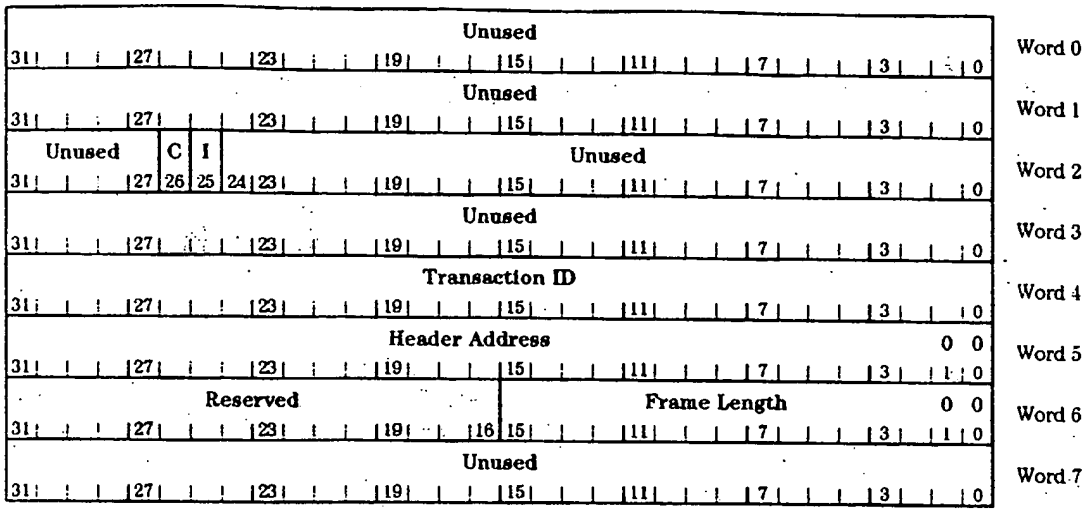

Word	Bit(s)	Field Label	Field Name	Description
0	31..0	Unused	Unused	Initialize to zero.
1	31..0	Unused	Unused	Initialize to zero.
2	31..27	Unused	Unused	Initialize to zero.
	26	C	Completion Message Disable	Set this bit to one when the host does not want the OSM to generate a completion mesage for this sequence.
	25	I	Completion Message Interrupt Disable	Set this bit to one when the host does not want an interrupt when the completion is posted back to the host.
	24..0	Unused	Unused	Initialize to zero.
3	31..0	Unused	Unused	Initialize to zero.
4	31..0	Transaction ID	Transaction ID	This identifier is used by the driver to match completions to requests. Tachyon returns this value unmodified on completions.
5	31..0	Header Address	Header Address	The address in host memory of the Fibre Channel frame to send. This address must by aligned on a 4-byte boundary.
6	31..16	Reserved	Reserved	Initialize to zero.

Table 6.12 High Priority Descriptor Block

PTI 172604

Word	Bit(s)	Field Label	Field Name	Description
	15..0	Frame Length	Frame Length	The length in bytes of the Fibre Channel frame to send. The maximum value of this field is 2084 bytes. This includes the Tachyon Header Structure (32 bytes) + the payload, any optional headers, and any Fill bytes from the F_CTL field of the Tachyon Header Structure (up to 2048 bytes) + the EOF (4 bytes). The length must be aligned on a 4-byte boundary.
7.	31..0	Unused	Unused	Initialize to zero.

Table 6.12 High Priority Descriptor Block (Continued)

HPDB Notes

1. Since frames leaving the High Priority Message Channel are managed differently from normal outbound frames, the amount of information which needs to be included in an ODB for a high priority frame is only a subset of a normal ODB.
2. Although the structure of the ODB for either type of message is the same, the host may elect not to fill in ODB fields which are unused for the High Priority Messsage Channel.

6.3.6 Inbound Read Status Frame Structure

For debugging purposes, the host can set the Status Request bit in the Tachyon Control register to one to request that Tachyon dump its internal state into an SFS buffer. When Tachyon detects that this bit is set, it temporarily suspends all other operations while sending the status information to the host. Tachyon DMAs the contents of various registers in the outbound and inbound paths to the host and then sends a read_status completion message and interrupt to the host.

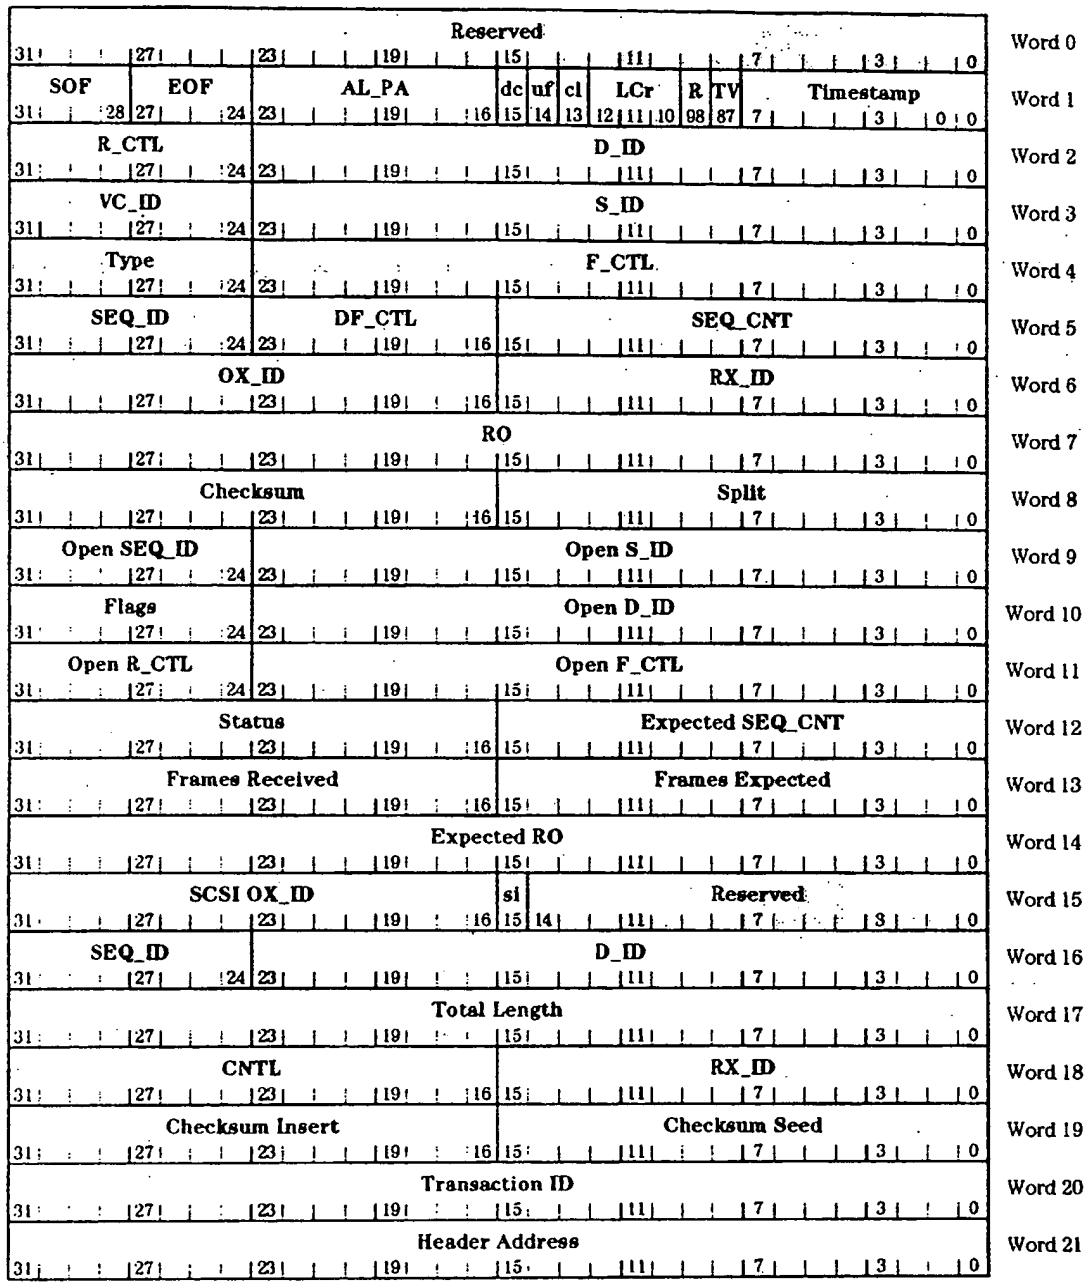

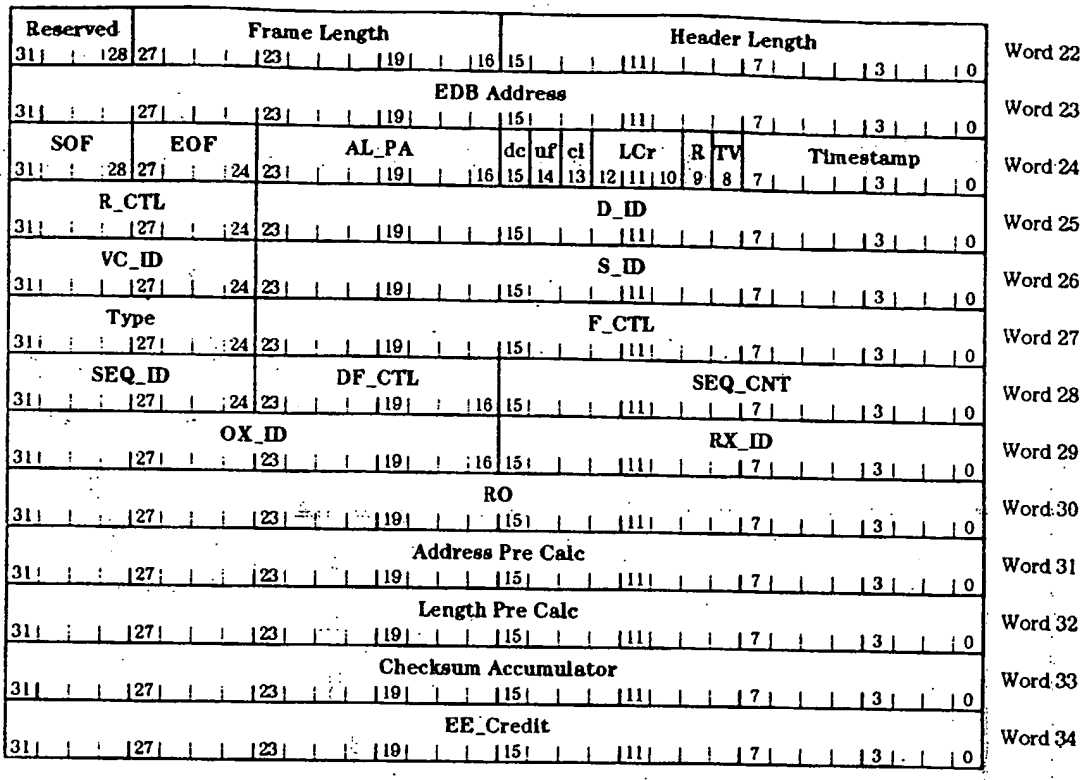

Word	Bit(s)	Field Label	Field Name	Description
0	31..0	Reserved	Reserved	Initialize to zero.
1	31..28	SOF	Start of Frame	The SOF value in the last Tachyon Header Structure received.
	27..24	EOF	End of Frame	The EOF value in the last Tachyon Header Structure received.
	23..16	AL_PA	FC-AL Physical Address	The AL_PA value in the last Tachyon Header Structure received.
	15	dc	Disable CRC	A copy of the Disable CRC bit from the last Tachyon Header Structure received.
	14	uf	Unfair Access	A copy of the Unfair Access bit from the last Tachyon Header Structure received.
	13	cl	Loop Close	A copy of the Loop Close bit from the last Tachyon Header Structure received.
	12..10	LCr	Loop Credit	A copy of this value from the last Tachyon Header Structure received.
	9	R	Reserved	Initialize to zero.

Table 6.13 Inbound Read Status Frame

PTI 172607

Word	Bit(s)	Field Label	Field Name	Description
	8	TV	Timestamp Valid	A copy of this bit from the last Tachyon Header Structure received.
	7..0	Timestamp	Timestamp	A copy of this value from the last Tachyon Header Structure received.
2	31..24	R_CTL	R_CTL	A copy of this value from the last Tachyon Header Structure received.
	23..0	D_ID	Destination ID	A copy of this value from the last Tachyon Header Structure received.
3	31..24	VC_ID	VC_ID	A copy of this value from the last Tachyon Header Structure received.
	23..0	S_ID	Source ID	A copy of this value from the last Tachyon Header Structure received.
4	31..24	Type	Type	A copy of this value from the last Tachyon Header Structure received.
	23..0	F_CTL	F_CTL	A copy of this value from the last Tachyon Header Structure received.
5	31..24	SEQ_ID	Sequence ID	A copy of this value from the last Tachyon Header Structure received.
	23..16	DF_CTL	DF_CTL	A copy of this value from the last Tachyon Header Structure received.
	15..0	SEQ_CNT	Sequence Count	A copy of this value from the last Tachyon Header Structure received.
6	31..16	OX_ID	OX_ID	A copy of this value from the last Tachyon Header Structure received.
	15..0	RX_ID	RX_ID	A copy of this value from the last Tachyon Header Structure received.
7	31..0	RO	Relative Offset	A copy of this value from the last Tachyon Header Structure received.
8	31..16	Checksum	Checksum	A copy of the IP checksum for the last frame received.
	15..0	Split	Split	A copy of the header/data split offset for the last frame received.
9	31..24	Open SEQ_ID	Open SEQ_ID	The value from the current MFS inbound sequence.
	23..0	Open S_ID	Open S_ID	The value from the current MFS inbound sequence or, if a Class 1 connection is open, the ID of the remote node.
10		Flags	Flags	This field comprises the following subfields.

Table 6.13 Inbound Read Status Frame (Continued)

Word	Bit(s)	Field Label	Field Name	Description
	31	ib_cl_op	Inbound Class 1 Open	Indicates that an inbound Class 1 connection is open.
	30	ib_mfs_op	Inbound MFS Open	Indicates that an inbound MFS reassembly is in progress.
	29	ob_cl_op	Outbound Class 1 Open	Indicates that an outbound Class 1 connection is open.
	28	n_ff_rcvd	Networking First Frame Received	Indicates that the first frame of an MFS was received.
	27	n_lf_rcvd	Networking Last Frame Received	Indicates that the last frame of an MFS was received.
	26..24	Reserved	Reserved	Initialize to zero.
	23..0	**Open D_ID**	Open D_ID	The value from the current MFS inbound sequence or, if a Class 1 connection is open, the ID of the local node.
11	31..24	**Open R_CTL**	Open R_CTL	The value from the current MFS inbound sequence.
	23..0	**Open F_CTL**	Open F_CTL	The value from the current MFS inbound sequence.
12	31..16	**Status**	Completion Status	The value from the current MFS inbound sequence.
	15..0	**Expected SEQ_CNT**	Expected SEQ_CNT	The sequence count expected in the next frame.
13	31..16	**Frames Received**	Number of Frames Received	The number of frames received in the open MFS sequence.
	15..0	**Frames Expected**	Number of Frames Expected	The number of frames expected in the open MFS sequence.
14	31..0	**Expected RO**	Expected RO	The Relative Offset expected in the next frame.
15	31..16	**SCSI OX_ID**	SCSI OX_ID	The currently active OX_ID.
	15	**si**	Stored OX_ID Invalid	Set this bit if the SCSI OX_ID stored in the ISM is invalid.
	14..0	**Reserved**	Reserved	Initialize to zero.
16	31..24	**SEQ_ID**	SEQ_ID	The value of this field from the current ODB.
	23..0	**D_ID**	D_ID	The value of this field from the current ODB.

Table 6.13 Inbound Read Status Frame (Continued)

PTI 172609

Word	Bit(s)	Field Label	Field Name	Description
17	31..0	Total Length	Total Length	The number of bytes left to send in the sequence after Tachyon requests the next A/L. Total Length is decremented by the value of the next A/L that is requested by Tachyon.
18	31..16	CNTL	CNTL	The value of this field from the current ODB.
	15..0	RX_ID	RX_ID	The value of this field from the current ODB.
19	31..16	Checksum Insert	Checksum Insert	The value of this field from the current ODB.
	15..0	Checksum Seed	Checksum Seed	The value of this field from the current ODB.
20	31..0	Transaction ID	Transaction ID	The value of this field from the current ODB.
21	31..0	Header Address	Header Address	The value of this field from the current ODB.
22	31..28	Reserved	Reserved	Initialize to zero.
	27..16	Frame Length	Frame Length	The value of this field from the current ODB.
	15..0	Header Length	Header Length	The value of this field from the current ODB.
23	31..0	EDB Address	EDB Address	The address of the next A/L pair within the current EDB that will be requested by Tachyon. Note that this address may never be fetched if the End Bit was set in a previous A/L pair withing this EDB.
24	31..28	SOF	Start of Frame	A copy of this value from the last Tachyon Header Structure sent.
	27..24	EOF	End of Frame	A copy of this value from the last Tachyon Header Structure sent.
	23..16	AL_PA	Arbitrated Loop Physical Address	A copy of this value from the last Tachyon Header Structure sent.
	15	dc	Disable CRC	A copy of this bit from the last Tachyon Header Structure sent.
	14	uf	Unfair Access	A copy of this bit from the last Tachyon Header Structure sent.
	13	cl	Close Loop	A copy of this bit from the last Tachyon Header Structure sent.
	12..10	LCr	Loop Credit	A copy of this value from the last Tachyon Header Structure sent.
	9	Reserved	Reserved	Initialize to zero.
	8	TV	Timestamp Valid	A copy of this bit from the last Tachyon Header Structure sent.

Table 6.13 Inbound Read Status Frame (Continued)

PTI 172610

Word	Bit(s)	Field Label	Field Name	Description
	7..0	Timestamp	Timestamp	A copy of this value from the last Tachyon Header Structure sent.
25	31..24	R_CTL	R_CTL	A copy of this value from the last Tachyon Header Structure sent.
	23..0	D_ID	D_ID	A copy of this value from the last Tachyon Header Structure sent.
26	31..24	VC_ID	VC_ID	A copy of this value from the last Tachyon Header Structure sent.
	23..0	S_ID	S_ID	A copy of this value from the last Tachyon Header Structure sent.
27	31..24	Type	Type	A copy of this value from the last Tachyon Header Structure sent.
	23..0	F_CTL	F_CTL	A copy of this value from the last Tachyon Header Structure sent.
28	31..24	SEQ_ID	SEQ_ID	A copy of this value from the last Tachyon Header Structure sent.
	23..16	DF_CTL	DF_CTL	A copy of this value from the last Tachyon Header Structure sent.
	15..0	SEQ_CNT	SEQ_CNT	A copy of this value from the last Tachyon Header Structure sent.
29	31..16	OX_ID	OX_ID	A copy of this value from the last Tachyon Header Structure sent.
	15..0	RX_ID	RX_ID	A copy of this value from the last Tachyon Header Structure sent.
30	31..0	RO	RO	A copy of this value from the last Tachyon Header Structure sent.
31	31..0	Address Pre Calc	Address Pre Calc	The address of the next data to fetch.
32	31..0	Length Pre Calc	Length Pre Calc	The length of the next data to fetch.
33	31..0	Checksum Accumulator	Checksum Accumulator	A snapshot of the current checksum accumulator value.
34	31..0	EE_Credit	EE_Credit	A snapshot of the current EE_Credit value.

Table 6.13 Inbound Read Status Frame (Continued)

PTI 172611

6.4 Completion Messages

6.4.1 Inbound Message Queue Entry (Completion Message)

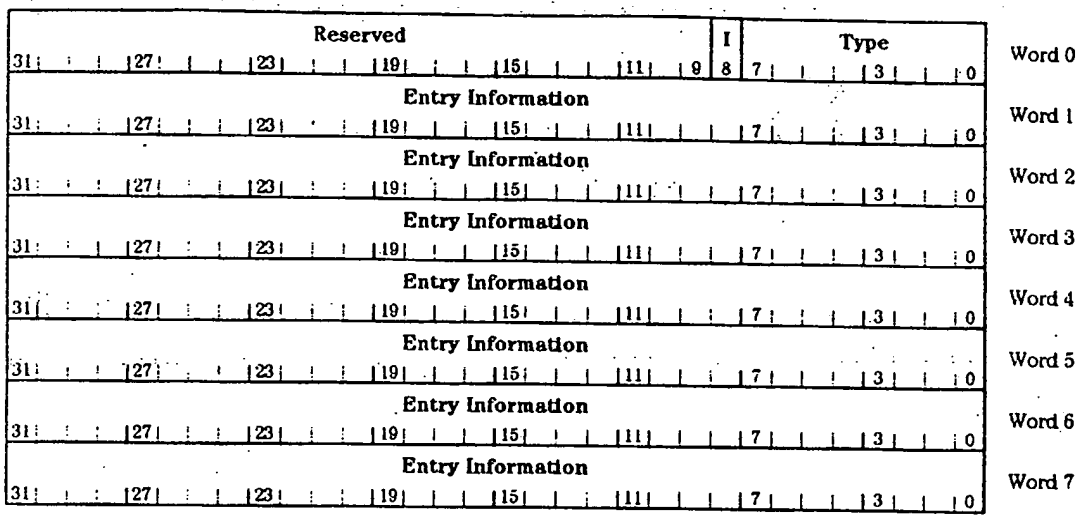

Word	Bit(s)	Field Label	Field Name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Desired	Set to one when an interrupt is generated for the completion message.
	7..0	Interrupt Type	Interrupt Type	Indicates the completion message type. Refer to the "Completion Message Types" table on the following page.
1..7	31..0	Entry Information	Entry Information	Additional information depending on the type of completion message.

Table 6.14 Inbound Message Queue Entry

Inbound Message Queue Entry Note

1. Even though completion messages vary in length from one to six words, they are always written to the host with WRITE8 transactions. Refer to "7.6.2 TSI Transaction Types" on page 207. Pad Words are added to the end of all completion messages so that their length is 8 words (or 32 bytes). The values of these Pad Words are undefined.

PTI 172612

6.4.2 Completion Message Types

Type	Page Number	IMQ Entry Word 0		
		Bits 31..9	Bit 8	Bits 7..0
outbound_completion	178	Reserved	0	0x00
outbound_completion_i	178	Reserved	1	0x00
out_hi_pri_completion	181	Reserved	0	0x01
out_hi_pri_completion_i	181	Reserved	1	0x01
inbound_mfs_completion	182	Reserved	1	0x02
inbound_ooo_completion	182	Reserved	0	0x03
inbound_sfs_completion	184	Reserved	1	0x04
inbound_c1_timeout	187	Reserved	1	0x05
inbound_busied_frame	184	Reserved	0	0x06
inbound_unknown_frame_i	184	Reserved	1	0x06
sfs_buf_warn	187	Reserved	1	0x07
mfs_buf_warn	187	Reserved	1	0x08
imq_buf_warn	187	Reserved	1	0x09
frame_mgr_interrupt	187	Reserved	1	0x0A
read_status	184	Reserved	1	0x0B
inbound_scsi_data_completion	186	Reserved	1	0x0C
inbound_scsi_command	184	Reserved	1	0x0D
bad_scsi_frame	184	Reserved	1	0x0E
inb_scsi_status_completion	184	Reserved	1	0x0F

Table 6.15 Completion Message Types

PTI 172613

6.4.3 Outbound Completion Message

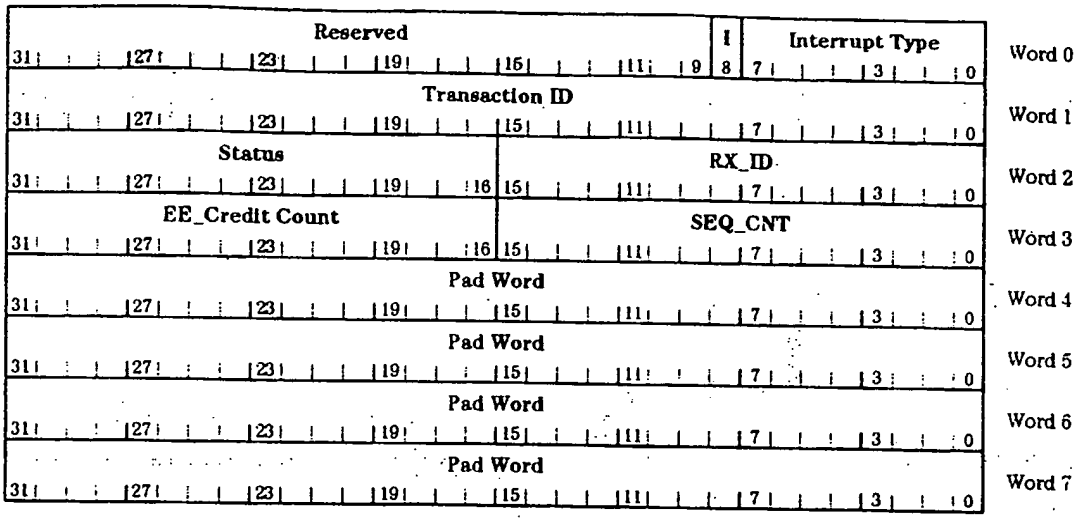

Word	Bit(s)	Field Label	Field name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Desired	Set to one when the host requests that an interrupt be generated for this completion message.
	7..0	Interrupt Type	Interrupt Type	outbound_completion = 0x00
1	31..0	Transaction ID	Transaction ID	Copy of the TRANS_ID in the ODB or Outbound SEST Entry for this sequence.
2		Status	Completion Status	This field comprises the following subfields.
	31	C1	Class 1 Error	The Class 1 connection experienced an error during sequence transmission. Class 1 errors occur with these combinations of conditions: 1. SOFc1 is set to one and the outbound C1 connection is already opened. 2. SOFc1 is cleared to zero and the outbound C1 connection is closed. 3. Tachyon is transmitting a frame or waiting for ACKs and the outbound connection is opened and Tachyon receives an EOFDTI. 4. Tachyon receives an EOFDT that does not match the current sequence. This causes the outbound connection to be closed.

Table 6.16 Outbound Completion Message

PTI 172614

Word	Bit(s)	Field Label	Field name	Description
				5. Tachyon receives an EOFDT or EOFDTI before the sequence was completely transmitted and the outbound connection is opened. 6. Tachyon receives an ACK_EOFDT before the sequence was transmitted and Tachyon has sent an SOFc1. 7. Tachyon is transmitting a frame or waiting for ACKs and it receives a RJT_EOFDT that matched the current sequence, and the C1 and R bits are both set to one. 8. Tachyon is transmitting a frame or waiting for ACKs and it receives an ACK_ABT_EOFDT that matched the current sequence, and the C1 and A bits are both set to one.
	30	LD	Link Down	This bit indicates that the link went down during the transmission of a sequence. If the Link Down (LD) bit is set to one, then the link is down and the connection is closed, even though the Class Connection Open (CO) bit may also be set to one.
	29	A	Abort Requested	The remote node requests that this sequence be aborted. An ACK with the abort bits set was received.
	28	AT	ACK Timeout	Set to one when a current sequence has exceeded the ED_TOV timeout for the return of the ACKs.
	27	OT	Frame Timeout	Set to one when the current sequence has exceeded the ED_TOV timeout for the transmission of an outbound frame. Also set to one when an outbound frame was sent to a bad AL_PA.
	26	R	Rejected	A reject frame was received for the current sequence.
	25	X	Retries Exceeded	If a BSY is received for the first frame of the sequence and the Sequence Interlock (Lck) bit or SOFc1 bit is set in the ODB, then the X bit is set to one if Tachyon transmitted the frame 16 times or the Retry Disable bit is set to one in the ODB.
	24	HE	Host Programming Error	Hardware detected an invalid programming combination in the ODB or EDB, for example:

Table 6.16 Outbound Completion Message (Continued)

PTI 172615

Word	Bit(s)	Field Label	Field name	Description
				1. The Header Length is zero. 2. The CONT_SEQ and E_C bits in the ODB are both set to one. 3. The End Bit (E) in the ODB is set to one and Total Sequence Length is greater than the Header Length minus 32. 4. The Buffer Length is zero. 5. The sum of the Buffer Lengths in the EDB is less than the Total Sequence Length in the ODB.
	23	CO	Class 1 Connection Open	This is a copy of the internal connection state. It indicates that the Class 1 connection is still open. If the Link Down (LD) bit is set, then the link is down and the connection is closed even though this CO bit may be set.
	22..16	Reserved	Reserved	Initialize to zero.
	15..0	**RX_ID**	RX_ID	The RX_ID is assigned by the remote node for this exchange. If this is not the first sequence of the exchange, the RX_ID should be the same one sent to Tachyon in the ODB. The RX_ID value is invalid if the Link Down (LD), Frame Timeout (OT), or Host Programming Error (HE) bits are set.
3	31..16	**EE_Credit Count**	EE_Credit Count	The EE_Credit remaining at the completion of the sequence. If an error occurred, this value may be less than the starting value. The difference is the number of frames or ACKs lost.
	15..0	**SEQ_CNT**	SEQ_CNT	The sequence count of the last frame transmitted.
4..7	31..0	**Pad Word**	Pad Word	Pad values are undefined.

Table 6.16 Outbound Completion Message (Continued)

PTI 172616

6.4.4 Outbound High Priority Completion Message

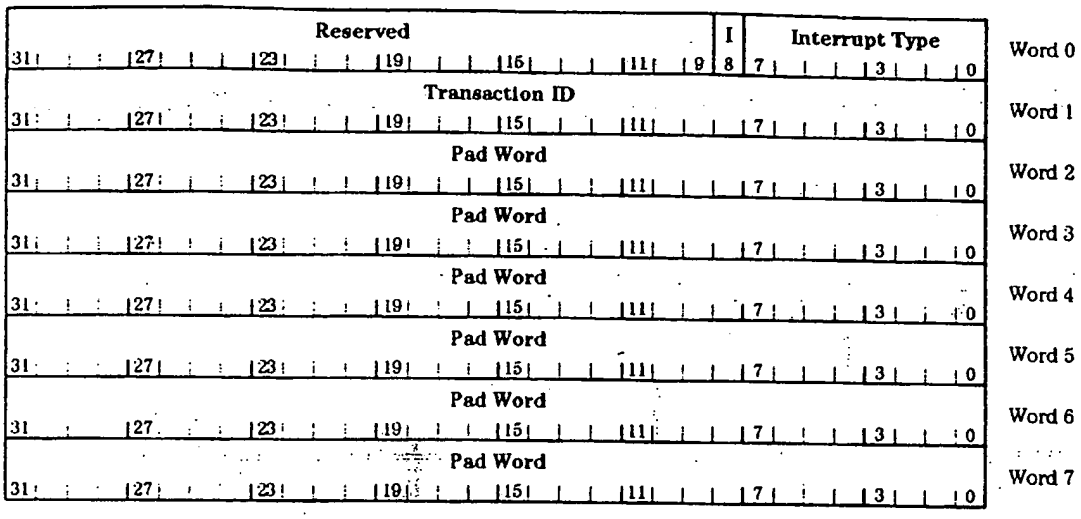

Word	Bit(s)	Field Label	Field Name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Desired	Set to one when the host has requested that an interrupt be generated for this completion message.
	7..0	Interrupt Type	Interrupt Type	out_hi_pri_completion = 0x01
1	31..0	Transaction ID	Transaction ID	A copy of the TRANS_ID in the HPDB
2..7	31..0	Pad Word	Pad Word	Pad values are undefined.

Table 6.17 Outbound High Priority Completion Message

PTI 172617

6.4.5 Inbound MFS and Inbound OOO Completion Messages

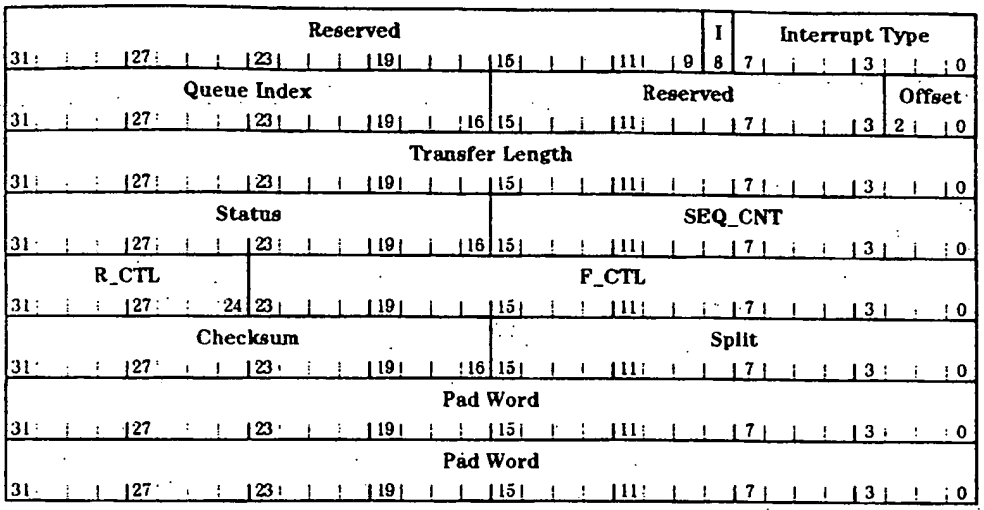

Word	Bit(s)	Field Label	Field Name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Generated	For inbound_mfs_completion messages, this bit is set to one to indicate that an interrupt was generated. For inbound_ooo_completion messages, this bit is cleared to zero.
	7..0	Interrupt Type	Interrupt Type	inbound_mfs_completion = 0x02 inbound_ooo_completion = 0x03
1	31..16	Queue Index	Circular Queue Index	The index value of the MFSBQ Entry that contained the last buffer address that was used for this sequence or fragment.
	15..3	Reserved	Reserved	Initialize to zero.
	2..0	Offset	Queue Entry Offset	Which of the eight addresses within the MFSBQ entry that was used for the last buffer of this sequence or fragment.
2	31..0	Transfer Length	Total Transfer Length	The total length, in bytes, of the data transferred to the host in this contiguous segment.
3		Status	Completion Status	This field comprises the following subfields.
	31	C1	Class 1 Error	The Class 1 connection experienced an error while a sequence was being received.
	30	LD	Link Down	The link went down while a sequence was being received.

Table 6.18 Inbound MFS and Inbound OOO Completion Message

PTI 172618

Word	Bit(s)	Field Label	Field Name	Description
	29	A	Abort Requested	The remote node is requesting that this sequence be aborted. An ABTS was received.
	28	T	Timeout	Current sequence has exceeded the ED_TOV timeout.
	27	DA	Deferred ACK	The final ACK for this MFS has not been sent. The host must use the informati in the FC header and the F_CTL field of this completion message for generating the final ACK which must be sent via the HPCQ.
	26..16	Reserved	Reserved	Initialize to zero.
	15..0	SEQ_CNT	SEQ_CNT	The expected sequence count of the next frame.
4	31..24	R_CTL	R_CTL	The value from the last frame header of this sequence.
	23..0	F_CTL	F_CTL	The value from the last frame header of this sequence.
5	31..16	Checksum	IP Checksum	An approximate checksum of the entire sequence. This value must be corrected by the host to remove the network header checksum.
	15..0	Split	First Buffer Length	The length, in bytes, of valid data in the first buffer. The first buffer usually contains just the FC header and is not filled. If TCP/UDP assists are enabled and the first frame of the sequence has a RO of zero and the Type field in the FC header is 5, for IS8802-2 LLC/SNAP, or matches the type specified in the Split Type field of the Tachyon Configuration register, then Tachyon places the entire first frame into the first buffer. If the RO is non-zero, or the TCP/UDP assists are disabled, or the type does not match, then Tachyon places only the Tachyon Header Structure into the first buffer and places the data payload into the next buffer. Refer to "3.10.3 Header/Data Splitting" on page 94.
6..7	31..0	Pad Word	Pad Word	Pad values are undefined.

Table 6.18 Inbound MFS and Inbound OOO Completion Message (Continued)

PTI 172619

6.4.6 Inbound SFS, Unknown Frame, Inbound Busied Frame, Read Status, Inbound SCSI Command, Bad SCSI Frame, and Inbound SCSI Status Completion Messages

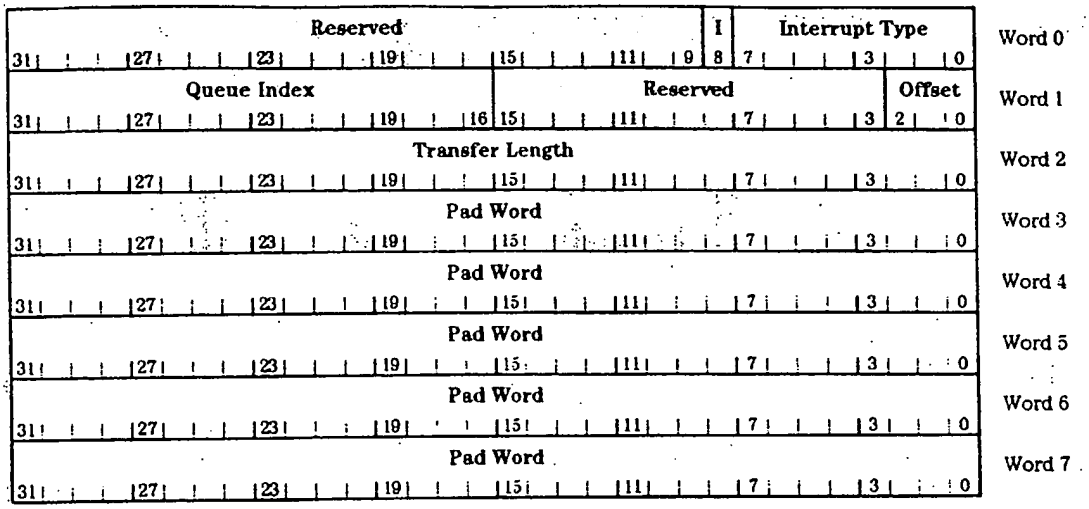

Word	Bit(s)	Field Label	Field Name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Generated	Set to one for all the completion messages indicated in the Interrupt Type field (below), except for the inbound_busied_frame completion messages. The I bit is cleared to zero if the interrupt type is an inbound_busied_frame completion message.
	7..0	Interrupt Type	Interrupt Type	The type of interrupt as follows: inbound_sfs_completion = 0x04 inbound_unknown_frame_i = 0x06 inbound_busied_frame = 0x06 read_status = 0x0B inbound_scsi_command = 0x0D bad_scsi_frame = 0x0E inb_scsi_status_completion = 0x0F
1	31..16	Queue Index	Circular Queue Index	The index of the SFSBQ entry that contained the last buffer address that was used for this sequence or fragment.
	15..3	Reserved	Reserved	Initialize to zero.
	2..0	Offset	Queue Entry Offset	Which of the eight addresses within the SFSBQ entry that was used for the last buffer of this sequence.

Table 6.19 Inbound SFS, Unknown Frame, Inbound Busied Frame, Read Status, Inbound SCSI Command, Bad SCSI Frame, and Inbound SCSI Status Completion Messages

Word	Bit(s)	Field Label	Field Name	Description
2	31..0	Transfer Length	Total Transfer Length	The total length, in bytes, of the data transferred to the host.
3..7	31..0	Pad Word	Pad Word	Pad values are undefined.

Table 6.19 Inbound SFS, Unknown Frame, Inbound Busied Frame, Read Status, Inbound SCSI Command, Bad SCSI Frame, and Inbound SCSI Status Completion Messages (Continued)

PTI 172621

6.4.7 Inbound SCSI Data Completion Message

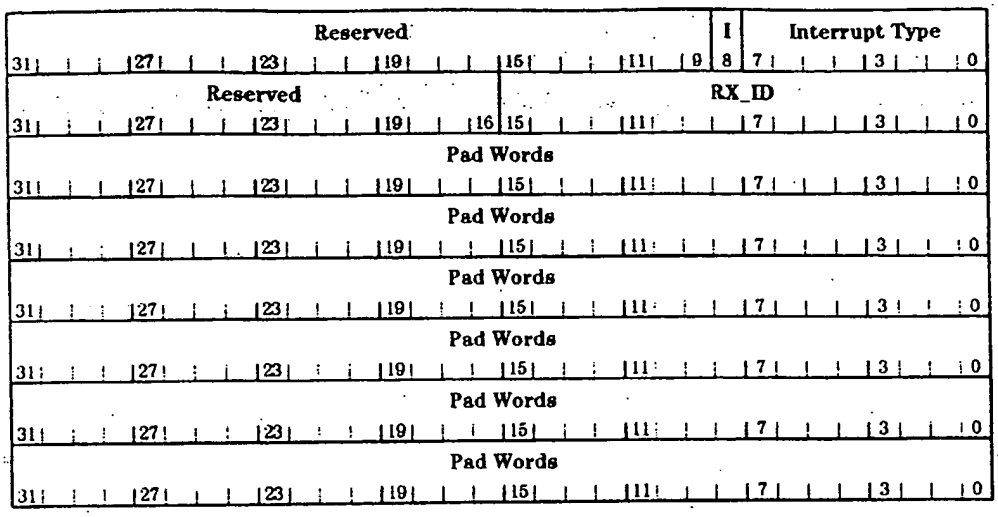

Word	Bit(s)	Field Label	Field Name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Desired	Set to one for the completion message indicated in the Interrupt Type field.
	7..0	Interrupt Type	Interrupt Type	inbound_scsi_data_completion = 0x0C
1	31..16	Reserved	Reserved	Initialize to zero.
	15..0	RX_ID	RX_ID	The responder exchange ID for the received sequence.
2..7	31..0	Pad Word	Pad Word	Pad values are undefined.

Table 6.20 Inbound SCSI Data Completion Message

PTI 172622

6.4.8 Inbound C1 Timeout, Buffer Warnings, and Frame Manager Interrupt Completion Messages

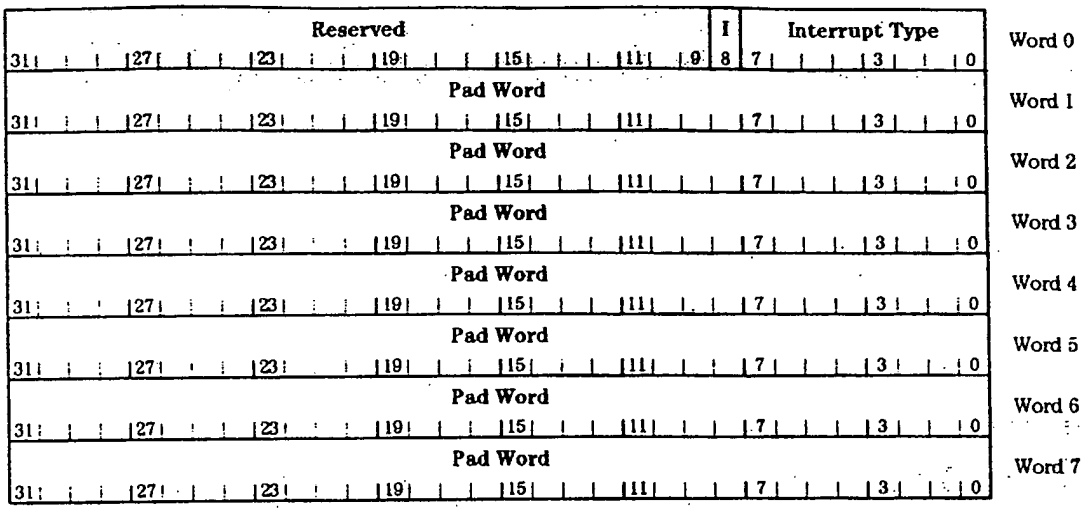

Word	Bit(s)	Field Label	Field Name	Description
0	31..9	Reserved	Reserved	Initialize to zero.
	8	I	Interrupt Generated	Set to one for all the completion messages indicated in the Interrupt Type field.
	7..0	Interrupt Type	Interrupt Type	The type of interrupt as follows: inbound_C1_timeout = 0x05 sfs_buf_warn = 0x07 mfs_buf_warn = 0x08 imq_buf_warn = 0x09 frame_mgr_interrupt = 0x0A
1..7	31..0	Pad Word	Pad Word	Pad values are undefined.

Table 6.21 Inbound C1 Timeout, Buffer Warnings, and Frame Manager Interrupt Completion Messages

PTI 172623

6.5 SCSI Data Structures

6.5.1 SEST

SEST Entry	Entry 0
SEST Entry	Entry 1
SEST Entry	...
SEST Entry	Entry n-1

Entry	Name	Description
0 .. n-1	SEST Entry	Each entry contains an Inbound or Outbound SEST Entry. Each entry is 32 bytes long.

Table 6.22 SEST

PTI 172624

6.5.2 Inbound SEST Entry

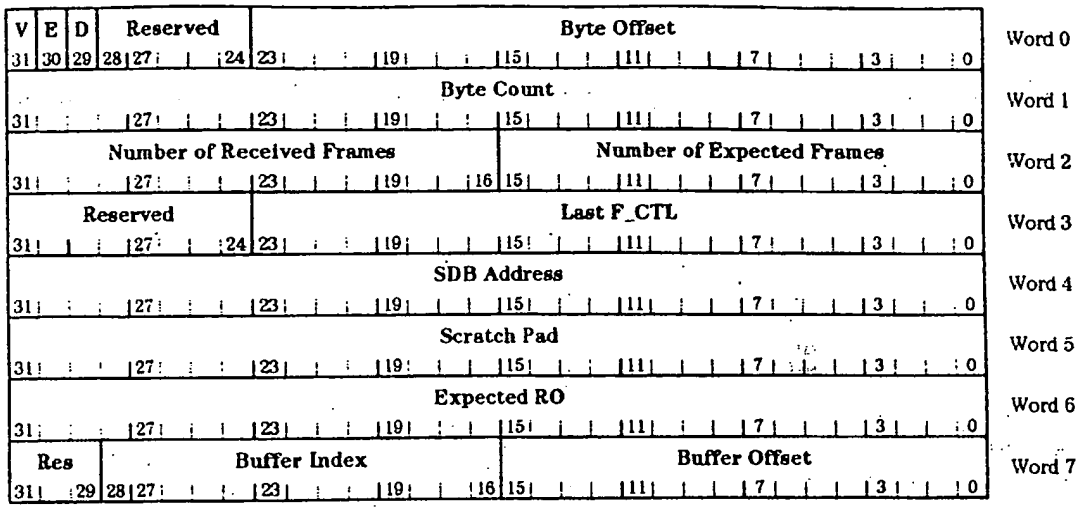

Word	Bit(s)	Field Label	Field Name	Description
0	31	V	SCSI Exchange State Table Entry (STE) Valid	Set to one by the host if this entry, which is indexed by the SEST entry is active. Refer to "3.8.4 Invalidating and Re-Using Inbound SEST Entries" on page 65. This bit is cleared by Tachyon when: 1. The Status sequence is received at the end of the exchange. 2. An Abort sequence (ABTS) frame is received for this SEST entry. 3. The host writes to the Tachyon Flush SEST Cache Entry register.
	30	E	SDB Error	A read beyond the end of the SDB was attempted with OOO Reassembly disabled.
	29	D	SCSI Direction	This bit indicates the direction of the exchange corresponding to the SEST entry. Set to one for an Inbound SEST entry.
	28..24	Reserved	Reserved	Initialize to zero.
	23..0	Byte Offset	Byte Offset	The offset, in bytes, within the first buffer address at which to place data. Used to displace the Relative Offset. Must be a multiple of four bytes. Must be cleared to zero if the OOO Reassembly Disable bit is set to one.
1	31..0	Byte Count	Byte Count	The number of bytes received. This field must be initialized to zero.

Table 6.23 Inbound SEST Entry

PTI 172625

Word	Bit(s)	Field Label	Field Name	Description
2	31..16	Number of Received Frames	Number of Received Frames	The number of frames received. This count is maintained by Tachyon. Initialize to zero.
	15..0	Number of Expected Frames	Number of Expected Frames	The number of frames expected, calculated by the difference of the last frame and first frame. Maintained by Tachyon. Initialize to zero.
3	31..24	Reserved	Reserved	Initialize to zero.
	23..0	Last F_CTL	Last F_CTL	The F_CTL value of the last frame. Maintained by Tachyon.
4	31..0	SDB Address	SDB Address	The address of the SCSI Descriptor Block.
5	31..0	Scratch Pad	Scratch Pad	May be used by the driver as a scratch pad for an 32-bit value, for example, the Transaction ID. This field remains untouched by Tachyon during the course of the exchange. This value is not copied into the completion message.
6	31..0	Expected RO	Expected Relative Offset	The expected relative offset of the next frame. Must be a multiple of four bytes and be aligned on a 4-byte boundary. Used for In Order Reassembly mode only.
7	31..29	Reserved	Reserved	Initialize to zero.
	28..16	Buffer Index	Buffer Index	The index in the SDB which points to the buffer where frame data can be placed. Used for In Order Reassembly mode only.
	15..0	Buffer Offset	Buffer Offset	The location in the buffer where frame data can begin to be placed. Used for In Order Reassembly mode only. This can be used to align the data received on boundaries other than sizeof(SCSI Buffer Length).

Table 6.23 Inbound SEST Entry (Continued)

Inbound SEST Entry Notes

1. When the host creates an inbound SEST entry, words 1-3, and 7 must be cleared to zero before starting the exchange.
2. The STE Valid (V) bit should not be set to one in the Inbound SEST Entry until all associated structures have been built.
3. All fields in the Inbound SEST Entry must be cleared to zero at initialization.
4. The Expected Relative Offset (Exp_RO) field is only used by targets that have to break up the data transfer phase into more than one sequence. Because the targets send out the FCP_XFER_RDYs requesting the data with the RO included, they can program the Exp_RO field before sending the FCP_XFER_RDY to the initiator.

PTI 172626

6.5.3 Outbound SEST Entry

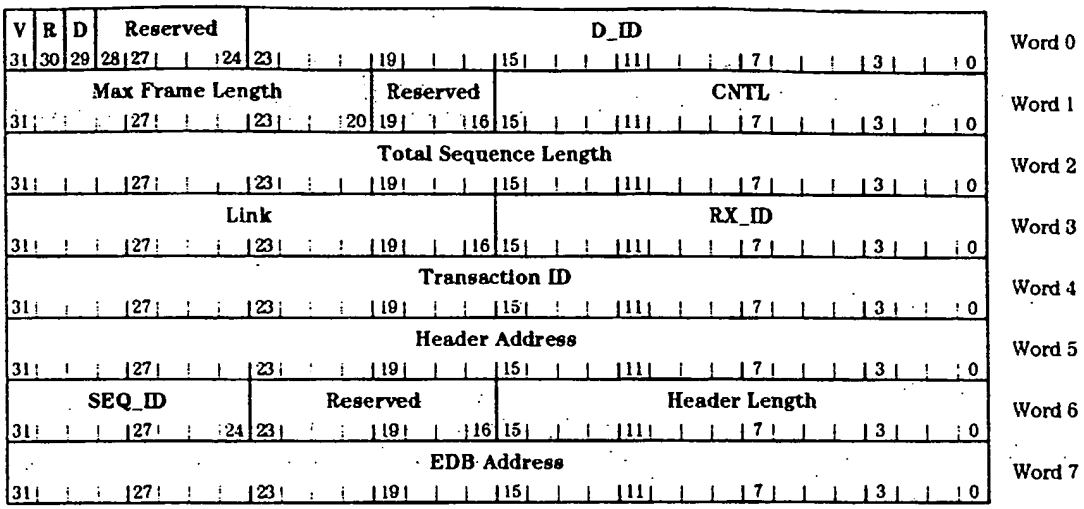

Word	Bit(s)	Field Label	Field Name	Description
0	31	V	SCSI Exchange State Table Entry (STE) Valid	Set to one if this entry, indexed by the OX_ID, is active. Never cleared by Tachyon. Refer to "3.8.3 Invalidating and Re-Using Outbound SEST Entries" on page 64.
	30	R	Reserved	Initialize to zero.
	29	D	SCSI Direction	This bit indicates the direction of the exchange corresponding to the SEST entry. Cleared to zero for an Outbound SEST entry.
	28..24	Rserved	Reserved	Initialize to zero.
	23..0	D_ID	D_ID	Tachyon writes this field with the S_ID returned in the FCP_XFER_RDY frame from the target device. Used in subsequent data transfers.
1	31..20	Max Frame Length	Maximum Frame Length	The maximum size, in bytes, of the frames that the sequence is segmented into as provided by the login parameters.
	19..16	Reserved	Reserved	Initialize to zero.
		CNTL	Control	This field comprises the following Tachyon control bits for this sequence.
	15..14	Cl	Class	FC class of service to use: 1 = Class 1 2 = Class 2 3 = Class 3 0 = Undefined

Table 6.24 Outbound SEST Entry

PTI 172627

Word	Bit(s)	Field Label	Field Name	Description
	13	Lck	Sequence Interlock	Set to zero for SCSI.
	12	SOFc1	Start Class 1 Connection	Set to one when the host want to send SCSI data on a Class 1 connection.
	11	E_C	End Connection	Must be set if the SOFc1 bit is set, otherwise this bit should be cleared to zero.
	10	No Comp	Completion Message Disable	Set to one when the host does not want the OSM to generate a completion message for this sequence. This bit should normally be set.
	9	No Int	Completion Message Interrupt Disable	Set to one when the host does not want an interrupt when the completion message is posted back to it after data is sent.
	8	ACK_0	ACK_0	Use the ACK_0 model for transmitting the sequence. Send all frames without waiting for individual frame ACKs, but wait for an ACK for the entire sequence.
	7..4	Reserved	Reserved	Initialize to zero.
	3..0	EE_Credit	EE_Credit	The EE_Credit, given during host login, to use to send the sequence.
2	31..0	**Total Sequence Length**	Total Sequence Length	The host initially fills in this field with the length, in bytes, of the data to be sent. Tachyon then overwrites this field with the length returned in the FCP_XFER_RDY frame from the target device.
3	31..16	**Link**	Link	This field is used by Tachyon to maintain a linked list of outbound SCSI oeprations. This field must be set to 0xFFFF by the host when building the STE.
	15..0	**RX_ID**	RX_ID	Tachyon fills in this field with the RX_ID returned in the FCP_XFER_RDY frame from the target device. This value is used in the subsequent data transfer.
4	31..0	**Transaction ID**	Transaction ID	This value is used by the driver to match outbound completions to requests. It is returned by Tachyon unmodified on completions, if generated.
5	31..0	**Header Address**	Header Address	The host address of the Tachyon Header Structure and Optional Headers, if they exist. Must be aligned on a 4-byte boundary.
6	31..24	**SEQ_ID**	SEQ_ID	The unique FC sequence identifier for the data transfer.

Table 6.24 Outbound SEST Entry (Continued)

PTI 172628

Word	Bit(s)	Field Label	Field Name	Description
	23..16	Reserved	Reserved	Initialize to zero.
	15..0	Header Length	Header Buffer Length	The length, in bytes, of the FC frame header. The length must be at least 32 bytes.
7	31..0	EDB Address	EDB Address	The address of the EDB. It must be 32-byte aligned.

Table 6.24 Outbound SEST Entry (Continued)

PTI 172629

6.5.4 SDB, In Order Reassembly

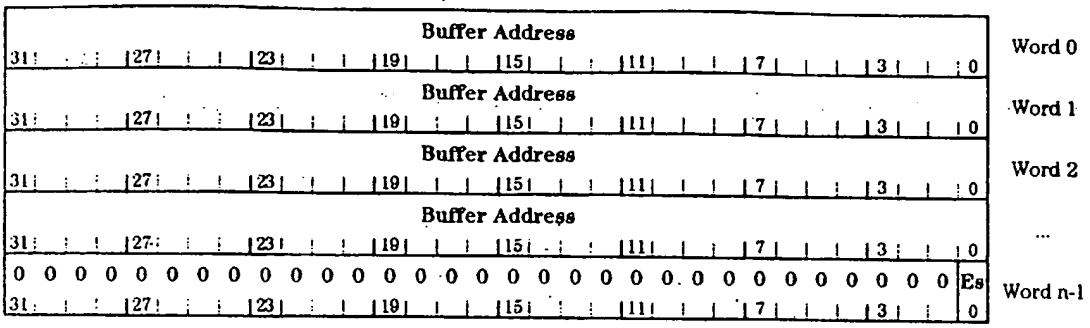

Word	Bit(s)	Field Label	Field Name	Description
0.. n-2	31..0	Buffer Address	Buffer Address	The address to a designated SCSI buffer. Must be a multiple of four bytes.
n-1	31..1	0		The last SDB entry is the End Delimiter. Bits 31..1 must be cleared to zeros.
	0	Es	End of SDB	This bit signals the end of the SDB for In Order Reassembly.

Table 6.25 SDB, In Order Reassembly

SDB, In Order Reassembly Notes

1. When the OOO Reassembly Disable bit in the Tachyon Configuration register is set to one, then inbound SCSI frames are received in order.
2. In the In Order Reassembly mode, the SDB may be any length. The host sets the entry following the last valid address (End Delimiter) to signal the end of the SDB. When the host sets the End of SDB (Es) bit to one and Tachyon attempts to write past the last valid buffer, Tachyon sets the SDB Error bit in the Inbound SEST Entry and continues accepting data. Any data that is received after the last buffer has been filled will start overwriting the first word of the first buffer.
3. The SDB must be aligned on a power of 2 boundary equal to or greater than the SDB length.
4. The SDB length is the number of entries multiplied by 4 bytes. The length must always be between 4 and 64k bytes.
5. The buffer addresses must be on a word-aligned boundary.
6. The Buffer Length does not have to be a power of 2.

PTI 172630

6.5.5 SDB, OOO Reassembly

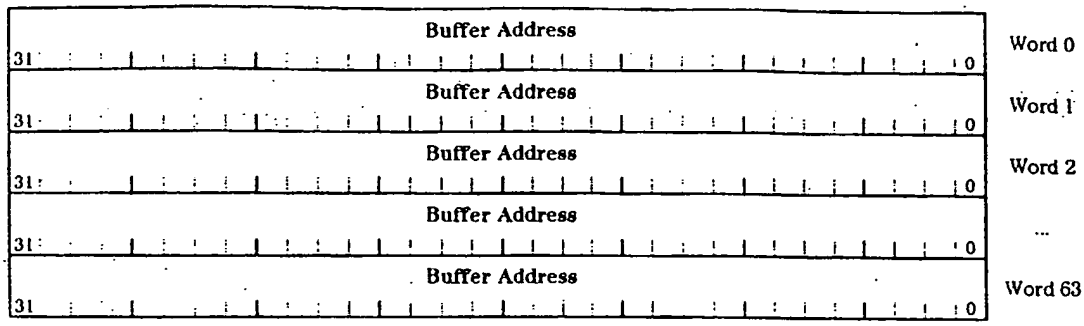

Word	Bit(s)	Field Label	Field Name	Description
0..63	31..0	Buffer Address	Buffer Address	The address to a designated SCSI data buffer.

Table 6.26 SDB, OOO Reassembly

SDB, OOO Reassembly Notes

1. When the OOO Reassembly Disable bit of the Tachyon Configuration register is cleared to zero, then inbound frames may be received out of order.
2. For OOO Reassembly, the SDB contains 64 address entries to completely define the sequence.
3. Typically, inbound SCSI transactions are in the 8K byte to 16K byte range and only require two to four entries in the SDB. However it is recommended that all 64 entries in the SDB be filled in. Unused locations in the SDB should be set to point to a safe area of host memory. If a frame arrives with an invalid relative offset (RO) value due to an error, then the incoming data is written harmlessly to the safe location. If the unused locations within the SDB are not filled in (perhaps to conserve host memory) then random host memory locations may be written to by incoming data with a invalid RO value.
4. The SDB must be aligned on a 256-byte boundary.
5. The buffer addresses in the SDB must begin on a memory address boundary that is multiple of the receive buffer length. To start the logically contiguous data buffer on an arbitrary word boundary, the host may define a buffer offset within the SEST Entry that tells Tachyon to start filling in data at a location offset from the first word of the first buffer. The maximum length for an inbound SCSI sequence is (64 * the buffer length). Therefore, if an inbound SCSI sequence of maximum length is received and the Buffer Offset in the SEST entry is non-zero, then the SDB buffer ends at ((64 * the buffer length) - the Buffer Offset) and the remaining data wraps to the first SDB buffer.
6. If an overwrite occurs past the 64th buffer on a SEST out of order operation, Tachyon wraps around and overwrites the n-64th buffer, that is, a write to the 65th buffer overwrites the first buffer, a write to the 66th buffer overwrites the second buffer, and so on. This feature may be useful for a SEST entry that uses the Byte Offset (B_Offset) field. If the transfer is greater than (64 * SCSI Buffer Length - B_Offset) but less than (64 * SCSI Buffer Length) then the extra data is written to the first buffer without corrupting it. This wrap around feature prevents Tachyon from writing to some undetermined place in host memory if an overwrite occurs.

7. Tachyon Signal Descriptions

7.1 Tachyon Logic Symbol

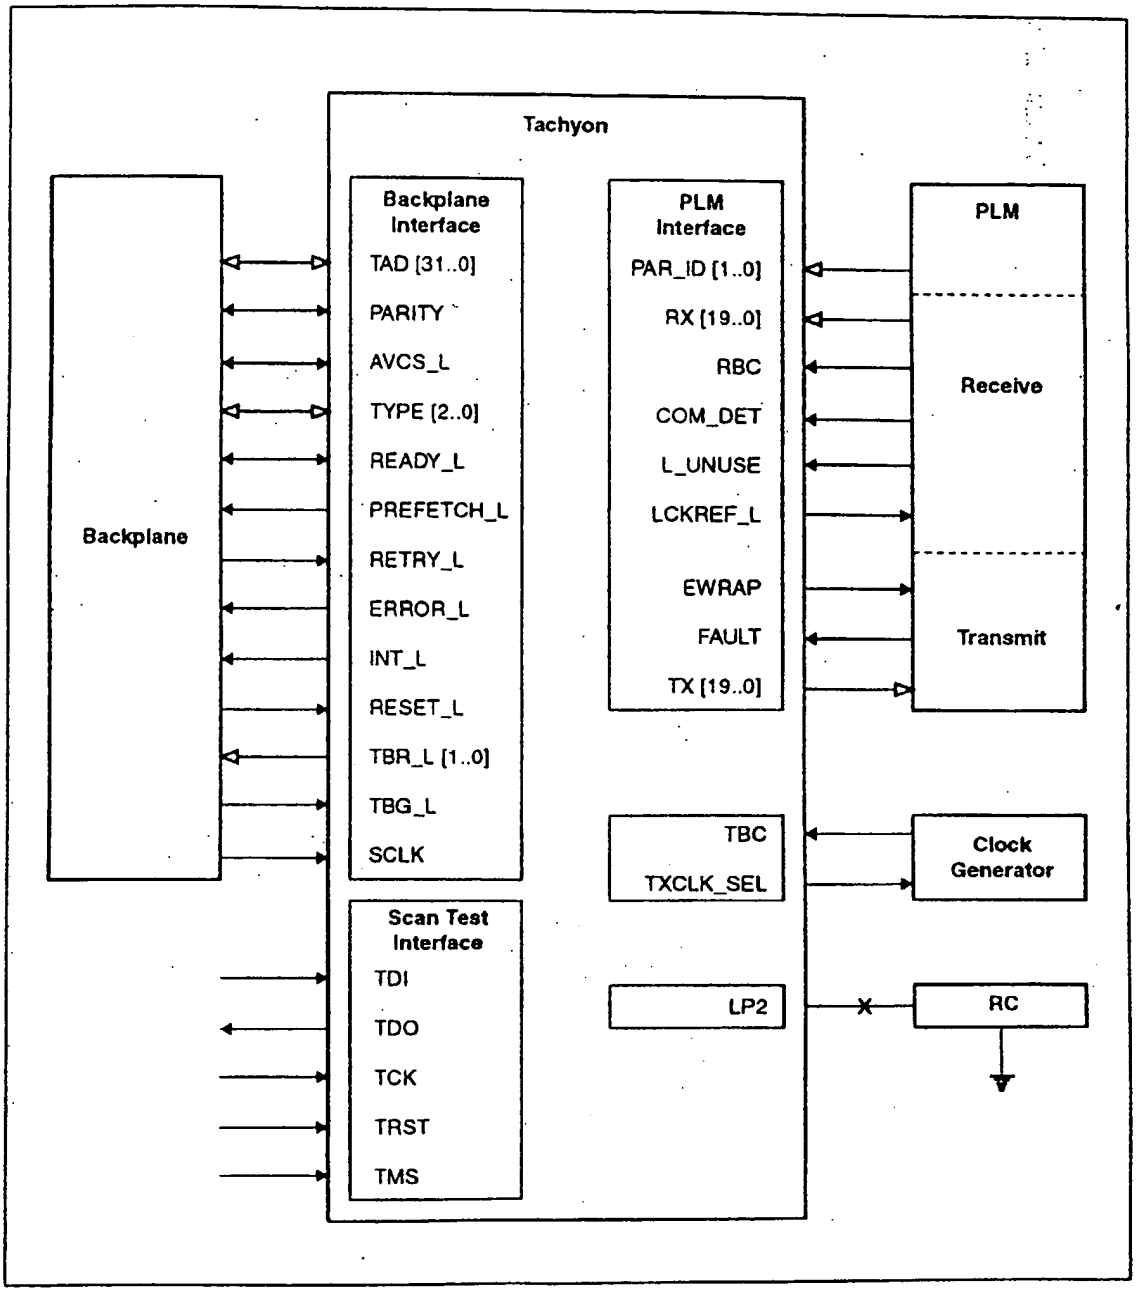

Figure 7.1 Tachyon Logic Symbol

PTI 172632

7.2 Tachyon Pin-out

Pin	Pad Label	Pin	Pad Label	Pin	Pad Label	Pin	Pad Label
1	Vdd2 (a)	53	Vdd	105	Vdd2 (a)	157	Vdd
2	Vss2 (e)	54	Vss2 (e)	106	TAD[26]	158	Vss
3	RX[8]	55	IDD_TEST (c)	107	TAD[27]	159	TX[17]
4	RX[9]	56	LP2	108	TAD[28]	160	TX[19]
5	RX[19]	57	PLLAGND	109	Vdd	161	L_UNUSE
6	Vdd	58	PLLVdd	110	Vss	162	TXCLK_SEL
7	Vss	59	PLLVss	111	TAD[29]	163	Vdd2 (a)
8	RX[17]	60	TAD[0]	112	TAD[30]	164	Vss
9	RX[15]	61	TAD[1]	113	TAD[31]	165	TBC
10	RX[13]	62	TAD[2]	114	Vdd2 (a)	166	Vdd2 (a)
11	Vdd2 (a)	63	Vdd	115	Vss2 (e)	167	Vss2 (e)
12	Vss2 (e)	64	Vss	116	TYPE[0]	168	TX[1]
13	RX[11]	65	TAD[3]	117	TYPE[1]	169	TX[3]
14	RX[10]	66	TAD[4]	118	TYPE[2]	170	TX[5]
15	EWRAP	67	TAD[5]	119	PARITY	171	Vss
16	Vdd2 (a)	68	Vdd2 (a)	120	Vdd	172	Vss
17	Vss2 (e)	69	Vss2 (e)	121	Vss	173	TX[7]
18	RX[18]	70	TAD[6]	122	TCK	174	TX[0]
19	RX[16]	71	TAD[7]	123	Vss2 (e)	175	TX[2]
20	RX[14]	72	Vdd	124	TMS	176	Vss
21	Vdd	73	Vss	125	RSTN	177	Vss
22	Vss	74	TAD[8]	126	TDI	178	TX[4]
23	RX[12]	75	TAD[9]	127	Vss2 (e)	179	TX[6]
24	PAR_ID[1]	76	TAD[10]	128	TDO	180	TX[8]
25	Vdd	77	Vss	129	Vss2 (e)	181	Vdd
26	Vss2 (e)	78	Vss	130	Vss2 (e)	182	Vss
27	Vss	79	Vdd	131	Vdd2 (a)	183	Vss
28	PREFETCH_L	80	TAD[11]	132	reserved (d)	184	TX[9]
29	TBG_L	81	TAD[12]	133	reserved (d)	185	PAR_ID[0]
30	TBR_L[1]	82	TAD[13]	134	reserved (d)	186	Vss
31	TBR_L[0]	83	TAD[14]	135	Vdd	187	Vss
32	Vdd2 (a)	84	Vdd2 (a)	136	SCAN_EN (d)	188	RBC
33	Vss2 (e)	85	Vss2 (e)	137	reserved (b)	189	Vdd
34	RESET_L	86	TAD[15]	138	reserved (b)	190	Vss
35	INT_L	87	TAD[16]	139	reserved (b)	191	RX[1]
36	ERROR_L	88	Vdd	140	TEST_MODE	192	RX[3]
37	Vdd2 (a)	89	Vss	141	reserved (d)	193	RX[5]
38	Vss2 (e)	90	TAD[17]	142	Vdd2 (a)	194	Vdd2 (a)
39	RETRY_L	91	TAD[18]	143	Vss2 (e)	195	Vss2 (e)
40	READY_L	92	TAD[19]	144	TX[12]	196	RX[7]
41	AVCS_L	93	Vdd2 (a)	145	TX[14]	197	LCKREF_L
42	Vdd	94	Vss2 (e)	146	TX[16]	198	FAULT
43	Vss	95	TAD[20]	147	Vdd2 (a)	199	Vdd2 (a)
44	PLL_RSTN	96	TAD[21]	148	Vss2 (e)	200	Vss2 (e)
45	reserved (c)	97	TAD[22]	149	TX[18]	201	COM_DET
46	PLL_IDD_TEST (b)	98	Vdd	150	TX[10]	202	RX[0]
47	PLL_TEST_DATA (b)	99	Vss	151	TX[11]	203	RX[2]
48	PLL_TEST_MODE (c)	100	TAD[23]	152	Vdd	204	Vdd2 (a)
49	Vdd	101	TAD[24]	153	Vss	205	Vss
50	Vss	102	TAD[25]	154	TX[13]	206	RX[4]
51	SCLK	103	Vss2 (e)	155	TX[15]	207	RX[6]
52	Vdd2 (a)	104	Vdd2 (a)	156	Vdd2 (a)	208	Vdd

Table 7.1 Tachyon Pinout

PTI 172633

Tachyon Pinout Notes

(a) For most applications, Vdd and Vdd2 internal power straps may be connected on the printed circuit board (PCB). Refer to "A.1 PCB Layout Suggestions" on page 313.
(b) Pin has internal pull-up and may be left unconnected (solder pad only) or pulled up to Vdd (3.6 V) with a 4.7kΩ resistor.

WARNING	Do not pull up reserved pins to a voltage greater than Vdd (3.6V)

(c) Pin has internal pull-down and may be left unconnected (solder pad only) or connected to Vss.
(d) Pin must be left unconnected.
(e) For most applications, Vss and Vss2 may be connected on the PCB. Refer to "A.1 PCB Layout Suggestions" on page 313.

PTI 172634

7.3 Physical Link Module (PLM) Interface

Pin #	I/O	Pad Label	Pad Name	Description
Various	O	TX[19..0]	Transmit Data	10-bit or 20-bit transmit data bus to the PLM. For 10-bit PLMs, only the TX[9..0] outputs are used. The PAR_ID bit determines whether a 10-bit or 20-bit PLM should be used.
165	I	TBC	Transmit Byte Clock	53.125 MHz or 26.5625 MHz clock (± 0.01% average tolerance) supplied by the oscillator. Note that for the TBC: 1) at 26 MHz, a maximum 40/60 duty cycle is allowed; 2) at 53 MHz, only a maximum 45/55 duty cycle is allowed.
Various	I	RX[19..0]	Receive Data	10-bit or 20-bit receive data bus from the PLM. For 10-bit PLMs, only the RX[9..0] outputs are used. The PAR_ID bit determines whether a 10-bit or 20-bit PLM should be used.
188	I	RBC	Receive Byte Clock	The receive byte clock signal is derived from the incoming data stream. The clock frequency is 53.125 MHz for 1063 Mbaud and 10-bit 531 Mbaud PLMs, and 26.5625 MHz for 10-bit 266 Mbaud PLMs (± 0.01%). 20-bit 531 Mbaud PLMs are not supported. Note that for RBC: 1) at 26 MHz, a maximum 40/60 duty cycle is allowed; 2) at 53 MHz, only a maximum 45/55 duty cycle is allowed.
15	O	EWRAP	Electrical Wrap Enable	When asserted high by a host Write to the Frame Manaager Configuration register, this signal tells the PLM to loopback the serialized transmit data to the receive deserializer. This signal is also used as a reset to internal logic on the PLM in response to a laser fault signal to perform recovery. An external loopback hood is not required for EWRAP.
197	O	LCKREF_L	Lock to Reference	When asserted low, this signal provides the mechanism for directing the receive clock generation circuits to obtain frequency lock on a multiple of the Transmit Byte Clock (TBC). Also, when asserted low, this signal causes the PLM to lock its PLL to the TBC.

Table 7.2 Physical Link Module (PLM) Interface

PTI 172635

<table>
<tr><th>Pin #</th><th>I/O</th><th>Pad Label</th><th>Pad Name</th><th>Description</th></tr>
<tr><td>201</td><td>I</td><td>COM_DET</td><td>Comma Detect</td><td>When asserted high by the PLM, this signal indicates when a K28.5 control character containing a comma of positive disparity (Refer to the "FC-PH") is detected on the link data stream at the same time as the data is presented on the Receive Data lines.</td></tr>
<tr><td>161</td><td>I</td><td>L_UNUSE</td><td>Link Unusable</td><td>When asserted high, this signal indicates that the link is unusable for data transfer, that is, light is no longer being received.</td></tr>
<tr><td>198</td><td>I</td><td>FAULT</td><td>Fault</td><td>When asserted high, this signal indicates that a fault has been detected on the module, for example, an improper power level has occurred on the laser.</td></tr>
<tr><td>162</td><td>O</td><td>TXCLK_SEL</td><td>Transmit Clock Select</td><td>Indicates which clock frequency the external clock multiplexing hardware should supply. This signal is derived from the PAR_ID lines.
TXCLK_SEL: 0 = 53.1250 MHz; 1 = 26.5625 MHz</td></tr>
<tr><td>24
185</td><td>I</td><td>PAR_ID[1]
PAR_ID[0]</td><td>Parallel ID</td><td>These two bits are asserted by the PLM to indicate the link rate and assumed interface width of the PLM. They are interpreted by Tachyon as follows.
PAR_ID[1..0]: 01 = 265.625 MHz, 10 bits; 10 = 531.25 MHz, 10 bits; 11 = 1062.5 MHz, 20 bits</td></tr>
</table>

Table 7.2 Physical Link Module (PLM) Interface (Continued)

Note	Tachyon supports GLMs that conform to the FCSI-301-Revision 1.0 GLM family.

PTI 172636

7.4 PLL External Connections

Tachyon contains a PLL which is used for internal functionality. This section describes how the external pins of this PLL should be connected.

Pin #	I/O	Pad Label	Pad Name	Description
44	I	PLL_RSTN	PLL Reset	When asserted low, this signal is the reset signal for the PLL. During normal operation this pin should be left unconnected. This pin has an internal pull-up.
46	I	PLL_IDD_TEST	PLL Quiescent Current (IDDQ) Test Enable	This pin has an internal pull-up. Refer to "LSI's LCB500K Design Manual, 1994, sections 3.24 and 3.6.3 (IDD Test)" for usage.
47	I	PLL_TEST_DATA	PLL Test Data Input	This pin controls the system clock when PLL_TEST_MODE is asserted high. This pin has an internal pull-up.
48	I	PLL_TEST_MODE	PLL Test Mode Select	When asserted high, the PLL's VCO output is bypassed, and PLL_TEST_DATA drives the system clock. This pin has an internal pulldown.
56	I	LP2	PLL External RC Network	This pin is used for external RC circuitry of the PLL. This pin should be connected to a 200Ω 5% resistor, 1nF 5% capacitor, and the PLLAGND pin in series. Refer to "8.6 External PLL Components" on page 252.
57	I	PLLAGND	PLL Analog Ground	This pin is used for external RC circuitry of the PLL. This pin should be connected to a 1nF 5% capacitor, 200Ω 5% resistor, and the LP2 pin in series. Refer to "8.6 External PLL Components" on page 252.
58	I	PLLVdd	PLL Vdd	This pin should be connected to a Vdd source via ferrite bead with capacitor to Vss. For additional details and alternative methods of connecting this pin, refer to "A.1 PCB Layout Suggestions" on page 313.
59	I	PLLVss	PLL Vss	This pin should be connected to a Vss source. Refer to "A.1 PCB Layout Suggestions" on page 313.

Table 7.3 PLL External Connections

PTI 172637

7.5 JTAG 1149.1 Scan Test Interface

Pin #	I/O	Pad Label	Pad Name	Description
55	I	IDD_TEST	Boundary Scan (JTAG) Cells' IDDQ Test Enable	This pin has an internal pulldown. Refer to LSI or "LSI's LCB500K Design Manual, 1994, p. 13-82" for details.
122	I	TCK	Test Clock	Transitions the Test Access Port (TAP) state machine to the next state on the rising edge of this signal. If not used, leave unconnected or pull up with a 4.7kΩ resistor to Vdd.
124	I	TMS	Test Mode Select	This signal determines the next state of the TAP state machine. If not used, leave unconnected or pull up with a 4.7kΩ resistor to Vdd.
125	I	RSTN	TAP Reset	When asserted low, the TAP controller and RAM BIST are reset. If not used, leave unconnected or pull up with a 4.7kΩ resistor to Vdd.
126	I	TDI	Test Data In	Used to scan data serially into the boundary register cells. If not used, leave unconnected or pull up with a 4.7kΩ resistor to Vdd.
128	O	TDO	Test Data Out	Used to scan data serially from the boundary register cells. If not used, leave unconnected or pull up with a 4.7kΩ resistor to Vdd.
136	I	SCAN_EN	Interanl Scan Chain Enable	This pin has an internal pulldown. Assert this pin high when performing internal scan chain tests.
140	I	TEST_MODE	Test Enable	When asserted high, this control signal tri-states all pins on the Tachyon package, with the exception of "reserved" pins 132, 133, 134, and 141. If not used, connect Vss through a 4.7kΩ resistor.

Table 7.4 JTAG Scan Test Interface

PTI 172638

7.5.1 JTAG Instructions

Instruction	Instruction Type	Encoding
bypass	Mandatory	11111
extest	Mandatory	00000
sample	Mandatory	00001
clamp	Optional	00110
high-z	Optional	00111

Table 7.5 JTAG Instructions

For information on the use of JTAG Boundary Scan, refer to the "IEEE 1149.1 Boundary Scan specification".

7.6 Tachyon System Interface (TSI)

This section describes the Tachyon System Interface for the Tachyon Fibre Channel chip. Included in this discussion is a description of the signals as well as a description of the protocol used across this interface.

7.6.1 TSI Signal List

Pin #	I/O	Pad Label	Pad Name	Description
Various	I/O	TAD[31..0]	Multiplexed Address/ Data	Address is driven by the transaction master during the address phase. The master drives only word addresses, so only bits TAD[31..2] contain useful information; however, parity must be generated for TAD[31..0]. During the data cycles(s), data is driven by the device being read or by the device mastering a write.
119	I/O	PARITY	Parity Information	This bi-directional signal carries parity information for address and data on the TSI. Parity is optional, and is defined by the Parity Even (bit 1) and Parity Enable (bit 2) bits in the Tachyon Configuration register.
11	I/O	AVCS_L	Address Valid Chip Select	This signal is driven by the master of a transaction during the address phase.
118 117 116	I/O	TYPE[2] TYPE[1] TYPE[0]	Transaction Types	Asserted low by the master of a transaction. TYPE indicates the size and direction of the transaction. Refer to "7.6.2 TSI Transaction Types" on page 207.
10	I/O	READY_L	Device Ready	Asserted by a transaction responder when a transaction is complete. Refer to "7.6.8 Driving Bi-Directional Signals" on page 211. More information about READY_L is available in "7.6 Tachyon System Interface (TSI)" and "7.7 TSI Functional Waveforms".
39	I	RETRY_L	Read Retry	Asserted by the host to terminate a read request if the read data cannot be made available to Tachyon.
28	O	PREFETCH_L	Prefetch Request	Asserted low by Tachyon to signal that the next sequential data will be read. This signal may be active during retry. Refer to "7.6.10 Read Transactions" on page 211.
36	O	ERROR_L	Error Out	Asserted low by Tachyon to indicate a bus parity error, protocol error, or internal parity error.
35	O	INT_L	Interrupt	Asserted low by Tachyon to signal that a message has been posted to the host.

Table 7.6 TSI Signal List

PTI 172640

Pin #	I/O	Pad Label	Pad Name	Description
34	I	RESET_L	Synchronous Reset	Asserted low by the host to perform a hard reset of Tachyon. RESET_L must be held low for a minimum of 10 clock periods. All configuration information is lost on reset.
30 31	O	TBR_L[1] TBR_L[0]	Bus Requests	One bus request signal is asserted when Tachyon needs to master a transaction. TBR_L[1] indicates a read using the prefetched channel. TBR_L[0] is used for writes and non-prefetched reads. If only one bus request signal is desired, these two active low signals should be ANDed together, external to the chip. Refer to "7.6.17 Arbitration" on page 216.
29	I	TBG_L	Bus Grant	Asserted low by the host to signal acceptance of Bus Request.
51	I	SCLK	System Clock	24-40 MHz system clock used to drive the back-plane side of Tachyon.

Table 7.6 TSI Signal List (Continued)

PTI 172641

7.6.2 TSI Transaction Types

The transaction type is encoded onto the TYPE[2..0] signals during the address phase. The master of a transaction asserts AVCS_L along with the TYPE [ ] and TAD [ ] busses to indicate that an address cycle is in progress. TSI transaction types are noted in the following table.

TYPE[2..0]		Transaction Type	Transaction Name
Binary	Hexadecimal		
000	0x0	Single-word read	READ1
001	0x1	Double-word read	READ2
010	0x2	Four-word read	READ4
011	0x3	Eight-word read	READ8
100	0x4	Single-word write	WRITE1
101	0x5	Double-word write	WRITE2
110	0x6	Four-word write	WRITE4
111	0x7	Eight-word write	WRITE8

Table 7.7 TSI Transaction Types

All transactions must be aligned to their size. For example, the three least significant word address bits for an 8 word transaction (TAD[4..2]) must all be zero. The two least significant bits (TAD[1..0]) are never used, and it is recommended they are cleared to zero for future compatibility.

In transactions where the host is master of the transaction and Tachyon is the responder, only single word Read and Write transaction types are allowed. Tachyon asserts ERROR_L in response to any other size transaction.

PTI 172642

7.6.3 Data Structure Transaction Size

Data Structure Transaction Size

The following table shows the size of transactions that Tachyon uses when accessing host-based and SCSI data structures.

Host Data Structure	Operation
Outbound Command Queue (OCQ)	READ8
High Priority Command Queue (HPCQ)	READ8
Inbound Message Queue (IMQ)	WRITE8
Single Frame Sequence Buffer Queue (SFSBQ)	READ8
Multiframe Sequence Buffer Queue	READ8
Outbound SEST Entry Read	READ4, and sometime later, a READ8 is performed
Inbound SEST Entry Read	READ4, and sometime later, a READ1 is performed
Outbound SEST Entry Writeback	WRITE4 of words 0 through 3
Inbound SEST Entry Writeback	WRITE1 for each of words 0, 2, and 3

Table 7.8 Host Data Structure Transaction Size

7.6.4 TSI Transaction Protocol

TSI provides a basic transaction protocol which uses two major operations: Write transactions and Read transactions. Every transaction has a master (which performs either a Write or a Read) and a responder. If the host is the master of a transaction, Tachyon is the responder in that transaction. Similarly, if Tachyon is the master of a transaction, then the host is the responder in that transaction.

The master of a transaction drives an address and transaction type onto the TAD [ ] and TYPE [ ] busses, respectively, while asserting AVCS_L to indicate the start of the transaction. If Tachyon masters a transaction, the host, as the responder, uses READY_L as its acknowledgment signal. Similarly, if the host masters the transaction, Tachyon, as the responder, uses READY_L as its acknowledgment signal. Transaction protocol, including timing and details of use of acknowledgment signals, is described in the following sections.

When Tachyon is the responder, only transactions of one word are allowed. Transaction sizes of other than one word, when Tachyon is the responder, causes Tachyon to assert ERROR_L.

7.6.5 Streaming

When Tachyon obtains mastership of TSI and has more than one transaction to perform, Tachyon may extend its bus tenancy and perform several TSI transactions (up to the maximum programmed limit) before releasing mastership. This is known as "streaming".

Streaming Rules

1. The maximum number of transactions that may occur during a stream is programmable. The Write Stream Size and Read Stream Size fields in the Tachyon Configuration register indicate the maximum number of Write or Read transactions to perform per bus tenancy. This maximum number can be different for Writes and Reads.
2. Streaming occurs in only one direction, i.e., a Write and a Read transaction never occurs during the same stream.
3. Streaming may occur with different transaction sizes, e.g., a Read of 8 words and a Read of 4 words can occur during the same stream.
4. Streaming does not cross non-contiguous address locations, i.e., Tachyon releases ownership of the bus whenever the next address to be accessed is not sizeof (last_transaction) past the address of the last transaction.
5. Streaming does not cross A/L pairs.
6. Tachyon does not stream across frames.

Tachyon Terminates TSI Bus Tenancy When:

1. Tachyon determines that it does not have more data to transfer (e.g., an end of frame occurs, end of sequence data occurs, Tachyon encounters an error, etc).
2. In the current tenancy, Tachyon completes the number of transactions in the programmed maximum stream size in the Tachyon Configuration register.
3. The host asserts RETRY_L during a Read transaction.
4. Tachyon becomes 32-byte aligned and additional Read transactions are queued. Specifically, when Tachyon transfers from a NON-READ8 transaction to a READ8 transaction, bus tenancy ends and a new stream starts in a new tenancy. Tachyon must start a new stream to start asserting PREFETCH_L, if appropriate.
5. The Internal Outbound Frame FIFO is full.

7.6.6 Data Extend Using PREFETCH_L

Tachyon uses the PREFETCH_L to notify the host that Tachyon will request the next sequentially addressed block of data on the next TBR_L[1] bus tenancy. This enables the host to read the data for the next memory address into its cache in order to minimize the read access delay on TSI.

Data Prefetch Rules

1. Tachyon only asserts PREFETCH_L if Tachyon will read the next 32-byte block.
2. Tachyon only asserts PREFETCH_L for READ8 (word) transactions.
3. Tachyon may assert PREFETCH_L for all or part of the bus tenancy, which can contain one or more transactions.
4. Tachyon initially asserts PREFETCH_L with TBR_L[1].
5. PREFETCH_L is valid during the address cycle of each transaction. The host must sample it during the address cycle to determine whether or not the Tachyon reads the next sequential block of 8 words.
6. PREFETCH_L does not guarantee that Tachyon requests the data immediately. PREFETCH_L only indicates that the next transaction from the TBR_L[1] channel is a READ8 from the next sequential address. Tachyon may begin a new bus tenancy using TBR_L[0] at any time and perform any number of non-prefetched reads on that channel, independent of the prefetching activity present on TBR_L[1].
7. Prefetching can wrap across bus tenancies, e.g., if Tachyon asserts PREFETCH_L during the final phase of a streamed transaction, then Tachyon reads the prefetched data the next time it requests the bus using TBR_L[1].
8. During a bus tenancy, if Tachyon deasserts PREFETCH_L in an address cycle, then Tachyon does not assert the signal again during that particular bus tenancy. Tachyon begins a new bus tenancy for each contiguous block of data.

7.6.7 Address Cycle

For the address cycle, the bus master asserts the following signals:

1. The address of the transaction on TAD [31..0]
2. The calculated value of parity on the PARITY line (parity can be even, odd, or disabled)
3. The type of transaction on TYPE [2..0]
4. AVCS_L to indicate the validity of the address

Tachyon ignores TAD [31..9] in an address cycle when it is the responder. However, if parity is enabled, Tachyon checks parity for these bits, so TAD [31..9] should be driven to known values with valid PARITY.

The host must drive the TYPE[ ] signal in the Address cycle of a host mastered transaction and may continue to drive it throughout the transaction. The host must release all TYPE [2..0] signals before granting TSI to Tachyon.

PTI 172645

7.6.8 Driving Bi-Directional Signals

Either the host or Tachyon can assert the AVCS_L and READY_L signals. The signals are both asserted low for just one cycle at a time. To operate one of these signals properly and guarantee timing margins, the host or Tachyon must drive the signal high for one cycle after it has been driven low. In the following cycle, it must be released (tri-stated). This is illustrated in all the timing diagrams provided in the examples in this chapter.

7.6.9 Write Transactions

A TSI Write transaction consists of an address cycle followed immediately by 1, 2, 4, or 8 data cycles. If the host masters the Write transaction, Tachyon asserts READY_L to acknowledge the transaction. If Tachyon masters the Write transaction, the host asserts READY_L to acknowledge the transaction . The responder may assert READY_L as soon as the cycle after the address cycle, or as late as after the last data cycle.

There is no timeout on TSI, so if Tachyon is the master, Tachyon waits to receive a READY_L or for a hard reset to occur. No other transactions are able to occur.

The Write transaction is not complete until the last data cycle or the acknowledgment cycle occurs, whichever occurs later. The master must stop driving TAD and PARITY the cycle after the last data cycle.

If Tachyon masters a Write transaction, it requests the bus using the TBR_L[0] line.

Only single word Write transactions are allowed when Tachyon is the responder.

7.6.10 Read Transactions

A TSI Read transaction consists of an address cycle, a turn cycle, zero or more wait cycles, the data (1, 2, 4, or 8 data cycles), and a recovery cycle. In the turn cycle, the master stops driving its TAD[ ], PARITY, and TYPE[ ] lines. After the turn cycle, the responder may drive data on the TAD[ ] lines and PARITY. READY_L must be asserted along with the first data cycle.

Only single word Read transactions are allowed when Tachyon is the responder.

For multiple word Read transactions, the remaining data must be transferred on consecutive bus cycles until all data is returned. If the responder is not ready to return Read data in the cycle after the turn cycle, it delays the assertion of READY_L, thereby inserting wait cycles on the bus until data is available. If many wait cycles are needed before the responder is ready to return the read data, the interface is locked during this time.

If the Read data cannot currently be made available to Tachyon, the host may assert RETRY_L for one cycle, instead of asserting READY_L. On the cycle after the host asserts RETRY_L, Tachyon releases TBR_L , waits for TBG_L to be deasserted for one cycle, and then reasserts TBR_L to re-arbitrate for the bus.

If Tachyon asserts the PREFETCH_L signal concurrently with the TBR_L[1] signal and the host then decides to retry the read transaction using the RETRY_L signal, then the PREFETCH_L signal will remain asserted even after Tachyon has given up the bus.

RETRY_L Rules

1. The host may assert RETRY_L only during DMA read operations.
2. RETRY_L and READY_L are mutually exclusive, i.e., the host may only assert one of the signals per Read transaction.
3. The host asserts RETRY_L only prior to the data cycle, which is signaled by the responder asserting READY_L. Once the data cycles have started, the host may not assert RETRY_L on that transaction.

7.6.11 TSI Transaction Window

A TSI transaction window starts when the master asserts AVCS_L. For a Read transaction, the transaction window ends one cycle after the last data cycle, i.e., the recovery cycle is the end of the transaction window. For a Write transaction, the transaction window ends the cycle after the responder asserts READY_L or the cycle after the last data cycle, whichever is later. If the host wanta to perform back-to-back Writes, it may do so without an idle cycle to separate the transactions. For example, the transaction window for the first transaction ends in the last data cycle or in the cycle in which the responder asserts READY_L, whichever is later.

7.6.12 TSI Transaction Ordering

Whenever more than one DMA channel is available between the host and a host bus adapter board, the possibility of an out of order transaction exists. For example, one block of data may arrive at the adapter board before an earlier requested block if the delays for the two channels are different.

Since Tachyon provides two DMA channels (via the two bus request lines), the following rules ensure that TSI transaction ordering is maintained.

Maintaining TSI Transaction Ordering

1. The order in which Tachyon initiates transactions is not altered if the host asserts RETRY_L. Thus, if the host aborts a Read operation by asserting RETRY_L, Tachyon immediately re-arbitrates for the bus in order to complete the aborted Read operation. No other Tachyon-mastered transactions, particularly other reads, occur until the original Read operation is completed.
2. Tachyon never asserts more than one bus request signal at a time. The two bus request lines, TBR_L[1:0] are mutually exclusive, i.e. the two are never asserted at the same time in an attempt to master two concurrent transactions.
3. Write transactions only occur on bus tenancies initiated with the TBR_L[0] (non-prefetch) signal, therefore Write transaction ordering is preserved.

Because Tachyon only initiates one Read and Write operation at a time, overall transaction ordering is maintained.

7.6.13 Endian-ness

Big Endian

Tachyon is big endian. On the address/data bus (TAD[ ]) and in all registers, bit 31 is the most significant bit and bit 0 is the least significant bit.

Little Endian

Refer to "A.2 Implementing Tachyon with Little Endian System" on page 319.

PTI 172647

7.6.14 Parity

Tachyon asserts the Error Out signal (ERROR_L) if a parity error is detected (if parity is enabled), or if an invalid transaction size is used on a host-mastered Read or Write. The host must determine what caused the error, log the error if necessary, and reset Tachyon to clear the error condition.

Parity is enabled or disabled via the Tachyon Configuration register. If enabled, Tachyon checks parity for both address and data. For a DMA Read that is mastered by Tachyon, if Tachyon detects a parity error, it asserts the ERROR_L signal two cycles after the cycle with bad parity. For a DIO Write transaction, if Tachyon detects a parity error, it asserts the ERROR_L signal one cycle after the cycle with bad parity for both address and data. For a DIO Read transaction, if Tachyon detects a parity error, it asserts the ERROR_L signal one cycle after the cycle with bad parity for address only. Refer to "7.6.15 Error Handling" on page 214.

Due to the delayed nature of data parity errors, assertion of ERROR_L may overlap subsequent transactions. After Tachyon asserts ERROR_L, the host should read the Tachyon Status register to determine the cause of the error. The host must reset Tachyon to clear the error condition and deassert ERROR_L.

EVEN parity is the total number of ones on the address/data bus (TAD [31..0]) and the parity bit (PARITY) equals an even number. For example, if there is an odd number of ones on TAD [31..0], PARITY is driven high for EVEN parity.

ODD parity is the total number of ones on the address/data bus (TAD[31..0]) and the parity bit (PARITY) equals an odd number. For example, if there is an odd number of ones on TAD [31..0], PARITY is driven low for ODD parity.

Tachyon generates and drives the parity bit when it drives address or data information on TSI. If a Tachyon-driven transaction has a parity error, the host is responsible for resetting Tachyon to clear the error. During cycles where data or address are not valid (as defined by AVCS_L, READY_L, and TSI timing rules), parity is not valid.

7.6.15 Error Handling

Assertion of ERROR_L

Assuming that parity is enabled, Tachyon asserts ERROR_L for error conditions which occur in different TSI bus cycles:

1. Address parity error to Tachyon: Asserted in the next cycle.
2. TYPE signal not equal to binary 000 or 100 during Address cycle to Tachyon: Asserted in the next cycle.
3. Write data parity error to Tachyon: Asserted in the next cycle.
1. DMA Read data parity error: Asserted in the second cycle after the cycle with error.
5. Internal parity error: Asserted in any cycle on TSI.

Tachyon asserts ERROR_L and sets the appropriate bits in the Fatal Error Status field of the Tachyon Status register. If another error occurs, the Fatal Error Status field does not change and the new error is not logged. To determine the cause of the error, the host should read the Tachyon Status register.

Once Tachyon asserts ERROR_L, it will keep ERROR_L asserted until the host resets Tachyon by asserting RESET_L or by writing 0x8000 0000 to the Tachyon Control register.

For any address parity error, TYPE error, or write data parity error for a transaction to Tachyon, Tachyon does not return READY_L. ERROR_L is the only indication of Tachyon's response.

Tachyon Mastership Under ERROR_L

While Tachyon asserts ERROR_L, Tachyon cannot master any more DMA transactions. Tachyon completes the current transaction, then terminates its bus tenancy, no matter what streaming or prefetch state it is in. If Tachyon is arbitrating for the bus when it asserts ERROR_L, Tachyon continues arbitrating, but it releases TBR_L as soon as it sees TBG_L asserted, and it does not proceed with the bus transaction.

Tachyon Fibre Channel Operation Under ERROR_L

When Tachyon asserts ERROR_L, the OSM freezes. An outbound frame in progress on the link at the time of ERROR_L assertion finishes transmitting, but no more frames follow. The link state machine is still active, so Tachyon still responds to some Fibre Channel link and loop primitive sequences, like OLS or NOS. The host has limited ability to manage the situation since many Tachyon functions are disabled by the error.

Tachyon Response Under ERROR_L

When Tachyon asserts ERROR_L, Tachyon enters a state where it can no longer master transactions on the TSI. Even though Tachyon can no longer master transactions on the TSI, all registers can be accessed. Tachyon does attempt to respond to slave accesses to status registers, which contain information about what type of error occurred. The slave accesses may fail if it experiences a fatal error.

If Tachyon asserts ERROR_L due to a host-mastered address parity error or host-mastered Write data parity error, Tachyon does not assert READY_L for that transaction.

If Tachyon is already asserting ERROR_L, and a host-mastered Write occurs, Tachyon does not accept data and asserts READY_L if no address or data parity errors exist. The only exception to this condition is the Tachyon Control register where the host can perform a soft reset. The Tachyon Control register accepts the data and asserts READY_L if no address or data parity error occurs.

If Tachyon is already asserting ERROR_L, and a host-mastered Read occurs, Tachyon asserts READY_L and returns data if no address parity error exists.

PTI 172649

When Tachyon is asserting ERROR_L, the host must provide a timeout of at least 20 clocks from its assertion of AVCS_L. If Tachyon does not assert READY_L within this time, the transaction had a parity error. The host may begin a new transaction.

Host Error Handling for DMA Reads

The host is responsible for detecting parity errors in addresses for DMA Reads initiated by Tachyon. When the host detects these errors, it has four response options:

1. The host does not respond to Tachyon. It just leaves the bus locked up. Tachyon waits for a response and does not issue any further transactions. The host asserts RESET_L to reset Tachyon.
2. The host returns READY_L and data, but does not grant mastership to Tachyon. Tachyon has no indication that there was a problem. Tachyon only knows that the transaction completed. If Tachyon is streaming, Tachyon continues streaming, and does not relinquish the bus. The master should continue responding until the stream is exhausted. Then the master can prevent another Tachyon tenancy by withholding TBG_L. If streaming is disabled, then Tachyon terminates its tenancy after the current transaction. However, Tachyon could send incorrect data on the Fibre Channel since the host received an incorrect memory address.
3. The host returns READY_L with data containing incorrect parity. Tachyon is forced to assert ERROR_L and stop mastering on the bus.
4. The host asserts RETRY_L, which forces Tachyon to relinquish the bus immediately. The host should not grant mastership to Tachyon afterwards.

Host Error Handling of DMA Writes

The host must check and handle parity errors on DMA addresses and data for DMA Writes initiated by Tachyon. The host has only two possible responses for these Write transactions, because Write transactions do not have RETRY_L support and the host does not return any data, as for DMA reads:

1. The host does not respond to Tachyon. It just leaves the bus locked up. Tachyon waits for a response and does not issue any further transactions. The host asserts RESET_L to reset Tachyon.
2. The host returns READY_L and data, but does not grant mastership to Tachyon. Tachyon has no indication that there was a problem. Tachyon only knows that the transaction completed. If Tachyon is streaming, Tachyon continues streaming, and does not relinquish the bus. The master should continue responding until the stream is exhausted. Then the master can prevent another Tachyon tenancy by withholding TBG_L. If streaming is disabled, then Tachyon terminates its tenancy after the current transaction.

7.6.16 Reset

The RESET_L signal performs a synchronous reset of Tachyon. The host asserts RESET_L for at least 10 SCLK cycles to ensure proper operation. Tachyon starts in an arbitrary state and RESET_L should be asserted while power is being applied. This forces Tachyon into a known state.

Following a reset, Tachyon is in an IDLE state. The power-on value of all Tachyon internal registers is defined in "5.1 Register Overview" on page 111.

7.6.17 Arbitration

Since only two devices are present on TSI, arbitration is not that complex. If Tachyon requires control of TSI, Tachyon asserts one of the request (TBR_L[ ]) lines. TBR_L[0] indicates reads and writes using the non-prefetched channel. TBR_L[1] indicates reads using the prefetch channel, except for the following situations:

1. At the beginning of a block of prefetched data TBR_L[1] is used, even though no data has been prefetched yet.
2. At the end of a prefetched block within a sequence the next A/L pair is read via the prefetch channel, even though the actual A/L pair data is not prefetched. The prefetching channel must be able to process these cases, in addition to the normal prefetched cases.

The two TBR_L signals are mutually exclusive, i.e., the two signals are never asserted simultaneously. Each TBR_L signal is guaranteed to be de-asserted for at least one clock cycle prior to other TBR_L signal being asserted.

Regardless of the type of request signal (TBR_L[0] or TBR_L[1], the host asserts the grant line (TBG_L) to grant the bus to Tachyon. The host must assert TBG_L for a minimum of one clock cycle. Afterwards, TBG_L may be de-asserted at any time without effect. To avoid latency, the host can de-assert the grant line as soon as possible.

After the host asserts TBG_L, Tachyon begins a transaction one cycle later. Tachyon begins the tranaction by asserting AVCS_L and starting the address cycle.

After Tachyon begins the transaction, it holds TBR_L[ ] asserted until at least the later of:

1. The last data cycle of the transaction
2. The cycle after READY_L or RETRY_L is asserted. The host can assert its READY_L at the conclusion of the transfer (on the cycle following the last data cycle) as a confirmation that the transfer is complete or to avoid latency, the host may elect to assert READY_L sooner.

Tachyon keeps the TBR_L signal asserted until it relinquishes its bus tenancy following its last transaction.

Tachyon stops driving all lines by the cycle when Tachyon de-asserts the TBR_L[ ] signal.

Once the host de-asserts TBG_L, it must not re-assert until after both of the TBR_L[ ] signals are de-asserted and after Tachyon asserts one of them again.

PTI 172651

7.6.18 Interrupts

When Tachyon places an entry into the IMQ, Tachyon generates an interrupt to notify the host, except for the following conditions:

1. The host requests Tachyon not to generate an interrupt for an outbound completion message.
2. A previous entry in the IMQ has not yet been consumed by the host. In this case, Tachyon does not generate another interrupt until the IMQ Consumer Index is updated by the host.

When Tachyon generates an interrupt, the interrupt signal, INT_L, is asserted for one TSI clock cycle. INT_L is then de-asserted and remains de-asserted until Tachyon generates the next interrupt.

Interrupts may not be generated when Tachyon is the owner of TSI, i.e., INT_L is always de-asserted when either TBR_L signal is asserted.

When software services interrupts, maintain the following order:

1. Software receives interrupt.
2. Software reads IMQ entry(ies).
3. Software re-enables the hardware interrupts (e.g. resets a latch which detects Tachyon interrupts).
4. Software updates IMQ Consumer Index, which re-enables interrupts within Tachyon.

WARNING	If Step 4 is performed before Step 3 above, a race condition occurs and a subsequent interrupt may be missed.

If the host software decides not to consume all IMQ entries before updating the IMQ Consumer Index, Tachyon immediately issues another interrupt to ensure that messages are not left on the IMQ. The normal mode of operation is to read all entries on the IMQ before updating the IMQ Consumer Index.

PTI 172652

7.6.19 Host Interface Design Notes

Arbitration

1. Host transactions may begin one delay cycle after Tachyon deasserts TBR_L[ ].
2. Tachyon may deassert TBR_L[ ] for as little as one cycle, then re-arbitrate if TBG_L is de-asserted.
3. If the host asserts TBG_L when Tachyon is asserting TBR_L[ ], Tachyon starts a transaction. Give a grant no sooner than one cycle before the host is finished with its transaction (for reads, at the start of the recovery cycle; for writes, in the last data cycle or the cycle after READY_L is asserted, whichever is later).
4. Do not anticipate the number of cycles from TBG_L that Tachyon asserts AVCS_L; wait for it.
5. Do not anticipate the number of cycles after READY_L that Tachyon de-asserts TBR_L[ ].
6. Tachyon bus tenancy can span up to 64 transactions. Individual applications can limit this number by programmable configuration.

Host-Mastered Transactions

1. Do not anticipate the number of cycles Tachyon takes to assert READY_L.
2. Tachyon checks parity on all 32 TAD[ ] signals during address cycles.
3. The responder may assert READY_L the cycle after Address for Writes.
4. A host-mastered transaction with address or TYPE error does not complete with READY_L.

Tachyon-Mastered Reads

1. Check PREFETCH_L in the Address cycle (optional).
2. Note the short time between streamed Reads.
3. The host may assert RETRY_L only during Reads. This forces Tachyon to terminate its tenancy.
4. Tachyon reads a prefetched address during the next transaction that it asserts TBR_L[1].

Tachyon-Mastered Writes

1. Responder may assert READY_L before the last data cycle.
2. Note the short time between streamed Writes.
3. There is no mechanism to stop a Write stream, unless the host asserts RESET_L. Withholding READY_L only suspends the Write stream.

PTI 172653

7.7 TSI Functional Waveforms

TSI Functional Waveform Signal Notes

- De-assertion of the TBR_L[ ] signal by Tachyon at the end of a transaction is determined by which internal Tachyon module is accessing the bus and the current state of streaming.
- The re-assertion delay of TBR_L [ ] depends on the state of TBG_L when TBR_L[ ] is de-asserted and other transactions that are queued in Tachyon.
- Transactions can start (assertion of AVCS_L) no sooner than one clock cycle after the de-assertion of the TBG_L signal.
- When the host or Tachyon drives either AVCS_L or READY_L; it asserts (driving low) the signal for one clock cycle, actively de-asserts (driving high) the signal for one cycle, and then tri-states the signal.
- The host interface must not attempt to drive either AVCS_L or READY_L when Tachyon is driving these signals either high or low.
- To meet timing parameters, it is recommended that the host actively de-assert AVCS_L and READY_L after driving them. All waveforms in this document show the host behaving in this manner.

TSI Functional Waveform Graphic Notes

- A tri-stated (high-impedance) signal is indicated by a dashed line between high and low.

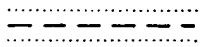

- An undefined signal, e.g., wait states, is indicated by a shaded area.

- A parity signal is indicated by a white boxed area.

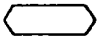

- A valid bus value is indicated in boxed areas.

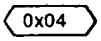

Note The following TSI waveform examples depict realistic possibilities for Tachyon Revision 2, however, the waveform examples do not reflect the exact timing of every possible transaction, nor for every possible version of Tachyon.

Note The following TSI functional waveform examples are meant to show functionality and are not intended to show any timing information. For timing information, refer to "7.8 TSI Timing Requirements" on page 244.

PTI 172654

7.7.1 Slave Reads and Writes

In these examples of Slave Reads and Writes, the host asserts three host-mastered, or slave, transactions, a slave WRITE1, a slave READ1, and an illegal slave READ4, shown in the following three figures.

During slave Write or Read transactions, only single word transactions are allowed. If the host requests a slave transaction of other than one word, Tachyon asserts ERROR_L.

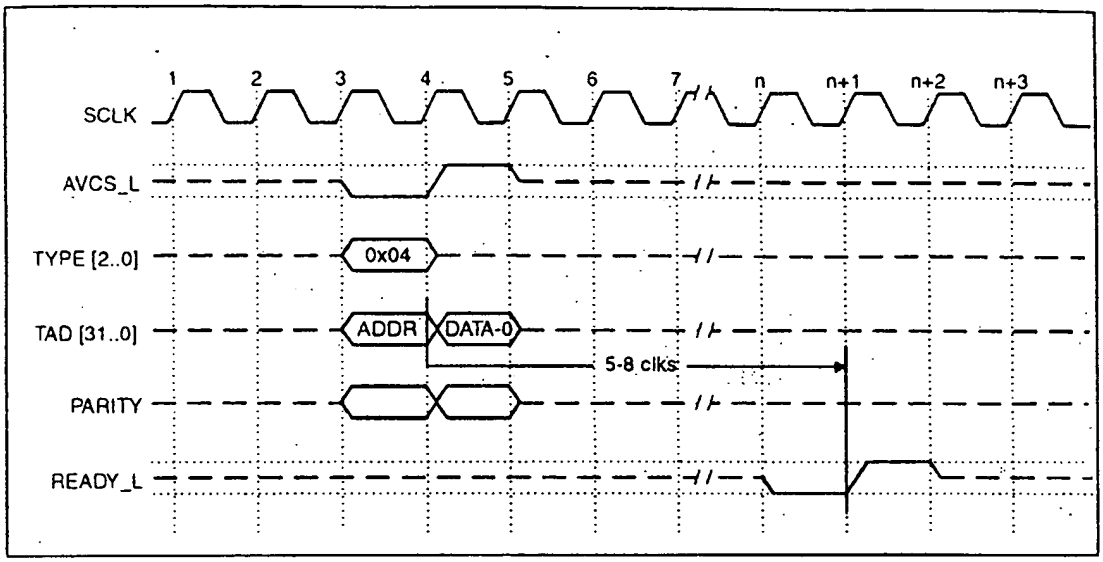

Figure 7.2 Slave WRITE1

This is a slave WRITE1 transaction, indicated by TYPE [2..0] = 0x04. The host initiates the transaction by strobing AVCS_L and driving the TYPE [2..0], TAD [31..0], and PARITY signals. The following clock cycle is the data cycle (DATA-0), with no wait states possible. Tachyon then ends the transaction by asserting READY_L as an acknowledgment.

The delay from the sampling of AVCS_L to the assertion of READY_L for a Write depends on which register is being written and the Fibre Channel speed selected. For Writes of non-Frame Manager registers, the delay is five clock cycles. If the fastest TSI clock (40 MHz) and the slowest Fibre Channel speed (266 MBaud) is selected , then the delay is between five and eight clock cycles.

Design should depend on the host waiting for the assertion of READY_L, rather than waiting for a specific number of clock cycles to occur. This ensures compatibility with future revisions of Tachyon.

PTI 172655

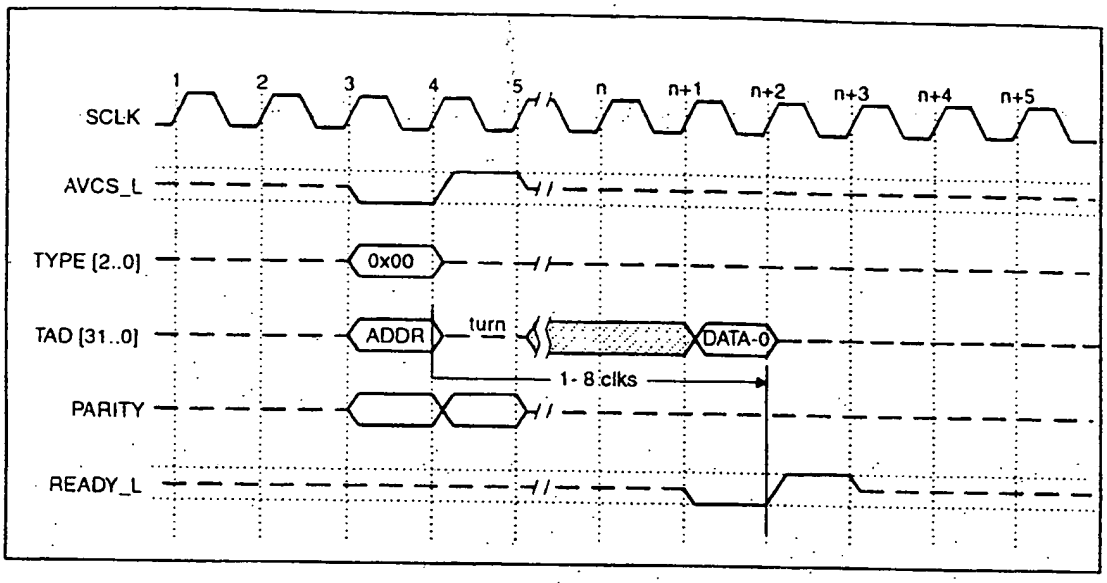

Figure 7.3 Slave READ1

This is a slave READ1 transaction, indicated by TYPE [2..0] = 0x00. Following the assertion of AVCS_L on any Read transaction, the next clock cycle is defined as a "turn" cycle. A turn cycle gives a finite period of time for the initiator of the transaction to stop driving the address on TAD [31..0] before the responder begins driving data on TAD [31..0]. After the turn cycle, Tachyon generates wait states before the data cycle (DATA-0) begins.

The delay from the sampling of AVCS_L to the assertion of READY_L for a Read depends on which register is being read and the Fibre Channel speed selected. For Reads of non-Frame Manager registers, the delay is one clock cycle. For Frame Manager registers, the delay is between one and eight clock cycles. If the fastest TSI clock (40 MHz) and the slowest Fibre Channel speed (266 MBaud) is selected, then the delay is between one and eight clock cycles.

Note	Parity is only depicted in these first two waveform examples, but may be used for all DMA transactions.

PTI 172656

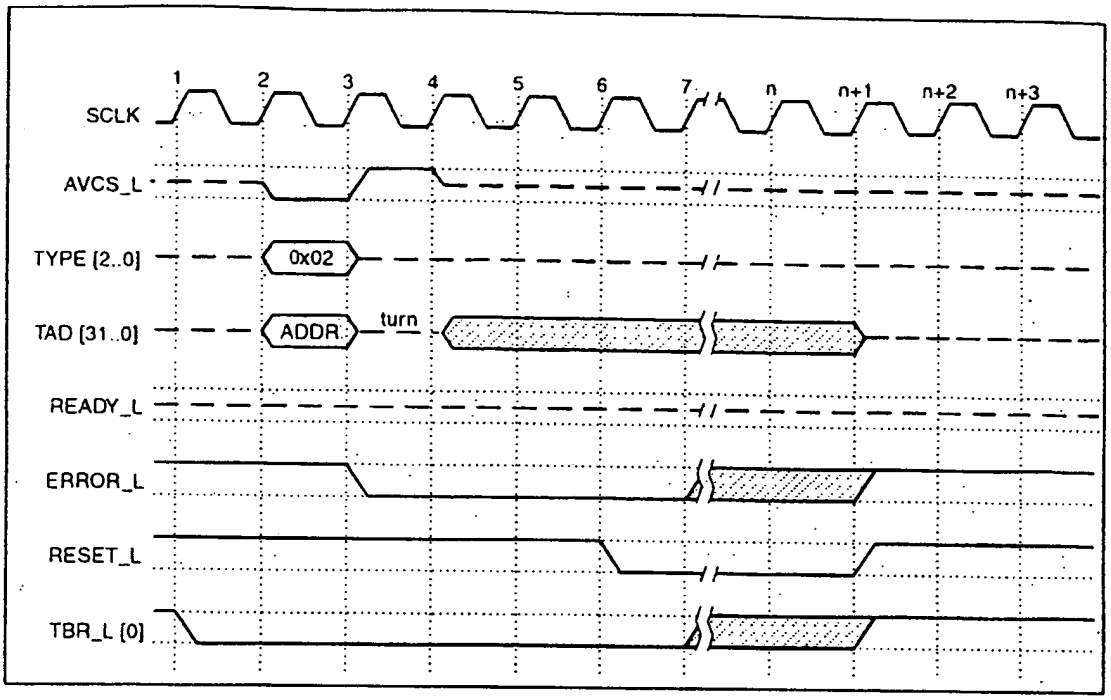

Figure 7.4 RESET_L of an Illegal Slave READ4

This third slave transaction demonstrates the effect of performing a slave transaction of other than one word. The host initiates an illegal slave READ4 (four words) transaction causing Tachyon to assert ERROR_L.

The host resets Tachyon by asserting RESET_L. RESET_L must remain asserted for a minimum of ten clock cycles. When the host resets Tachyon, asserted signals are released before the end of the RESET_L assertion. In this case, the asserted signals, ERROR_L, TBR_L [0], and TAD[31..0] are released.

TBR_L [0] is completely independent of the slave transaction. However, if it is asserted prior to an illegal slave transaction, it remains asserted while ERROR_L is asserted, until the reset of Tachyon occurs.

PTI 172657

7.7.2 DMA Writes

In this example of a DMA Write, Tachyon has 28 bytes of data to write to host memory. Tachyon writes this data with a series of three transactions; a WRITE4, a WRITE2, and a WRITE1, shown in the following three figures.

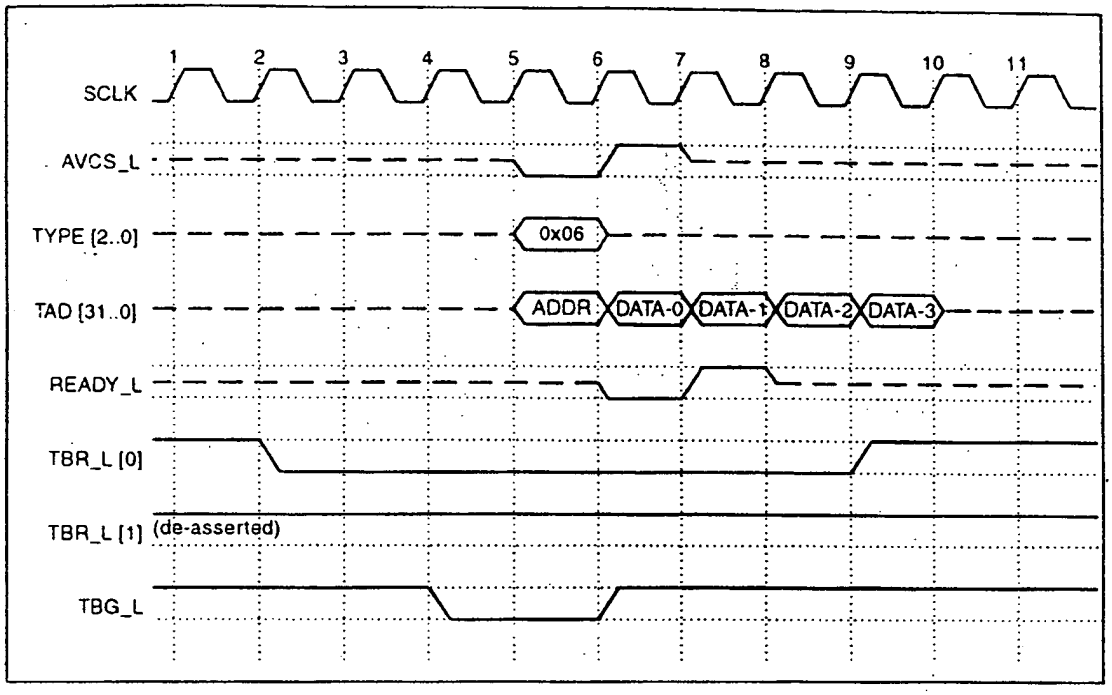

Figure 7.5 DMA Write, WRITE4 Transaction

This a WRITE4 transaction indicated by TYPE [2..0] = 0x06.

PTI 172658

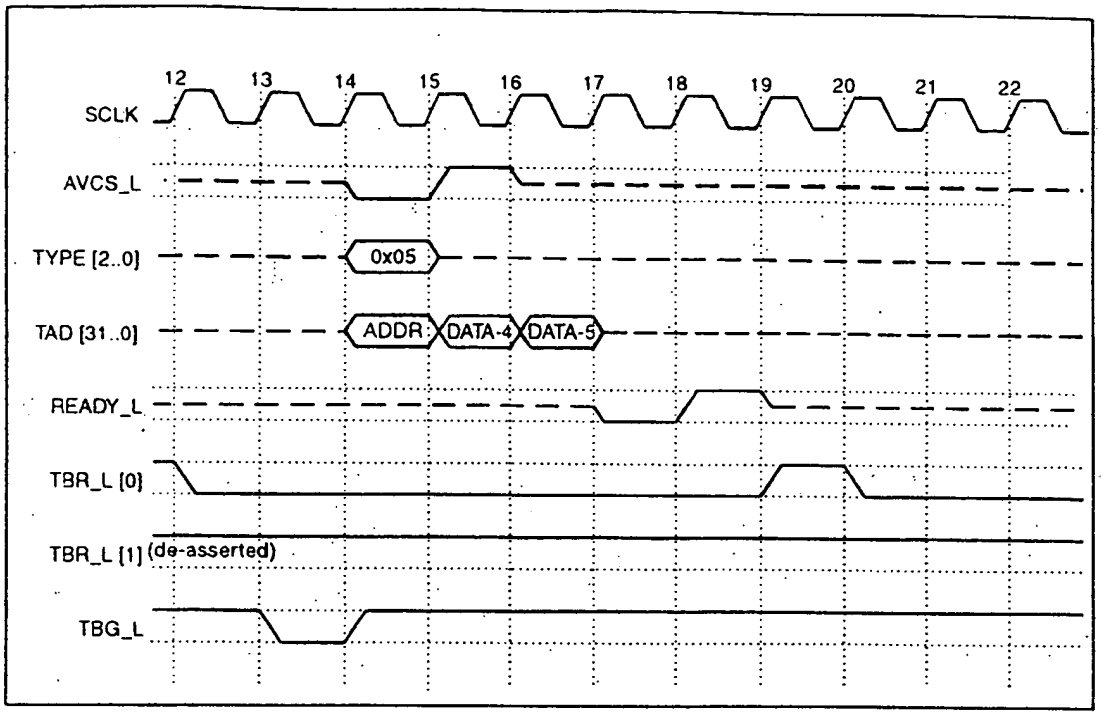

Figure 7.6 DMA Write, WRITE2 Transaction

This is WRITE2 transaction indicated by TYPE [2..0] = 0x05.

Since TBG_L is already de-asserted when TBR_L[0] is de-asserted, Tachyon can assert TBR_L [0] again without waiting.

PTI 172659

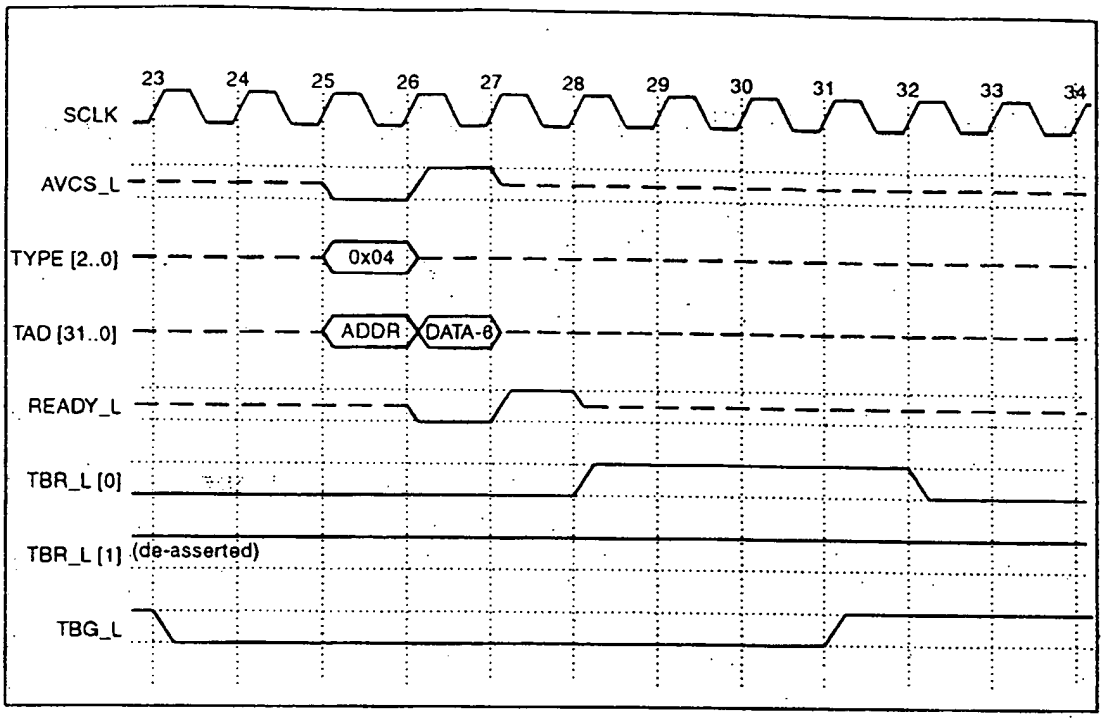

Figure 7.7 DMA Write, WRITE1 Transaction

In this example, the de-assertion of TBG_L is delayed and Tachyon cannot arbitrate for the next bus tenancy (assert TBR_L[0]) until TBG_L is de-asserted.

PTI 172660

7.7.3 DMA Write Streaming

When streaming, i.e., performing multiple transactions per bus tenancy, a new transaction cannot begin until two clock cycles after the last data cycle has finished and READY_L has been asserted.

In this example of a DMA Write Stream, Tachyon has 28 bytes of data to write to host memory. Tachyon writes this data with a series of three write transactions; a WRITE4, a WRITE2, and a WRITE1 shown in the following two figures.

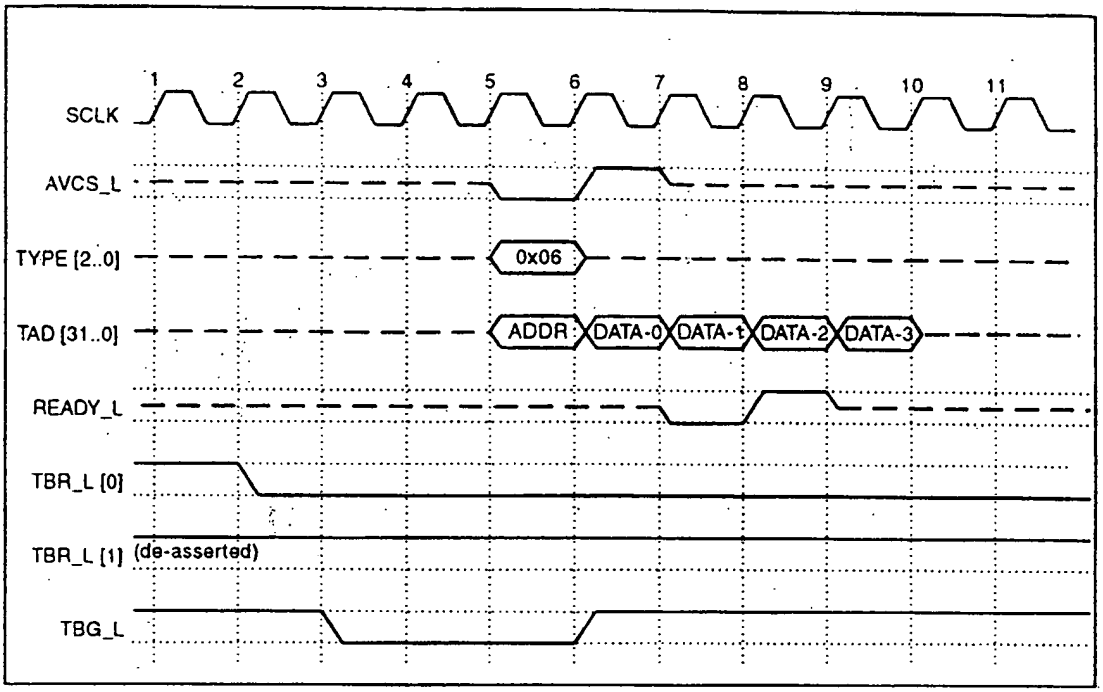

Figure 7.8 DMA Write Stream, WRITE4 Transaction

In the figure above, READY_L is asserted during the data cycles of the first transaction (WRITE4), therefore, the second transaction (WRITE2) can begin two cycles after the last data cycle of the first transaction. Refer to "Figure 7.9 DMA Write Stream, WRITE2 and WRITE1 Transactions" on page 227.

PTI 172661

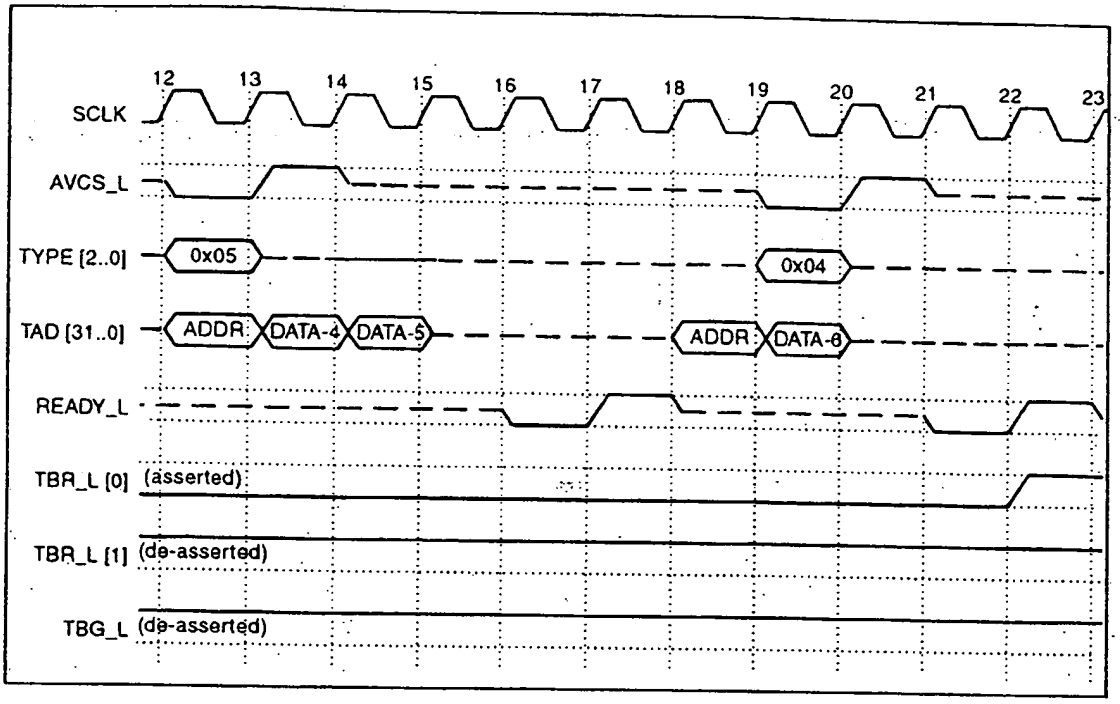

Figure 7.9 DMA Write Stream, WRITE2 and WRITE1 Transactions

This figure shows the second (WRITE2) and third (WRITE1) transactions of the DMA Write Stream.

The third transaction cannot occur two cycles after the last data cycle of the second transaction because it must wait until the host asserts READY_L.

After the host asserts READY_L, Tachyon ends the bus tenancy by de-asserting TBR_L[0].

PTI 172662

7.7.4 DMA Reads

In this example, two of the four DMA read transactions, a READ2 and a READ4, are shown in the following two figures.

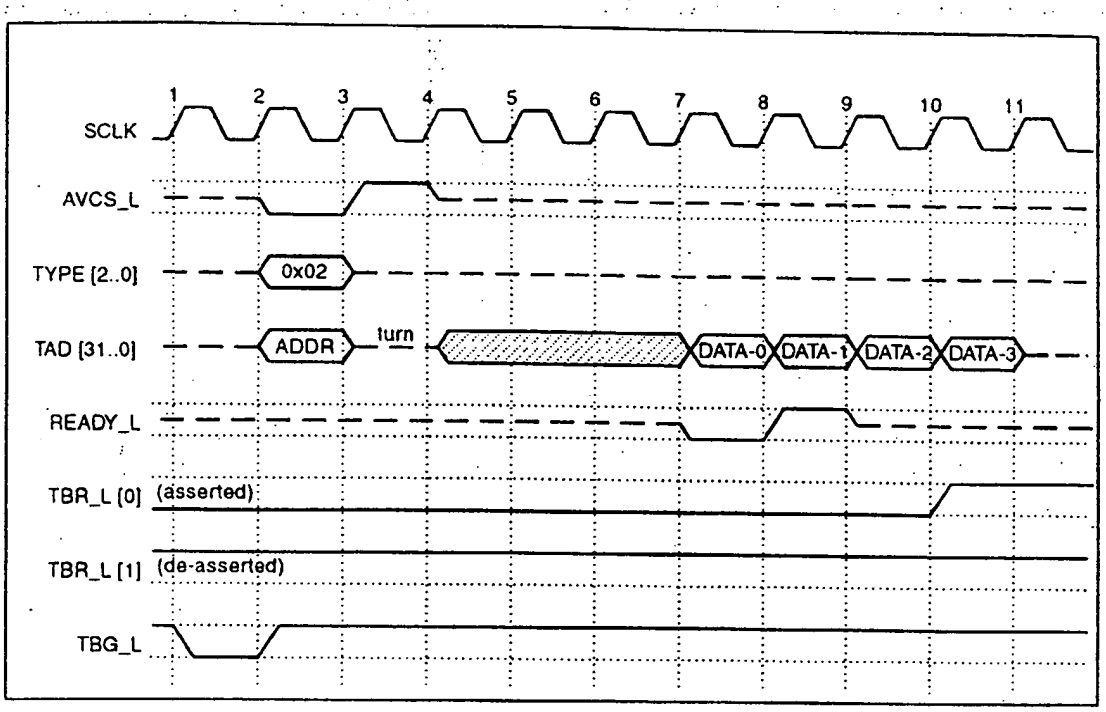

Figure 7.10 DMA Reads, READ4 Transaction

The bus request signal (TBR_L[0] in this example) is de-asserted just before the last data cycle.

If streaming is enabled, the de-assertion of TBR_L may occur just before the last data cycle or up to four cycles after that.

PTI 172663

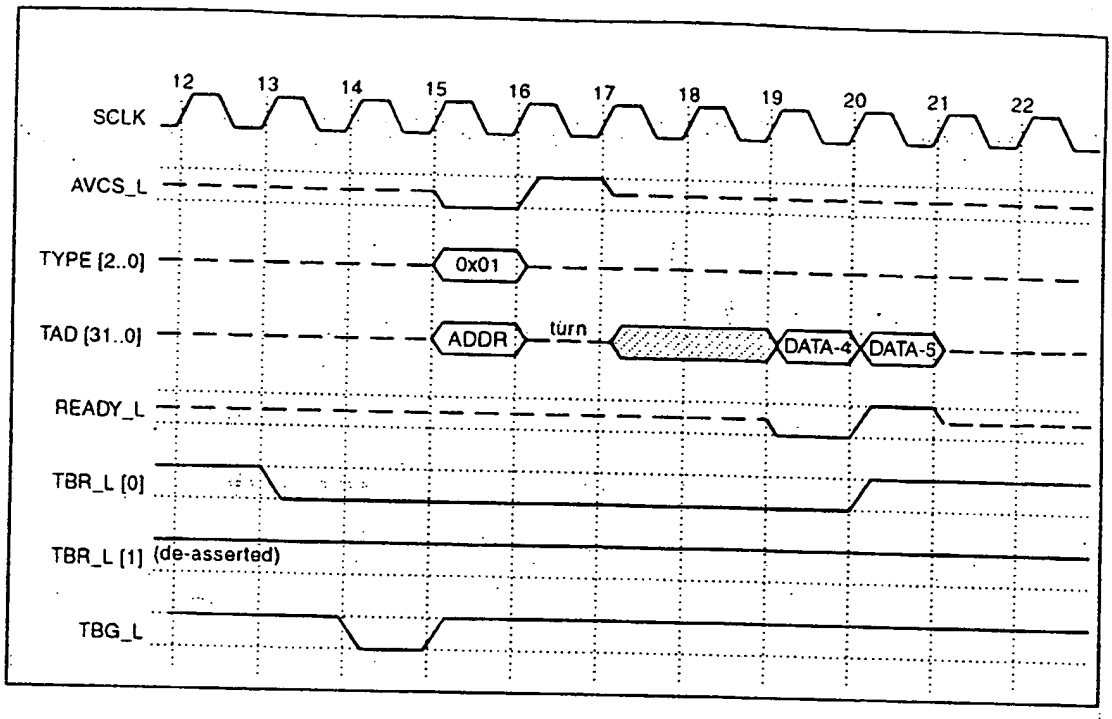

Figure 7.11 DMA Reads, READ2 Transaction

The above figure shows a READ2 transaction indicated by TYPE [2..0] = 0x01.

PTI 172664

7.7.5 DMA Read Prefetching

In this example, the basic behavior of prefetch-indication is depicted in the following two transactions, READ8 at ADDR-1 and READ8 at ADDR-2, shown in the following two figures.

In this example, streaming is not enabled.

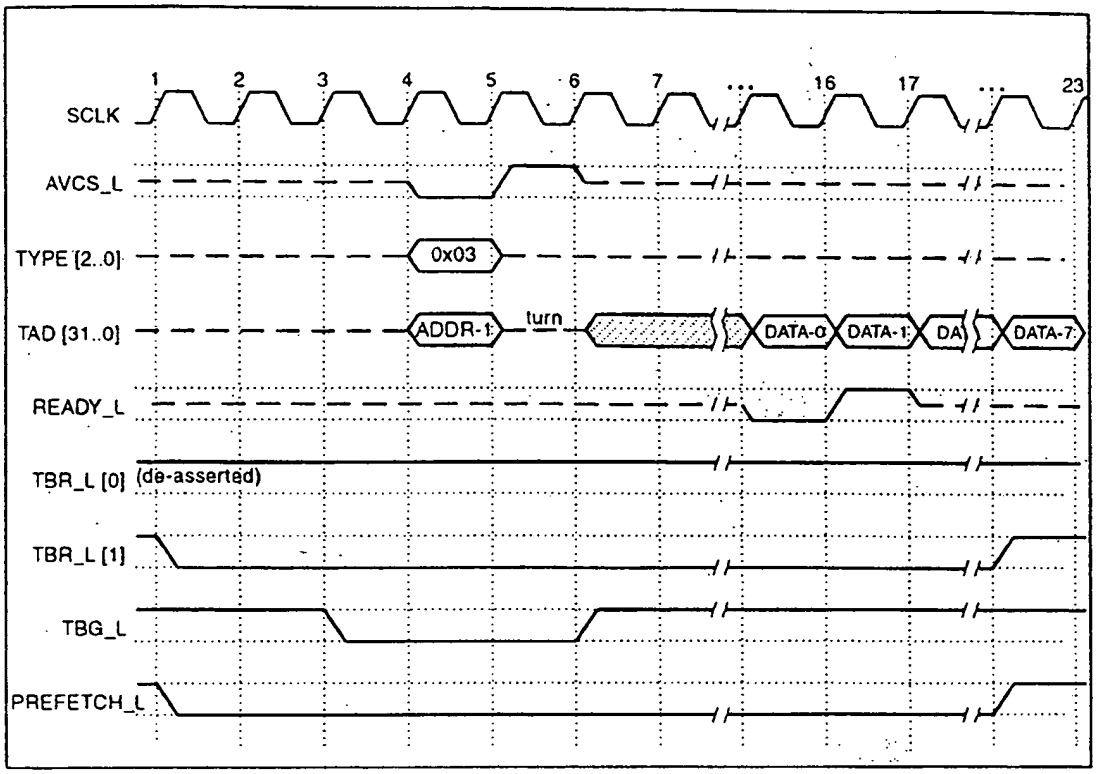

Figure 7.12 DMA Read Prefetching, READ8 at ADDR-1

The READ8 at address ADDR-1 is the first transaction of a block of 64 bytes that is read from contiguous locations. As the bus is requested for this transaction, PREFETCH_L is also asserted, indicating that the next Read on this channel is sequential to the current 32 bytes. A large number of wait states are needed, due to the latency of the host memory, in this example.

As the host interface retrieves the Read data for the first transaction memory, it may want to act upon the prefetch-indication that is given during the first transaction and valid during the address cycle. If the host wants to act upon the prefetch-indication, it should retrieve and buffer the next 32 bytes sequential to the first transaction's data, since it is guaranteed that Tachyon requests this data on the next TBR_L [1] bus tenancy.

PTI 172665

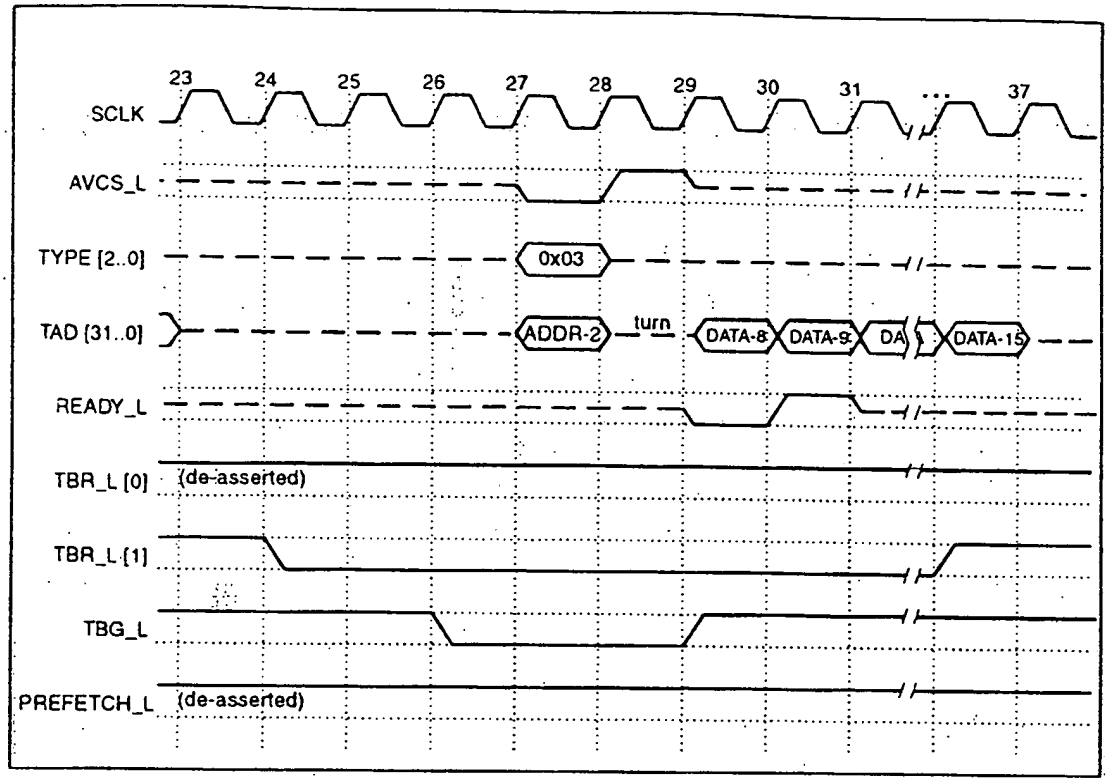

Figure 7.13 DMA Read Prefetching, READ8 at ADDR-2

When the second transaction has started, the host interface has already successfully prefetched the next block of 32 bytes, so the data cycles for the second transaction may start as soon as one cycle after the address cycle, as shown above. A turn cycle is always required immediately following the address cycle in any Read transaction.

In this example, ADDR-2 = ADDR-1 + 32 bytes.

PTI 172666

7.7.6 DMA Read Channels

When outbound data (headers and EDBs) are DMAed from host memory, Tachyon always requests the host bus using the prefetch read channel, TBR_L[1]. For all other data that is read from host memory, such as ODBs or inbound buffer blocks, Tachyon requests the host bus using the non-prefetch read channel, TBR_L[0].

This example shows the interaction between the two different read channels, TBR_L[0] and TBR_L[1], using the following transactions:

1. A READ8 at address ADDR-1 using the prefetch channel
2. A non-contiguous READ8 at address ADDR-X using a non-prefetch channel
3. A final READ8 at address ADDR-2 using the prefetch channel, where ADDR-2 = ADDR-1 + 32 bytes

These transactions are shown in the following three figures.

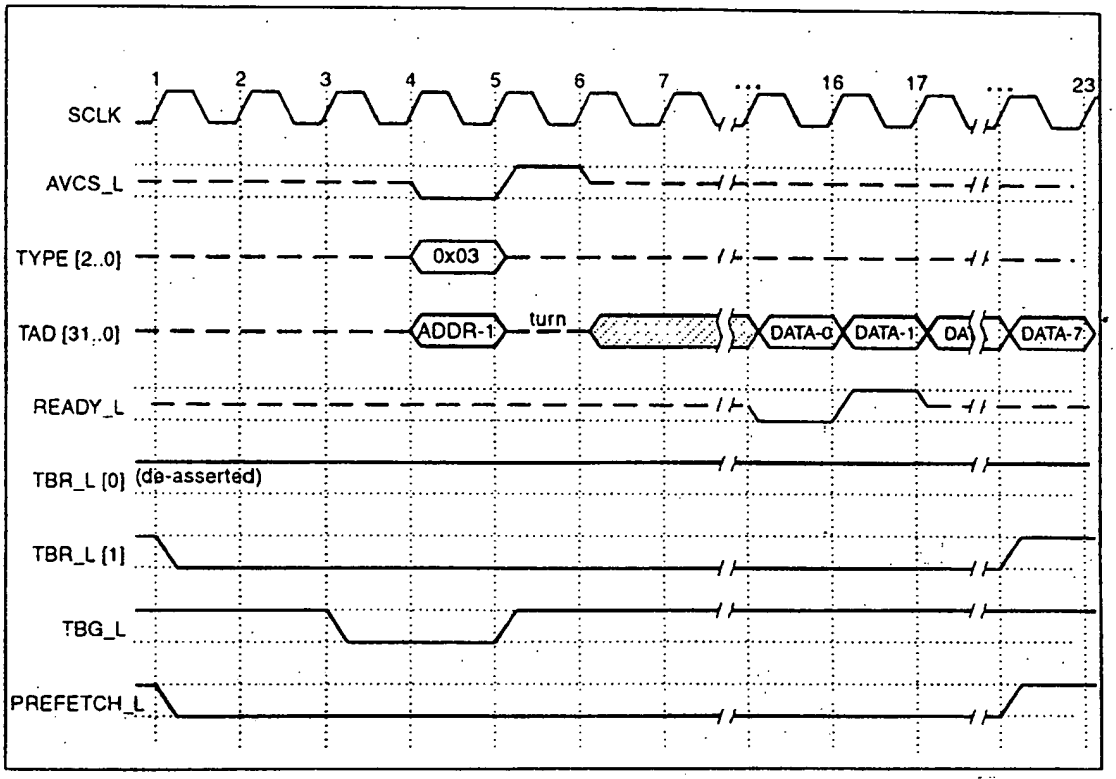

Figure 7.14 DMA Read Channels, READ8 at ADDR-1

This is a READ8 transaction at ADDR-1, using the prefetch channel, TBR_L [1].

As in the previous example, prefetching is indicated during the first transaction, so the next transaction on the TBR_L[1] channel (the READ8 at ADDR-2) is a Read to an address sequential to the current read data.

PTI 172667

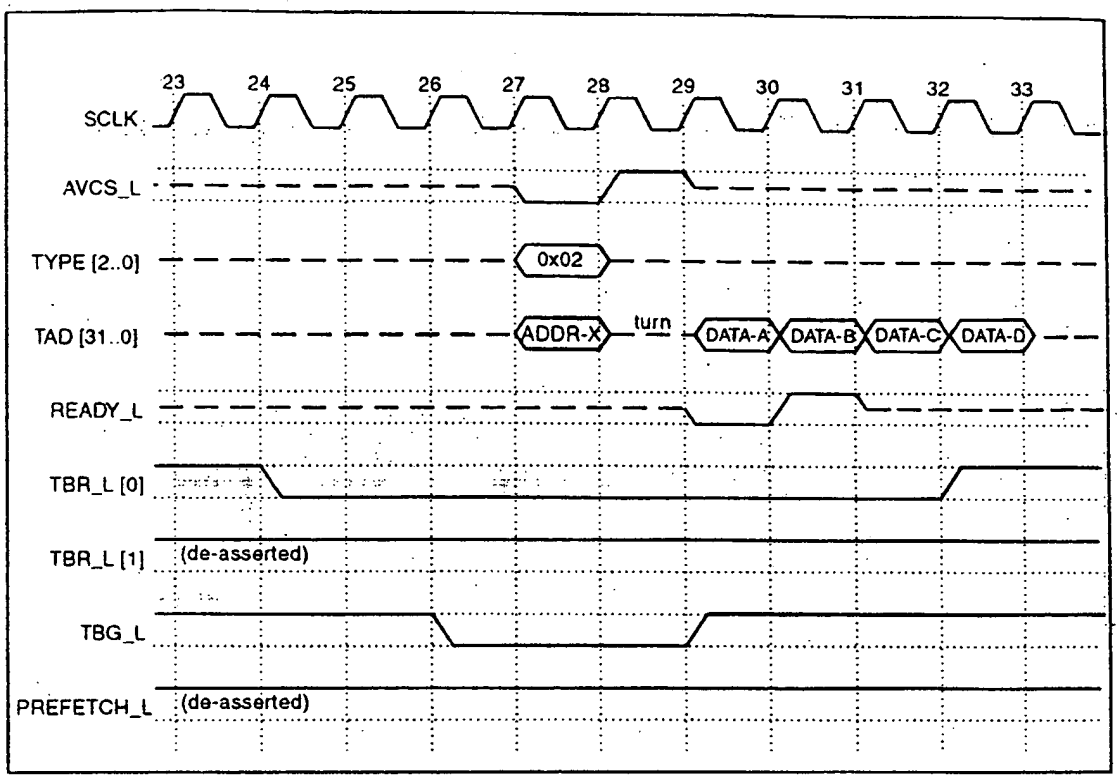

Figure 7.15 DMA Read Channels, READ8 at ADDR-X

The second transaction (a READ8 at ADDR-X) is non-sequential to the first and is read using the non-prefetched channel, TBR_L[0].

PTI 172668

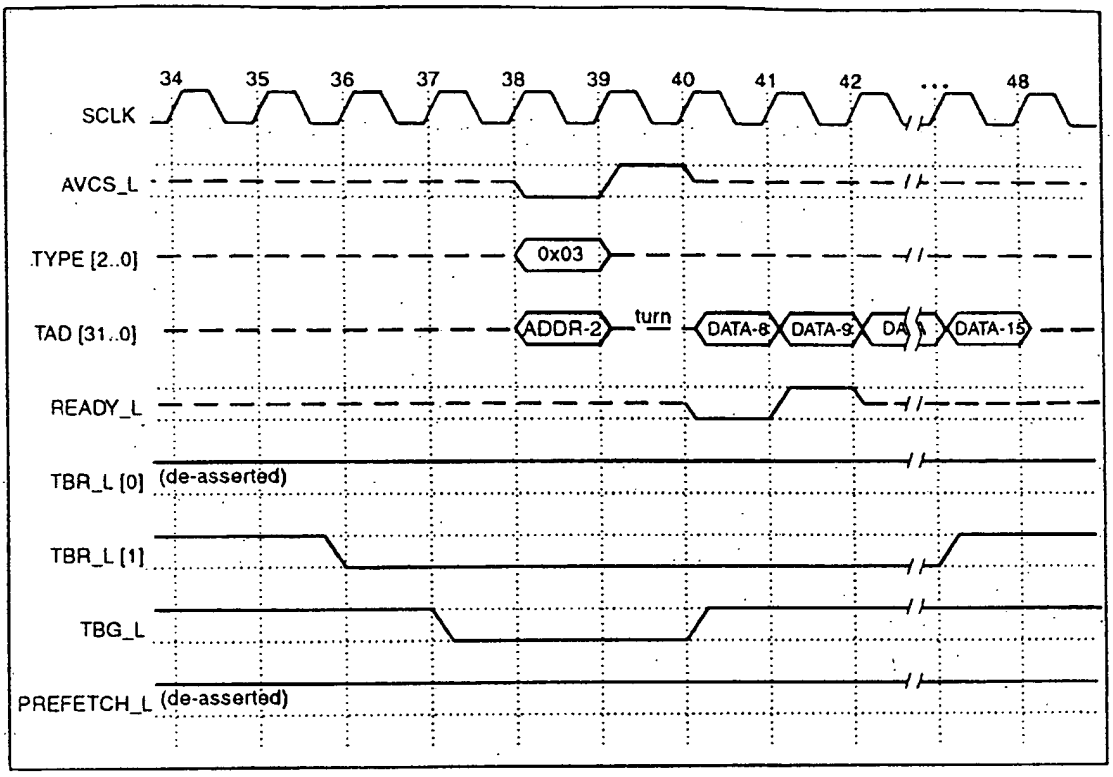

Figure 7.16 DMA Read Channels, READ8 at ADDR-2

Like the first READ8 transaction, the third READ8 transaction uses the prefetch read channel, TBR_L[1].

Because this transaction was prefetch-indicated during the first transaction, the data was buffered on the host interface and is readily available when requested. Also, because this was the last of the 64-byte block to read, PREFETCH_L is not asserted during the address cycle of the third transaction.

In this example, ADDR-2 = ADDR-1 + 32 bytes.

PTI 172669

7.7.7 Streamed Block Reads

In this example, Tachyon reads a block of 96 bytes from host memory. Tachyon reads this block in three 32-byte transactions; READ8 at ADDR-1, READ8 at ADDR-2, and READ8 at ADDR-3 shown in the following three figures.

Tachyon has been configured to allow at least four Read transactions per bus tenancy, and therefore, can stream the first block of data, i.e., Tachyon can read all 96 bytes in one bus tenancy.

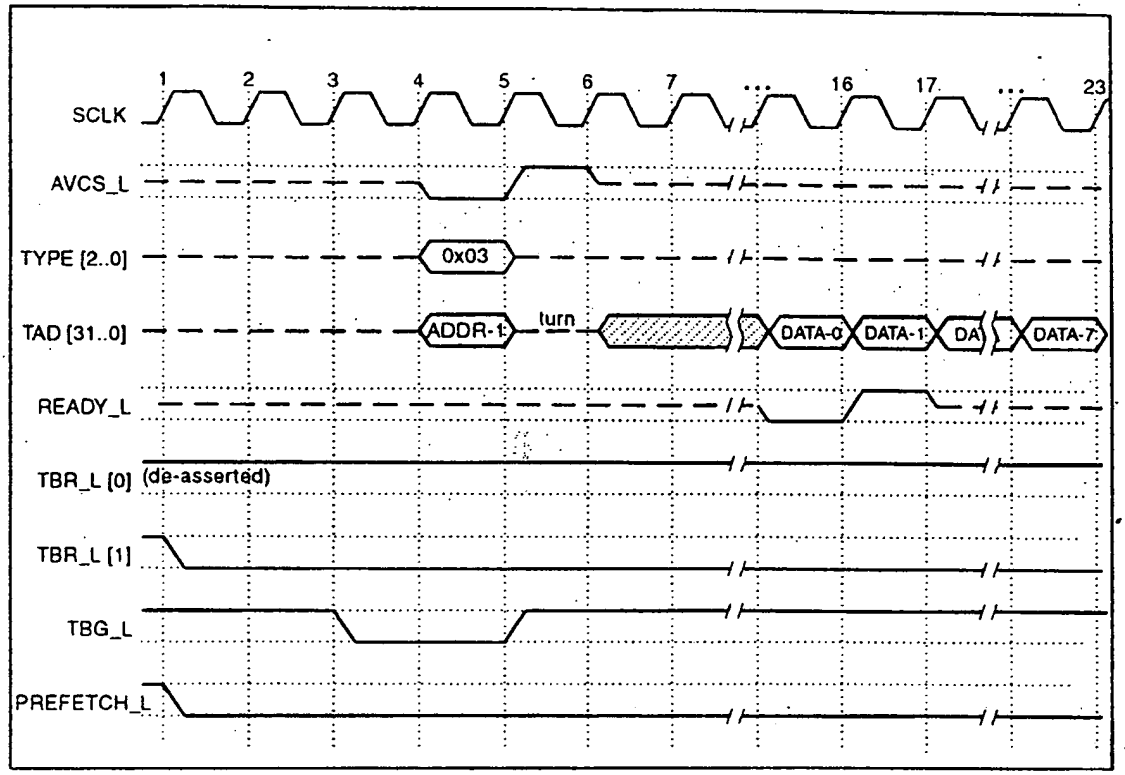

Figure 7.17 Streamed Block Read, READ8 at ADDR-1

Address ADDR-1 incurs a large number of wait states due to latency of the host memory.

PTI 172670

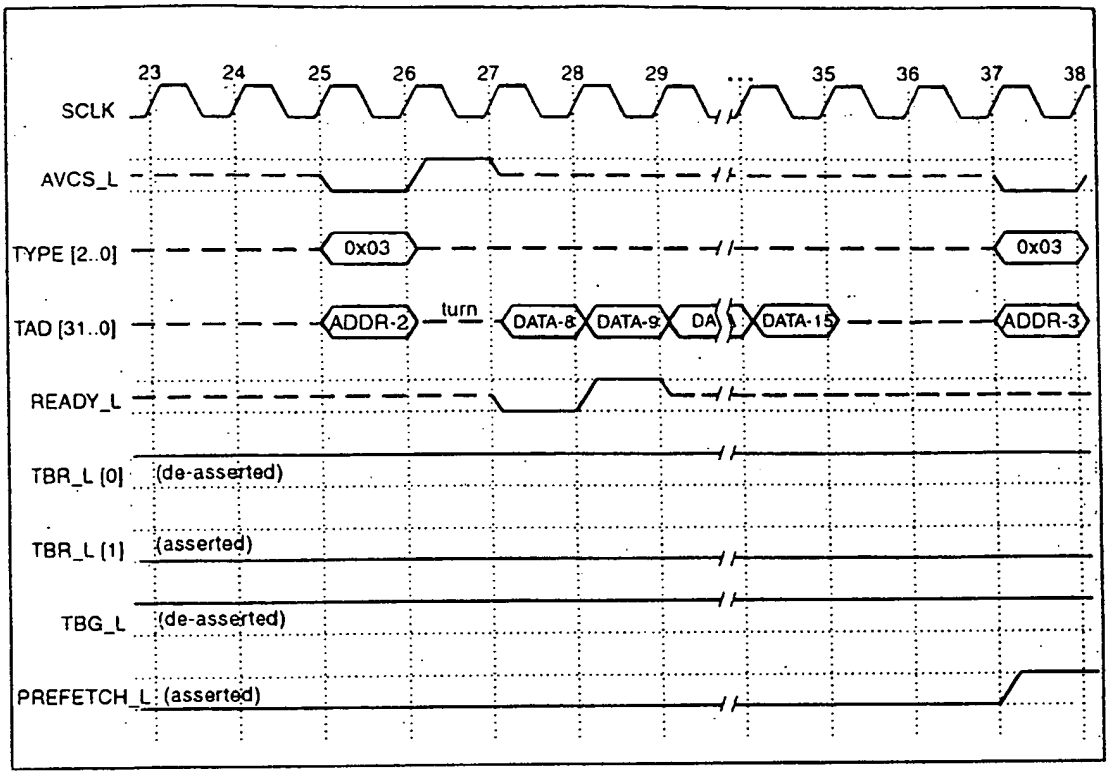

Figure 7.18 Streamed Block Read, READ8 at ADDR-2 and ADDR-3

Since PREFETCH_L is sampled during the address phase of the first transaction (ADDR-1), the data associated with address ADDR-2 can be prefetched and ADDR-2 does not incur many wait states while the read data returns.

The data associated with address ADDR-3 can also be prefetched, since PREFETCH_L is sampled during the address phase of the first transaction (ADDR-1). ADDR-3 does not incur many wait states while the read data returns. However, since the ADDR-3 transaction represents the last data in this block, PREFETCH_L is de-asserted prior to the address phase of the third transaction.

PTI 172671

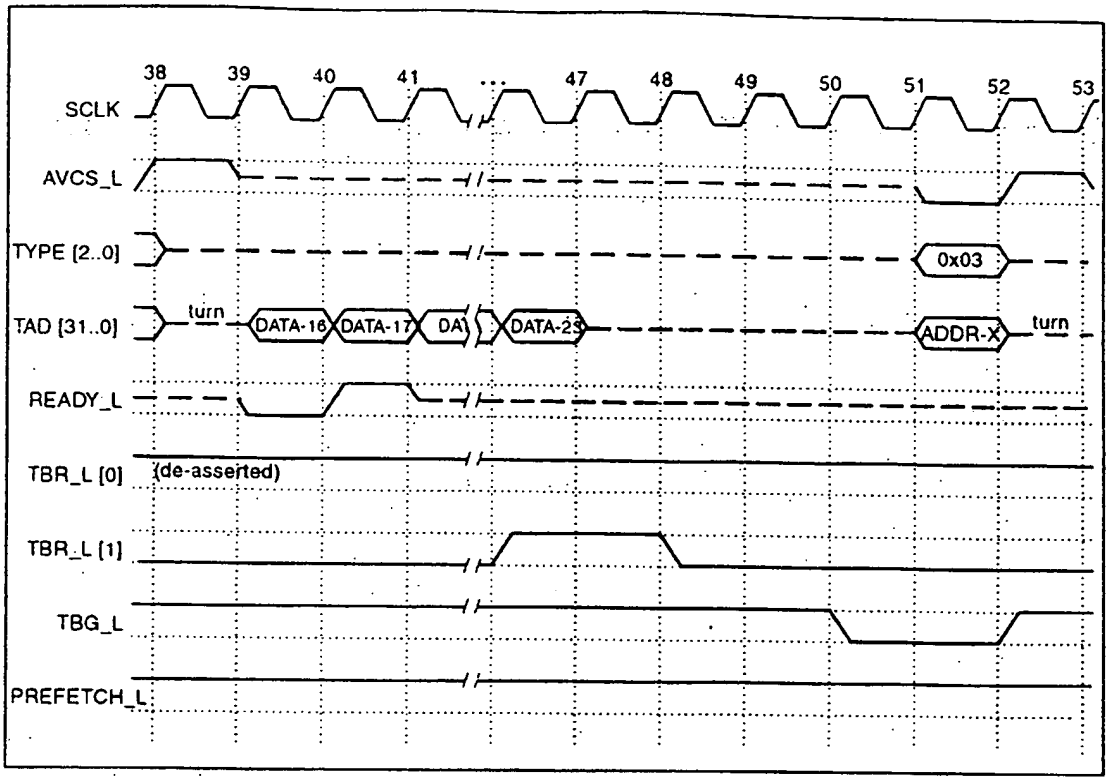

Figure 7.19 Streamed Block Read, READ8 at ADDR-3 and ADDR-X

On the last data cycle of the third transaction (ADDR-3), Tachyon ends the bus tenancy by de-asserting TBR_L[1].

Tachyon always begins a new bus tenancy when starting on a new block of data as shown here starting at address ADDR-X.

ADDR-X does not equal ADDR-3 + 32 bytes.

PTI 172672

7.7.8 Prefetching Across Back-To-Back Read Streams

In this example, Tachyon reads a block of data from memory with streaming on and prefetching active, similar to the previous example. The Back-to-Back Stream is depicted in the following transactions, READ8 at ADDR-1, READ8 at ADDR-2, READ8 at ADDR-3, shown in the following three figures.

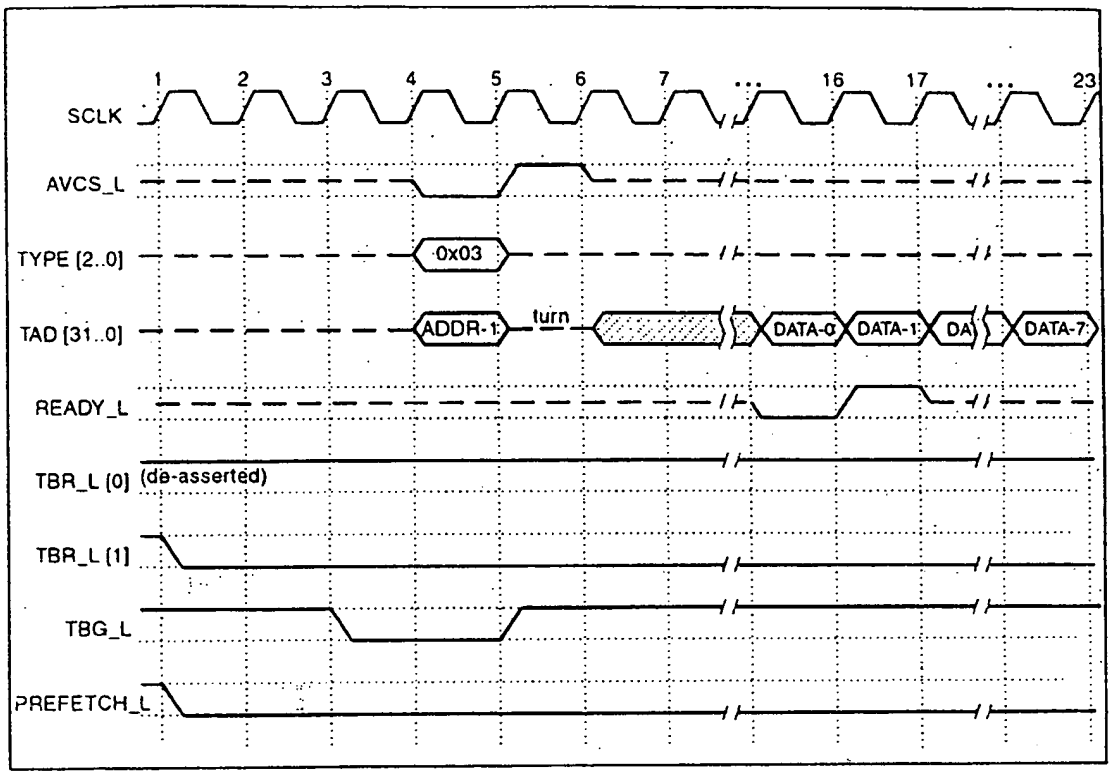

Figure 7.20 Prefetching Across Back-to-Back Read Stream, ADDR-1

Address ADDR-1 incurs a large number of wait states due to latency of the host memory.

PTI 172673

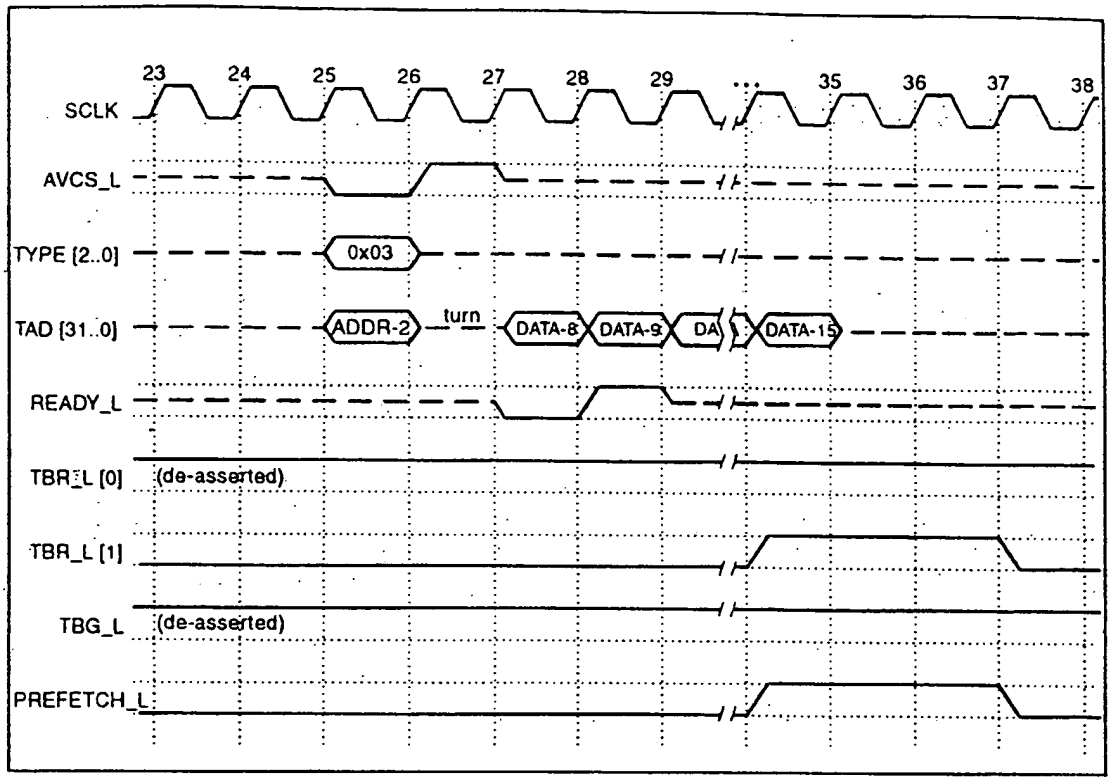

Figure 7.21 Prefetching Across Back-to-Back Read Stream, ADDR-2

Since PREFETCH_L is sampled before the address phase of the first transaction, the data associated with address ADDR-2 can be prefetched, therefore ADDR-2 does not incur many wait states while the read data returns.

In this example, tenancy is not lost and Tachyon knows that it has more sequential Reads to follow, so it gives a prefetch-indication with ADDR-2, even though it is going to lose tenancy.

PREFETCH_L and TBR_L[1] are de-asserted just before the last data cycle of the ADDR-2 transaction. Then, Tachyon immediately requests the bus again. ADDR-3 is the next sequential address, the subject of the prefetch-indication of the second transaction.

PTI 172674

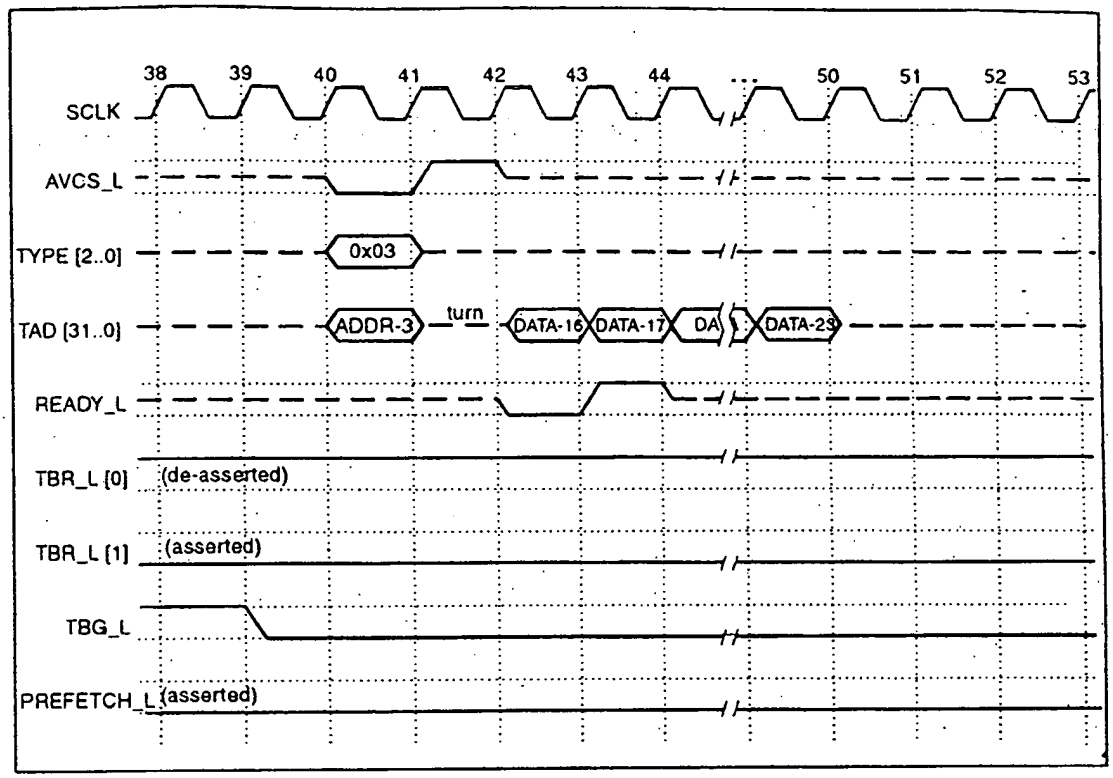

Figure 7.22 Prefetching Across Back-to-Back Read Stream, ADDR-3

The data associated with address ADDR-3 can be returned quickly because of its prefetch-indication in ADDR-2. Since the ADDR-3 transaction is the continuation of sequential accesses, PREFETCH_L is asserted for this transaction.

7.7.9 Retried Read Transaction

In this example, Tachyon retries a READ8 transaction. This is shown in the following two figures.

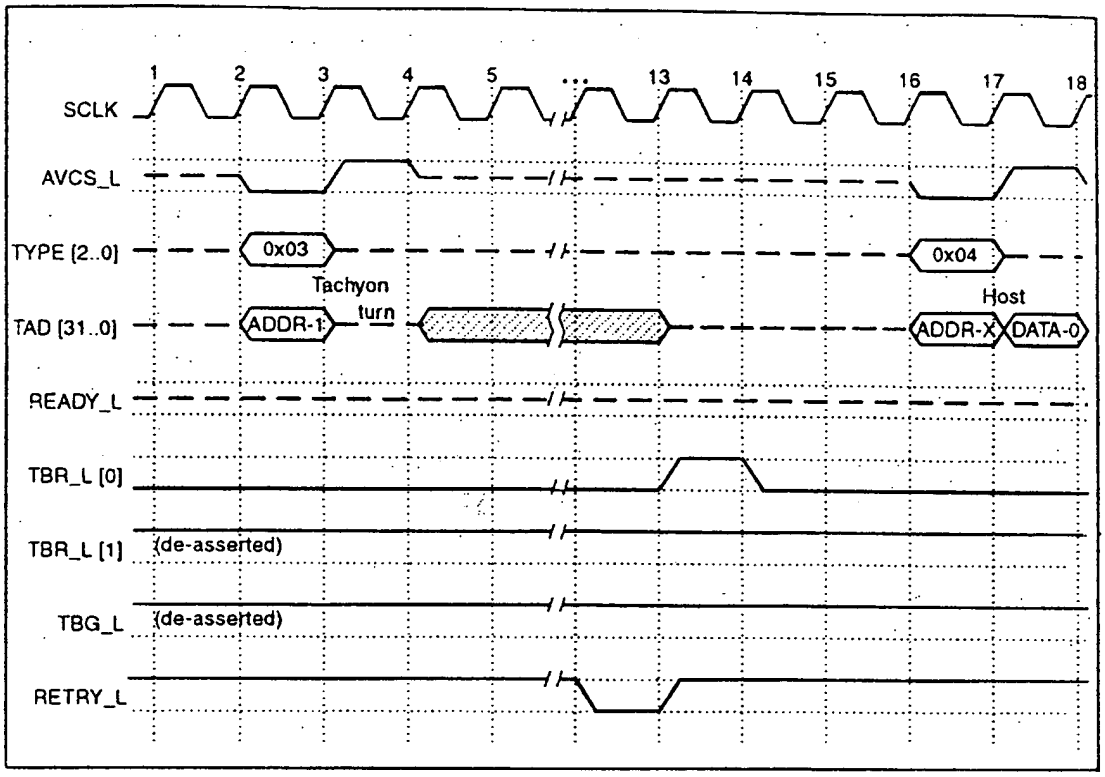

Figure 7.23 Retried Read Transaction, READ8 at ADDR-1 Attempted

When Tachyon requests the bus and attempts to perform a Read transaction at address ADDR-1, the host determines that it does not want to service the Read transaction at this time. As such, the host asserts RETRY_L, which instructs Tachyon to terminate both the transaction and the bus tenancy. Tachyon de-asserts the bus request line, but then immediately re-asserts it to retry the Read transaction.

The second transaction is a slave Write (WRITE1) to Tachyon, which, in this example, may have been deadlocked with the first transaction, which forces the retry.

Note	If Tachyon asserts the PREFETCH_L signal concurrently with the TBR_L[1] signal and the host then decides to retry the read transaction using the RETRY_L signal, then the PREFETCH_L signal will remain asserted even after Tachyon has given up the bus.

PTI 172676

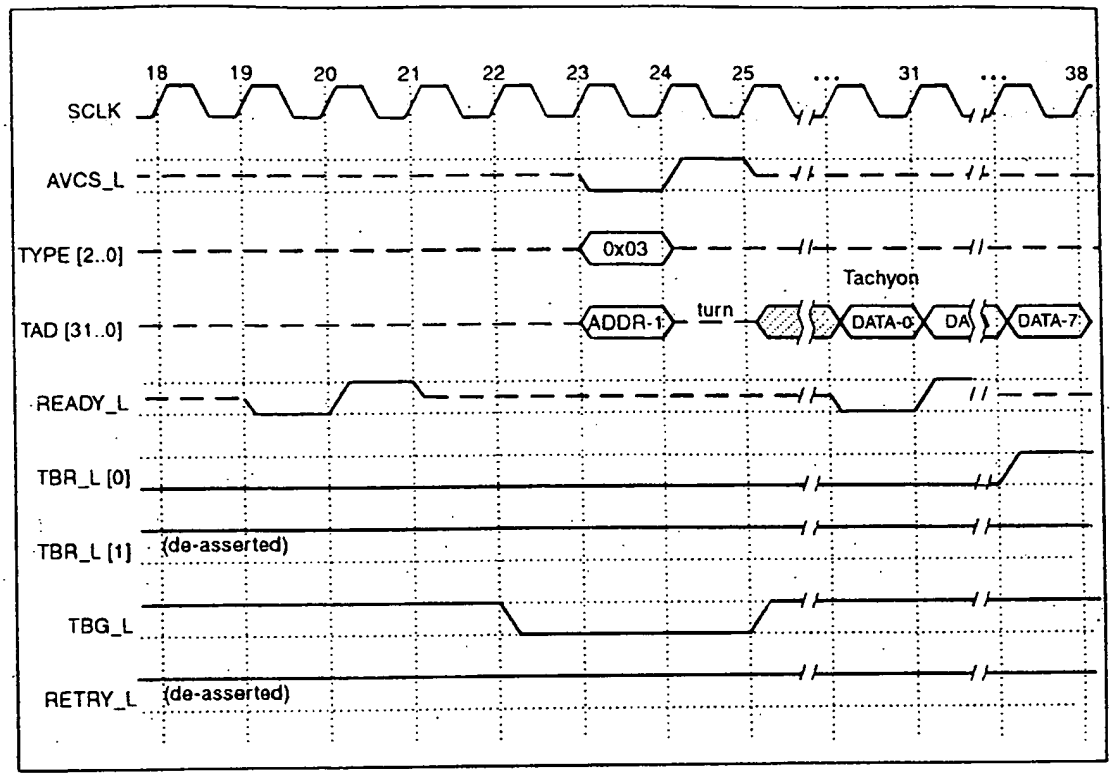

Figure 7.24 Retried Read Transaction, READ8 at ADDR-1 Retried

When the host grants Tachyon the bus again, Tachyon immediately attempts the Retried Read transaction at the same address, ADDR-1. No other Tachyon-mastered transaction may occur until the Retried Read transaction is completed.

PTI 172677

7.7.10 Interrupt Signal

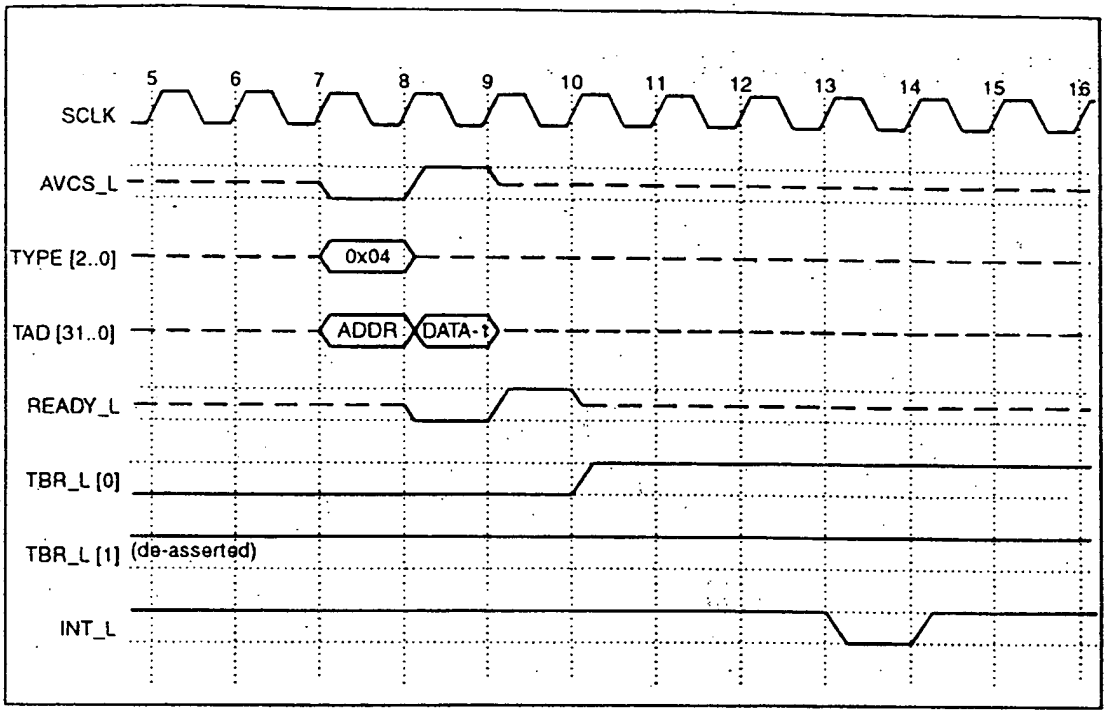

Figure 7.25 Interrupt Signal

This figure shows Tachyon generating an interrupt by asserting the INT_L signal for one clock cycle. The INT_L cannot be asserted if TBR_L [0] or TBR_L [1] is asserted.

PTI 172678

7.8 TSI Timing Requirements

7.8.1 TSI Input Signal Timing Requirements

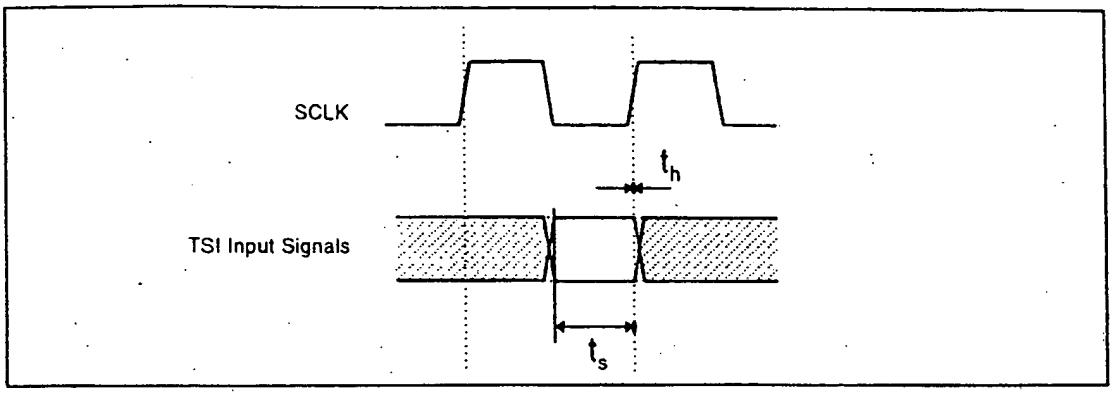

Figure 7.26 TSI Input Signal Timing

Signal	Minimum Setup Time to the Rising Edge of SCLK (t_s)	Minimum Hold Time with repect to the Rising Edge of SCLK (t_h)	Units
READY_L	12.5	0	ns
AVCS_L	12.5	0	ns
TYPE [2..0]	12.5	0	ns
TAD [31..0]	12.5	0	ns
PARITY	12.5	0	ns
RESET_L	14.0	0	ns
RETRY_L	12.5	0	ns
TBG_L	12.5	0	ns

Table 7.9 TSI Input Signal Timing Requirements

The minimum setup time (t_s) applies from the signal being valid to the rising edge of SCLK. The minimum hold time (t_h) applies from the rising edge of SCLK to the signal becoming invalid. There is no hold requirement at the rising edge of SCLK in which a signal is undriven.

PTI 172679

7.8.2 TSI Output Signal Timing Requirements

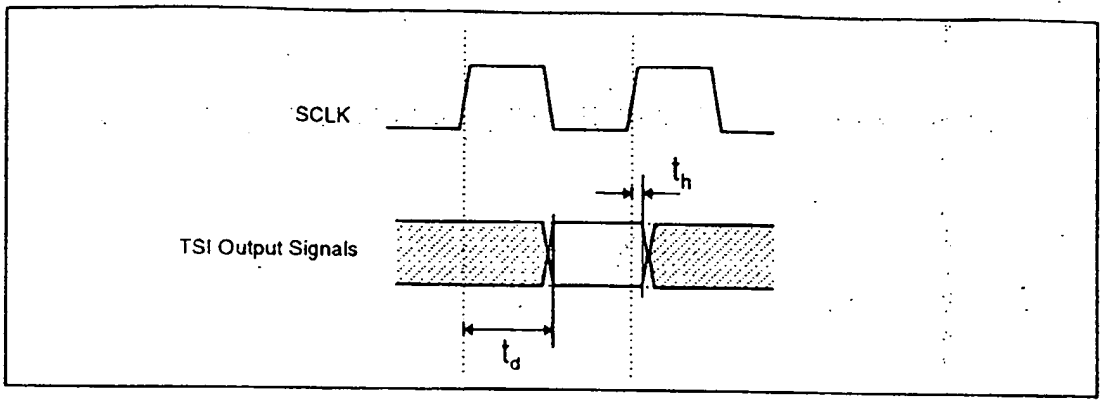

Figure 7.27 TSI Output Signal Timing

Signal	Maximum Delay from the Rising Edge of SCLK (t_d)		Minimum Delay from the Rising Edge of SCLK (t_h)	Units
	25 pF	50pF	0 pF	
READY_L	11.5	12.5	1.0	ns
AVCS_L	11.5	12.5	1.0	ns
TYPE [2..0]	11.5	12.5	1.0	ns
PREFETCH_L	11.5	12.5	1.0	ns
ERROR_L	11.5	12.5	1.0	ns
INT_L	11.5	12.5	1.0	ns
TBR_L [1..0]	11.5	12.5	1.0	ns
TAD [31..0]	13.0	14.0	1.0	ns
PARITY	16.75	17.75	1.0	ns

Table 7.10 TSI Output Signal Timing Requirements

All delays are from the rising SCLK edge to the signal being valid or undriven. Uni-directional signals, i.e., PREFETCH_L, ERROR_L, and TBR_L[ ], are never undriven. AVCS_L and READY_L are always driven low, then driven high, then undriven.

Conditions for Minimum Delays

1. Loading capacitance = 0 pF
2. Best-case chip process

Conditions for Maximum Delays

1. Loading capacitance = 50 pF
2. Worst-case chip process

PTI 172680

7.9 GLM Signal Information

7.9.1 GLM Transmit Signals

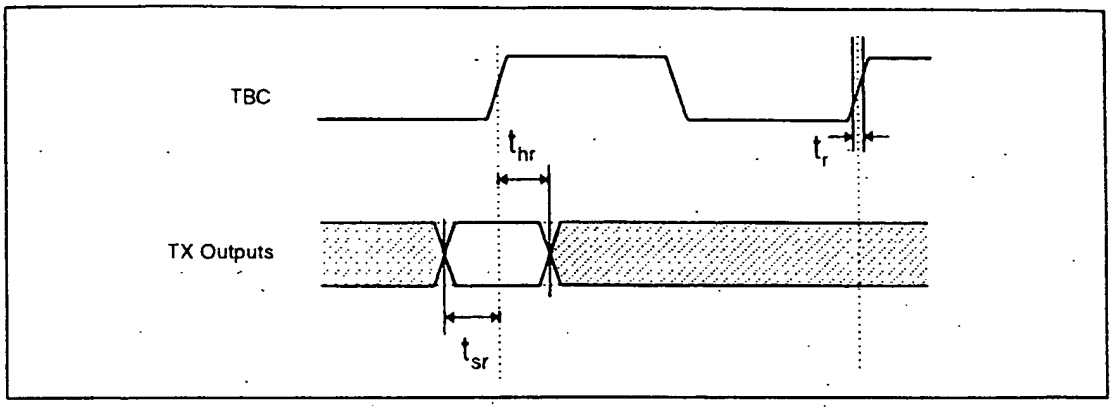

Figure 7.28 GLM Transmit Signal Information

Label	Description	At 26.56 MHz	At 53.13 MHz	Units
C	Capacitance (Refer to Table Note 1, below)	10 / 25 (min / max)	10 / 25 (min / max)	pF
dc_{TBC}	Duty Cycle of TBC	40 / 60	45 / 55	%
t_{sr}	GLM Setup to Rising Edge TBC	6.0	2.0	ns
t_{hr}	GLM Hold from Rising Edge TBC	3.3	3.3	ns
t_r	Rise Time of TBC from 0.8 V to 1.5 V	1.4	1.4	ns
j_{TBC}	Maximum Jitter of TBC	0.005	0.005	% of frequency

Table 7.11 GLM Transmit Signal Information

GLM Transmit Signal Information Table Notes

1. Tachyon meets GLM specifications; therefore a 10 pF output loading is assumed.
2. Table information assumes that transmit signals meet the conditions stated in the "Adapter Board Layout Requirements" section. Refer to "A.1 PCB Layout Suggestions" on page 313.

PTI 172681

7.9.2 GLM Receive Signals

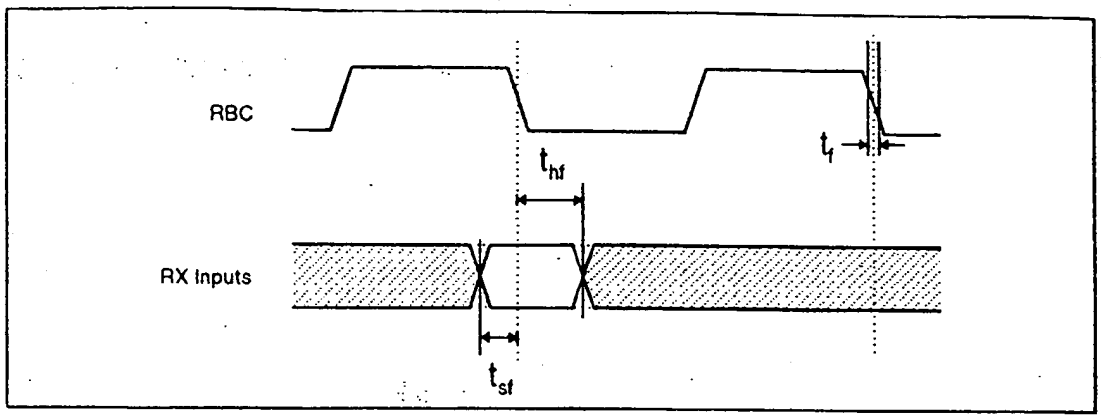

Figure 7.29 GLM Receive Signal Information

Label	Description	At 26.66 MHz	At 53.13 MHz	Units
dc_{RBC}	Duty Cycle of RBC	40 / 60	45 / 55	%
t_{sf}	GLM Setup to Falling Edge RBC	2.5	2.5	ns
t_{hf}	GLM Hold from Falling Edge RBC	6.0	6.0	ns
t_f	Fall Time of RBC from 2.0 V to 1.5 V	1.2	1.2	ns
j_{RBC}	Maximum Jitter of RBC	0.01	0.01	% of frequency

Table 7.12 GLM Receive Signal Information

GLM Receive Signal Information Table Notes

1. Table information assumes that receive signals meet the conditions stated in the "Adapter Board Layout Requirements" section. Refer to "A.1 PCB Layout Suggestions" on page 313.

PTI 172682

8. Electrical Descriptions

8.1 Absolute Maximum Ratings

Label	Parameter	Minimum	Maximum	Units
Vdd	Absolute Supply Voltage (Refer to the Warning below)	-0.3	+3.9	Volts
T_{op}	Ambient Operating Temperature (Refer to "A.3 Limited Airflow Applications" on page 321.)	0	+50 (0 m/s airflow) +70 (1.5 m/s airflow)	° C
T_{stg}	Storage Temperature	-40	+100	° C

Table 8.1 Absolute Maximum Electrical Ratings

WARNING The Absolute Minimum/Maximum Supply Voltage (Vdd) specification should not be used as a limit for normal device operation. Sustained operation exceeding the limits of the Recommended Supply Voltage (Refer to "Table 8.2 Recommended Operating Conditions" on page 249.) could result in permanent device damage or impaired device reliability.

8.2 Recommended Operating Conditions

Label	Parameter	Minimum	Typical	Maximum	Units
Vdd	Recommended Supply Voltage	3.0	3.3	3.6	Volts
f_{clk}	TSI Clock Frequency, SCLK	24	-	40	MHz
dc_{clk}	Duty Cycle, SCLK (Refer to Note 1 below)	40	-	60	%
$t_{r/f}$	Rise/Fall Time, SCLK (Refer to Note 2 below)	-	2.5	3.5	ns
j_{clk}	Maximum Jitter of SCLK	-	-	100	ps
TBC_r	TBC Rise Time	0.5	-	2.8	ns
RBC_f	RBC Fall Time	0.7	-	2.4	ns

Table 8.2 Recommended Operating Conditions

Recommended Operating Conditions Notes

1. For Duty Cycle for TBC and RBC:
   a. at 26 MHz, a maximum 40/60 Duty Cycle is allowed **PTI 172683**
   b. at 53 MHz, a maximum 45/55 Duty Cycle is allowed
2. The Rise/Fall Time of SCLK should be kept under 3.5 ns. A Rise/Fall Time at 2.5 ns or lower is preferred to minimize insertion delays and phase error.

8.3 Electrical Parameters

Label	Parameter	Signal (S)	Condition	Min	Typ	Max	Units
Vdd	Supply Voltage	Vdd	-	3.0	3.3	3.6	V
I_{dd}	Supply Current	Vdd	SCLK = 33 MHz RBC = 53 MHz TBC = 53 MHz Vdd = 3.6V T_{amb} = 50° C	-	-	1.0	A
V_{il}	Input Low Voltage	all inputs	-	Vss -0.5	-	0.8	V
V_{ih}	Input High Voltage	all inputs	-	2.0	-	5.5	V
I_{in}	Input Current	PLL_RSTN (pin 44)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	PLL_IDD_TEST (46)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	PLL_TEST_DATA (47)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	IDD_TEST (pin 55)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	TCK (pin 122)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	TMS (pin 124)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	RSTN (pin 125)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	TDI (pin 126)	V_{in} = Vss or Vdd	-35	-115	-214	µA
I_{in}	Input Current	PLL_TEST_MODE (48)	V_{in} = Vss or Vdd	35	115	222	µA
I_{in}	Input Current	SCAN_EN (pin 136)	V_{in} = Vss or Vdd	35	115	222	µA
I_{in}	Input Current	TEST_MODE (pin 140)	V_{in} = Vss or Vdd	35	115	222	µA
I_{in}	Input Current	all other inputs	V_{in} = Vss or Vdd	-10	±1	+10	µA
V_{oh}	Output High Voltage	all outputs	I_{ol} = -6 mA	2.4	-	Vdd	V
V_{ol}	Output Low Voltage	all outputs	I_{ol} = 6 mA	-	-	0.4	V
I_{oz}	3-State Output Leakage Current	-	V_{oh} = Vss or 5.5V	-10	±1	+10	µA
I_{os}	Output Short Circuit Current	-	V_{out} = Vdd	-	-	140	mA
I_{os}	Output Short Circuit Current	-	V_{out} = Vss	-	-	-40	mA
C_{in}	Pin Input Capacitance	-	-	3.0	-	5.0	pF
C_{out}	Pin Output Capacitance	-	-	3.0	-	5.0	pF

Table 8.3 Electrical Parameters

PTI 172684

8.4 Pull-Up Values

The following table lists the TSI signals that require pull-up resistors. All signals listed here must be pulled up to Vdd. A recommended resistor value is listed. Tolerances for these resistors are all 10%, unless stated otherwise.

TSI Signal Name	Pull-Up Value (Ω)
TAD[31..0]	10K x 32
PARITY	10K
AVCS_L	10K
TYPE[2..0]	10K x 3
READY_L	10K
RETRY_L	10K
PREFETCH_L	10K
ERROR_L	10K
INT_L	10K
TBR_L[1..0]	10K x 2

Table 8.4 Pull-up Values

8.5 Pull-Down Value

The following test pin must be pulled down to ground via a pulldown resistor.

Test Signal Name	Pin #	Pull-DownValue (Ω)
TEST_MODE	140	4.7K

Table 8.5 Test Mode Pin

PTI 172685

8.6 External PLL Components

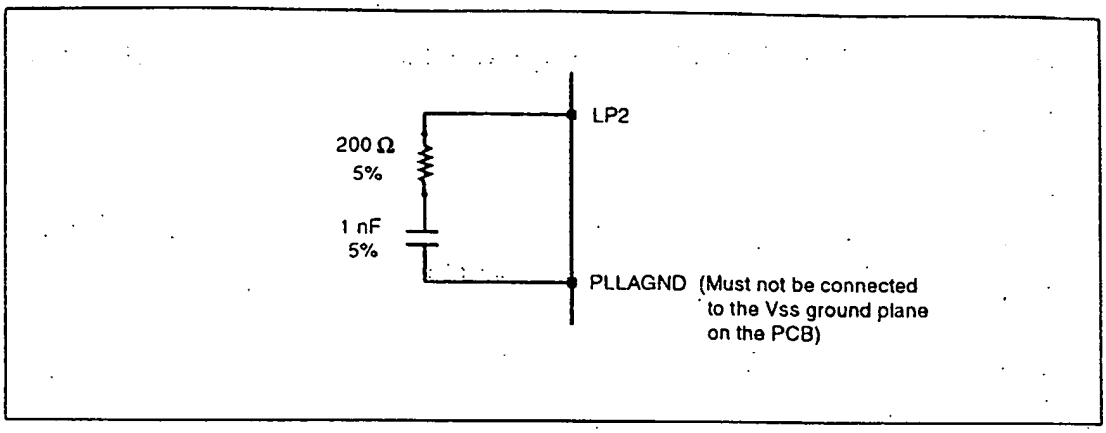

Figure 8.1 External PLL Components

PTI 172686

9. Mechanical Descriptions

9.1 General Information

Package Type	208 Metal Quad Flat Pack (MQUAD)
Power Dissipation	3.5 Watts
Airflow Requirements	1.5 meters/second

Table 9.1 Mechanical Information

9.2 Thermal Specifications

Conditions for Thermal Specifications:

1. No heat sink
2. Power dissipation = 3.5 Watts
3. Tj (max) = 110 ° C

Parameter	Conditions				Units
Airflow	0.0	1.0	2.0	3.0	meters/second
Package Thermal Resistance (θja)	15	12	11	10	° C/Watt
Package Thermal Resistance (θjc)	2.6	2.6	2.6	2.6	° C/Watt
Maximum Ambient Temperature	58	68	72	75	° C

Table 9.2 Thermal Specifications

PTI 172687

9.3 Dimensions

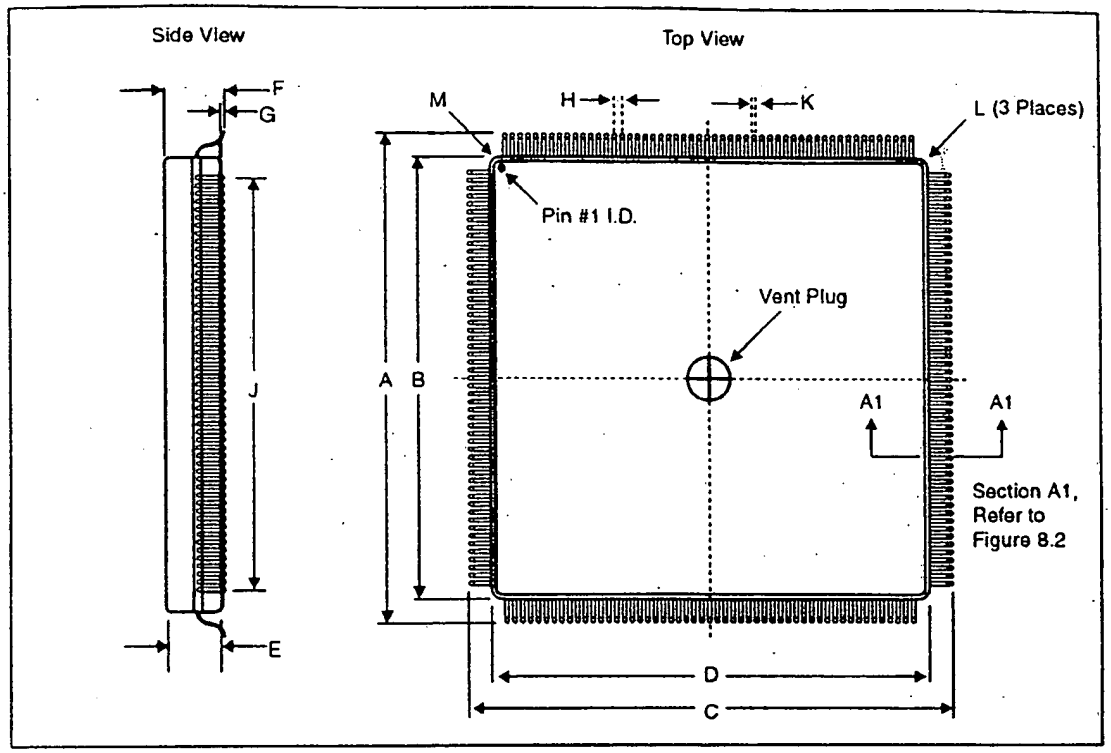

Figure 9.1 Side View and Top View

Dimension	Millimeters*		Inches	
	Minimum	Maximum	Minimum	Maximum
A	30.40	30.80	1.197	1.213
B	27.56	27.72	1.085	1.091
C	30.40	30.80	1.197	1.213
D	27.56	27.72	1.085	1.091
E	3.17	3.43	0.125	0.135
F	3.50	3.86	0.138	0.152
G	0.25	0.51	0.010	0.020
H	0.50 BSC (Basic)		0.019685 BSC	
J	25.50 BSC		1.0039 BSC	
K	0.23 TYP (Typical)		0.009 TYP	
L	0.20 x 45° TYP		0.008 x 45° TYP	
M	0.89 x 45° TYP		0.035 x 45° TYP	
* Note that the controlling dimensions for this part are metric.				

Table 9.3 Side View and Top View Dimensions

PTI 172688

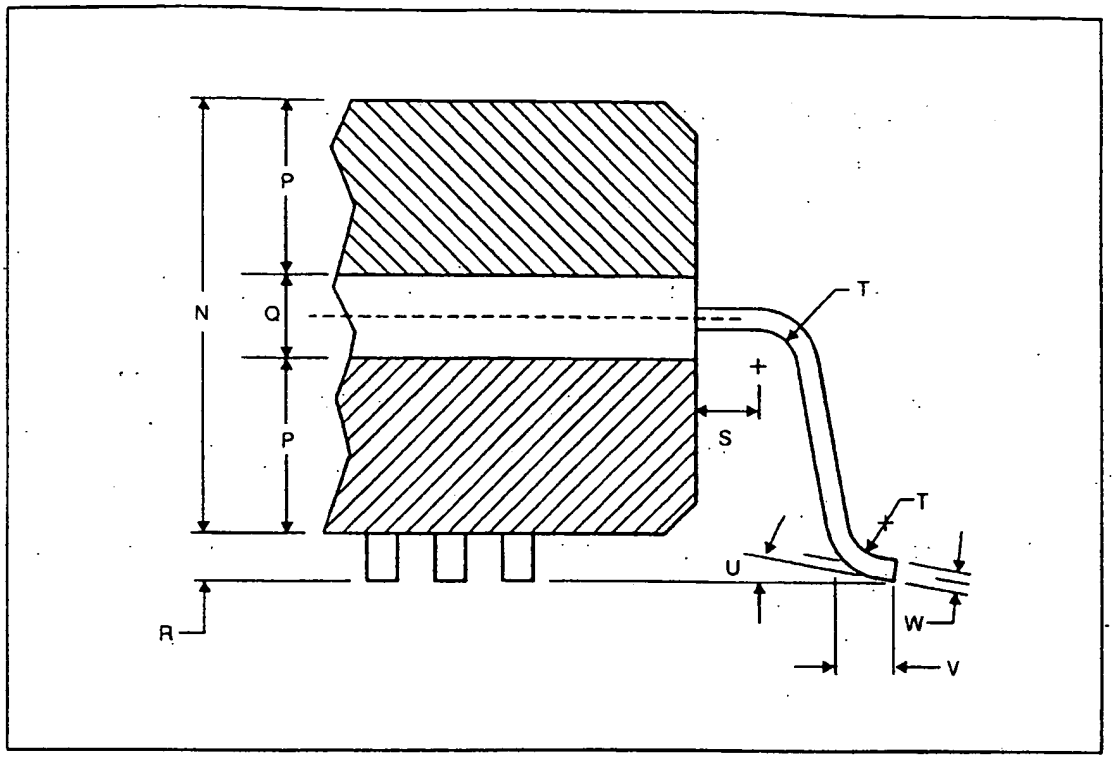

Figure 9.2 Section A1 View

Dimension	Millimeters		Inches	
	Minimum	Maximum	Minimum	Maximum
N	3.17	3.43	0.125	0.135
P	1.32		0.052	
Q	0.66		0.026	
R	0.25	0.51	0.010	0.020
S	0.58	-	0.023	-
T	0.25 R		0.010 R	
U	0°	7°	0°	7°
V	0.40	0.60	0.016	0.024
W	0.15		0.006	

Table 9.4 Section A1 Dimensions

PTI 172689

10. Error Information

10.1 Fatal Errors

Tachyon detects five different types of errors:

1. TSI address parity error
2. TSI DMA Read data parity error
3. TSI Slave Write data parity error
4. TSI Invalid slave transaction type
5. Tachyon internal data path parity error (also indicated by the Transmit Parity Error bit in the Frame Manager Status register)

For the above five errors, the following actions occur:

1. Tachyon asserts the ERROR_L signal on TSI. The ERROR_L signal remains asserted until the host performs a hard reset of Tachyon (the host asserts RESET_L).
2. Tachyon no longer arbitrates for TSI. The only TSI transactions which may occur after an error occurs are slave reads and writes to Tachyon's internal registers. All these errors are fatal to Tachyon operation.
3. Tachyon updates the Fatal Error Status in its Tachyon Status register to indicate to the host the error that caused Tachyon to assert the ERROR_L signal. If Tachyon asserts ERROR_L and another of the five error conditions occurs, Tachyon ignores the subsequent error condition. Once the host detects that Tachyon asserted ERROR_L, the host should perform any actions necessary on the current outstanding I/O. The host should perform slave reads to Tachyon to determine the cause of the error and then log the condition. Finally the host should perform a hardware reset of Tachyon (the host asserts RESET_L).

10.2 OSM Freeze

10.2.1 OSM Freeze States

The OSM freezes (stops processing ODBs) for the following reasons:

1. The link goes down anytime after it comes up initially.
2. Tachyon sets the Bad AL_PA (ba) bit in the Frame Manager Status register, indicating that an outbound frame was sent to a bad AL_PA.
3. Tachyon sets any of the following bits in the outbound_completion message: Class 1 error (C1), Link Down (LD), Abort Requested (A), ACK Timeout (AT), Frame Timeout (OT), Rejected (R), Retries Exceeded (X), Host Programming Error (HE).

Tachyon may send a completion message in each of the cases listed above. Refer to "10.2.2 OSM Recovery Using OCQ Reset" on page 258.

10.2.2 OSM Recovery Using OCQ Reset

After the OSM is frozen (but before the host unfreezes the it), the host may remove and/or re-order entries on the OCQ. There may be prefetched ODB entries inside Tachyon waiting to be processed. To remove these prefetched ODBs, the OCQ must be reset to its initial empty state by setting the OCQ Reset (or) bit in the Tachyon Control register.

To unfreeze the OSM, the host sets the Error Release bit in the Tachyon Control register. The OSM will now continue normal operation of the OCQ.

10.2.3 OSM Freeze Notes

1. Tachyon can retrieve an OCQ entry and buffer it for the OSM even if the OSM is in the frozen state.
2. If an outbound Class 1 connection is opened and Tachyon receives an unexpected frame with EOFdt, the OSM closes the connection, sets the Class 1 Error bit (in the outbound_completion message), and freezes. If the host sends Class 2 or Class 3 Multiframe Sequences (MFSs) while this Class 1 connection is opened, then these MFSs are aborted if the Class 1 connection closes due to an unexpected frame with EOFdt.

10.3 Error Recovery Procedures

10.3.1 Blocked Outbound Frame FIFO error recovery

If the host determines that it needs to send a high priority frame for error recovery purposes while the OSM is frozen, it should first determine the status of the Outbound Frame FIFO. The Outbound Frame FIFO bit of the Tachyon Status register indicates whether or not the Outbound Frame FIFO is blocked. The Outbound Frame FIFO is blocked when a frame is in the Outbound Frame FIFO and the Frame Manager is unable to pull out the frame. Since frames going out the High Priority Channel must also go through the Outbound Frame FIFO, sending an error recovery frame in this case is ineffective.

To clear the blocked Outbound Frame FIFO, use the following error recovery process:

1. For N_Ports, the host instructs the Frame Manager, via the Frame Manager Control register, to send a Link Reset. After the Frame Manager sends a Link Reset, the ISM and OSM reset both Inbound and Outbound FIFOs, thus clearing out the blocked condition. For FC-AL, a Link Reset cannot be send; the host instructs Frame Manager to send an Initialize command.
2. When the host sees that the FIFOs are clear, it may send any needed error recovery frames via the High Priority Channel.

10.3.2 ABTS Protocol

Generic description of the processing required to perform the Abort Sequence (ABTS) protocol.

Sequence Initiator	Sequence Recipient
Send ABTS frame via the High Priority channel.	Receive ABTS via the SFS channel.
Wait for ACK via SFS completion.	Send ACK for ABTS.
Wait for Basic Accept frame.	Send BA_ACC to initiator via OCQ.
Perform host sequence cleanup as required.	Perform host sequence cleanup as required.

Table 10.1 ABTS Protocol

10.3.3 Class 1 Connection Recovery Procedure

The host uses the Class 1 connection recovery procedure to check if the state of the Class 1 connection is correct for possible subsequent sequences that are processed by Tachyon. The state of the Class 1 connection may be determined by issuing a "Status Request" command to the Tachyon Control register.

If an inbound Class 1 connection is open after an inbound sequence error or a inbound_c1_timeout completion is received indicating an idle inbound connection, the host may either send a request to the remote node via the High Priority channel to close the Class 1 connection, or issue the Link Reset command to the Frame Manager Control register to close the Class 1 connection.

PTI 172692

The following pseudo-code describes the outbound Class 1 connection recovery procedure:

```
Check the state of Class 1 connection.
If (Class 1 connection is up)
{
        if (the host has no more Class 1 sequences in the OCQ to send)
        {
                reset OCQ (set OCQ Reset bit in the Tachyon Control register)
                post a RMC sequence to terminate connection or perform Link
                        Reset
                repost the sequences left on queue by error after the RMC or
                        Link Reset
        }
}
else (if Class 1 connection is down)
{
        if (host had more than one Class 1 sequence in queue)
        {
                reset OCQ (set OCQ Reset bit in the Tachyon Control register)
                modify next sequence to establish the class 1 connection
                repost the sequences left on queue by error
        }
}
Unfreeze OSM (set Error Release bit in Tachyon Control register)
```

It may be difficult to determine the exact state of the Class 1 connection if the connection is left open between sequences and an error occurs. The state of the connection may be determined via the "Status Request" command in the Tachyon Control register or the host may force the connection closed by performing a link reset via the Frame Manager Control register.

10.3.4 Sequence Recovery Procedure

The host uses the sequence recovery procedure to prevent sequences from being transmitted out of order when the upper layer protocol cannot support receiving them this way. When the host is given an indication that an error occurred in sending a sequence, the host must check if there are other sequences of this type. If there are sequences that shouldn't be sent yet, the host must remove them from the OCQ.

10.3.5 Link Down Recovery Procedure

When a link down is detected (by reading the Frame Manager Status register), the host must determine if any steps are required to restore the link. Under most conditions, the link returns to the link up state on its own. However, if the Link Failure bit is set or the Port State Machine is in either the LF1 or LF2 states, the host must issue the Clear LF command to the Frame Manager Control register.

10.3.6 Link Down and Login Recovery

Tachyon informs the host when the link is up and when it is down. Unfortunately, the FC_PH standard indicates that some link down conditions (such as receiving an OLS primative sequence) cause an explicit logout. The Frame Manager does provide some indications that the host must re-login to the fabric and/or remote ports via the Frame Manager Status register's LIPf, Fabric Login Required, Link Failure and Port State Machine bits. Unfortunately, not all link down conditions that may have caused an explicit logout (such as receiving an OLS) may be indicated. Therefore, the host should always verify that it is still logged in with the fabric and/or remote ports before proceeding with normal operation after a link down condition.

10.3.7 Link Failure Management

A link failure condition occurs when a Laser Fault, Out of Syncronization or Loss of Signal condition occurs within Frame Manager. Frame Manager sends a frame_mgr_interrupt completion message to the host to indicate that the host should read the Frame Manager Status register. After the host reads the Frame Manager Status register to determine the cause of the frame_mgr_interrupt completion message, the bit(s) that caused the interrupt are cleared. If the condition persists, the Frame Manager may again set the bit(s) in the Frame Manager Status register, causing another frame_mgr_interrupt completion message. Should the failure condition persist for a long period of time, many completion messages could be sent to the host. To prevent the constant interruptions, the host should write the Offline Command to the Frame Manager Control register. This stops the Frame Manager from sending frame_mgr_interrupt completion messages to the host. The host should then periodically read the Frame Manager Status register to see if the failure condition has been corrected. Once the condition is corrected, the host should write the Initialize Command to the Frame Manager Control register.

10.3.8 OSM Recovery

Refer to "10.2.2 OSM Recovery Using OCQ Reset" on page 258.

PTI 172694

10.4 Error Condition Tables

The following Error Condition tables show error conditions that can occur in the various error configurations that Tachyon supports. Detailed descriptions of error conditions are located in the section that follows the tables.

10.4.1 How to use the Error Condition Tables

1. Locate the error condition label in the top row of one of the error condition tables.
2. Locate the type of configuration in the leftmost column.
3. Scroll down the error condition column to the type of configuration selected. The intersection of the error condition and the type of configuration refers to an error number (Ex.x). Refer to the error number and its associated page number for a detailed description of the error condition.

Example

If Tachyon is the Initiator operating in a Class 2 Point-to-Point configuration and a time-out occurs due to a Lost ACK:

1. Locate the error condition label ("Lost ACK") in the top row of one of the error condition tables. "Lost ACK" is found in the first error condition table, "Table 10.2 Tachyon Fabric/Pt.-Pt. Configuration Error Conditions 1 - 6", page 263.
2. Locate the type of configuration in the leftmost column. In this case, Tachyon is a "Non-SCSI Initiator in a Class 2 Point-to-Point Configuration."
3. Scrolling down the "Lost ACK" column to the "Non-SCSI Initiator Class 2 Point-to-Point Configuration," check the error condition number. The error condition number in this example is E1.1. Refer to the detailed description in E1.1, page 276.

PTI 172695

Configurations	Error Conditions											
	1. Lost ACK		2. No EE_Credit for Tx		3. No BB_Credit for Tx		4. ACK with Abort bits = 01, 10, 11		5. ACK with history bit = 0		6. ACK with history bit = 1	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	E1.1	276	E1.1	276	-	-	E1.3	277	E1.4	277	E1.5	277
Class 2	E1.1	276	E1.1	276	E1.2	276	E1.3	277	E1.4	277	E1.5	277
Class 3	-	-	-	-	E1.2	276	-	-	-	-	-	-
Responder												
Class 1	-	-	-	-	-	-	-	-	-	-	-	-
Class 2	-	-	-	-	E1.2	276	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
SCSI Initiator Cmd Phase												
Class 1	Class 1 not supported for Command Sequence											
Class 2	E1.1	276	E1.1	276	E1.2	276	E1.3	277	E1.4	277	E1.5	277
Class 3	-	-	-	-	E1.2	276	-	-	-	-	-	-
Data Phase												
Class 1	E1.1	276	E1.1	276	-		E1.3	277	E1.4	277	E1.5	277
Class 2	E1.1	276	E1.1	276	E1.2	276	E1.3	277	E1.4	277	E1.5	277
Class 3	-	-	-	-	E1.2	276	-	-	-	-	-	-
Responder Data Phase												
Class 1	-	-	-	-	-	-	-	-	-	-	-	-
Class 2	-	-	-	-	E1.2	276	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Status Phase												
Class 1	Class 1 not supported for Status Sequence											
Class 2	-	-	-	-	E1.2	276	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-

Table 10.2 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 1 - 6

PTI 172696

Configurations	Error Conditions											
	7. ACK_1 on out frame; expecting an ACK_0		8. RJT		9. BSY		10. BSY retries exceeded		11. Unexpected ACK		12. Unexpected EOFdt	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	E1.6	278	E1.7	278	E1.8	279	E1.9	279	E2.1	282	E2.2	282
Class 2	E1.6	278	E1.7	278	E1.8	279	E1.9	279	E2.1	282	E2.1	282
Class 3	-	-	-	-	-	-	-	-	E2.1	282	E2.1	282
Responder												
Class 1	-	-	E1.12	281	E1.12	281	-	-	E2.1	282	E2.3	283
Class 2	-	-	E1.12	281	E1.12	281	-	-	E2.1	282	E2.1	282
Class 3	-	-	-	-	-	-	-	-	E2.1	282	E2.1	282
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	E1.6	278	E1.10	280	E1.8	279	E1.9		E2.1	282	E2.1	282
Class 3	-	-	-	-	-	-	-	-	E2.1	282	E2.1	282
Data Phase												
Class 1	E1.6	278	E1.10	280	E1.8	279	E1.11	280	E2.1	282	E2.2	282
Class 2	E1.6	278	E1.10	280	E1.8	279	E1.11	280	E2.1	282	E2.1	282
Class 3	-	-	-	-	-	-	-	-	E2.1	282	E2.1	282
Responder Data Phase												
Class 1	-	-	E1.12	281	E1.12	281	-	-	E2.1	282	E2.3	283
Class 2	-	-	E1.12	281	E1.12	281	-	-	E2.1	282	E2.1	282
Class 3	-	-	-	-	-	-	-	-	E2.1	282	E2.1	282
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	-	-	E1.12	281	E1.12	281	-	-	E2.1	282	E2.1	282
Class 3	-	-	-	-	-	-	-	-	E2.1	282	E2.1	282

Table 10.3 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 7 - 12

PTI 172697

Configurations	Error Conditions											
	13. Unexpected EOFt		14. EOFdti		15. EOFa		16. EOFni		17. Link Down		18. Link Fail (LOS, OLS, LF)	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	E2.4	283	E2.6	283	E2.9	285	E2.10	285	E2.11	285	E2.14	287
Class 2	E2.4	283	E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Class 3	-		E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Responder												
Class 1	E2.5	283	E2.7	284	E2.9	285	E2.10	285	E2.11	285	E2.14	287
Class 2	E2.5	283	E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Class 3	E2.5	283	E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	-	-	E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Class 3	-	-	E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Data Phase												
Class 1	E2.4	283	E2.6	283	E2.9	285	E2.10	285	E2.11	285	E2.14	287
Class 2	E2.4	283	E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Class 3	-		E2.8	284	E2.9	285	E2.10	285	E2.12	286	E2.15	288
Responder Data Phase												
Class 1	E2.5	283	E2.7	284	E2.9	285	E2.10	285	E2.13	287	E2.16	289
Class 2	E2.5	283	E2.8	284	E2.9	285	E2.10	285	E2.13	287	E2.16	289
Class 3	E2.5	283	E2.8	284	E2.9	285	E2.10	285	E2.13	287	E2.16	289
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	E2.5	283	E2.8	284	E2.9	285	E2.10	285	E2.13	287	E2.16	289
Class 3	E2.5	283	E2.8	284	E2.9	285	E2.10	285	E2.13	287	E2.16	289

Table 10.4 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 13 - 18

PTI 172698

Configurations	Error Conditions											
	19. ABTS		20. ABTX		21. RMC Received		22. LCR		23. SOFc1 with connection		24. SOFi,n1 with connection	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	E2.17	289	E2.20	290	E2.23	292	E2.26	293	E3.1	294	E3.2	294
Class 2	E2.17	289	E2.20	290	-	-	E2.26	293	E3.1	294	-	-
Class 3	E2.17	289	E2.20	290	-	-	E2.26	293	E3.1	294	-	-
Responder												
Class 1	E2.18	290	E2.21	291	E2.24	292	E2.26	293	E3.1	294	E3.2	294
Class 2	E2.18	290	E2.21	291	-	-	E2.26	293	E3.1	294	-	-
Class 3	E2.18	290	E2.21	291	-	-	E2.26	293	E3.1	294	-	-
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	E2.17	289	E2.20	290	-	-	E2.26	293	E3.1	294	-	-
Class 3	-	-	-	-	-	-	E2.26	293	E3.1	294	-	-
Data Phase												
Class 1	E2.17	289	E2.20	290	E2.23	292	E2.26	293	E3.1	294	E3.2	294
Class 2	E2.17	289	E2.20	290	-	-	E2.26	293	E3.1	294	-	-
Class 3	E2.17	289	E2.20	290	-	-	E2.26	293	E3.1	294	-	-
Responder Data Phase												
Class 1	E2.19	290	E2.22	291	E2.25	293	E2.26	293	E3.1	294	E3.2	294
Class 2	E2.19	290	E2.22	291	-	-	E2.26	293	E3.1	294	-	-
Class 3	E2.19	290	E2.22	291	-	-	E2.26	293	E3.1	294	-	-
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	-	-	-	-	-	-	E2.26	293	E3.1	294	-	-
Class 3	-	-	-	-	-	-	E2.26	293	E3.1	294	-	-

Table 10.5 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 19 - 24

PTI 172699

Configurations	Error Conditions											
	25. SOFi,n1 without connection		26. Unexpected inbound frame with E_C		27. Bidirectional SOFc1		28. NOP		29. Other BLS		30. ELS	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	E3.3	295	E3.4	295	E3.6	295	-	-	-	-	-	-
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Responder												
Class 1	E3.3	295	E3.5	295	E3.7	296	E3.8	296	E3.9	296	E3.10	296
Class 2	-	-	-	-	-	-	E3.8	296	E3.9	296	E3.10	296
Class 3	-	-	-	-	-	-	E3.8	296	E3.9	296	E3.10	296
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Data Phase												
Class 1	E3.3	295	E3.4	295	E3.6	295	-	-	-	-	-	-
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Responder Data Phase												
Class 1	E3.3	295	E3.5	295	E3.7	296	E3.8	296	E3.9	296	E3.10	296
Class 2	-	-	-	-	-	-	E3.8	296	E3.9	296	E3.10	296
Class 3	-	-	-	-	-	-	E3.8	296	E3.9	296	E3.10	296
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	-	-	-	-	-	-	E3.8	296	E3.9	296	E3.10	296
Class 3	-	-	-	-	-	-	E3.8	296	E3.9	296	E3.10	296

Table 10.6 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 25 - 30

PTI 172700

Configurations	Error Conditions											
	31. Rx Timeout		32. C1 Connection Activity Timeout		33. BB_Credit violation		34. Frame too long		35. Expiration Header		36. Second MFS	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	-	-	-	-	-	-	-	-	-	-	-	-
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Responder												
Class 1	E4.1	297	E4.2	297	-	-	E4.4	298	E4.5	298	E4.6	298
Class 2	E4.1	297	-		E4.3	297	E4.4	298	E4.5	298	E4.6	298
Class 3	E4.1	297	-		E4.3	297	E4.4	298	E4.5	298	E4.6	298
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Data Phase												
Class 1	-	-	-	-	-	-	-	-	-	-	-	-
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Responder Data Phase												
Class 1	-	-	E4.2	297	-	-	E4.4	298	E4.5	298	-	-
Class 2	-	-	-	-	E4.3	297	E4.4	298	E4.5	298	-	-
Class 3	-	-	-	-	E4.3	297	E4.4	298	E4.5	298	-	-
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	-	-	-	-	E4.3	297	E4.4	298	E4.5	298	-	-
Class 3	-	-	-	-	E4.3	297	E4.4	298	E4.5	298	-	-

Table 10.7 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 31 - 36

PTI 172701

Configurations	Error Conditions											
	37. C2, C3 MFS on C1 connection		38. OOO Reassembly cleared		39. OOO Reassembly set		40. Unknown Frame		41. Frame from unlogged in node		42. Frame for invalid exchange	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	-	-	-	-	-	-	-	-	-	-	-	-
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Responder												
Class 1	-	-	E4.9	299	E4.9	299	E4.13	301	E4.15	302	E4.16	302
Class 2	E4.8	299	E4.10	299	E4.10	299	E4.13	301	E4.15	302	E4.16	302
Class 3	E4.8	299	E4.10	299	E4.10	299	E4.13	301	E4.15	302	E4.16	302
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Data Phase												
Class 1	-	-	-	-	-	-	-	-	-	-	-	-
Class 2	-	-	-	-	-	-	-	-	-	-	-	-
Class 3	-	-	-	-	-	-	-	-	-	-	-	-
Responder Data Phase												
Class 1	-	-	E4.11	300	E4.12	300	E4.14	302	E4.15	302	E4.17	302
Class 2	-	-	E4.11	300	E4.12	300	E4.14	302	E4.15	302	E4.17	302
Class 3	-	-	E4.11	300	E4.12	300	E4.14	302	E4.15	302	E4.17	302
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	-	-	-	-	-	-	E4.14	302	E4.15	302	E4.17	302
Class 3	-	-	-	-	-	-	E4.14	302	E4.15	302	E4.17	302

Table 10.8 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 37 - 42

PTI 172702

Configurations	Error Conditions											
	43. SCSI, ULP Timeout		44. SCSI, xfr_rdy queued at exch term		45. SCSI OX_ID direction bit wrong		46. OB timestamp expiration (frame in OB FIFO too long)		47. Host Programming Error		48. Elastic Store Error	
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	-	-	-	-	-	-	E5.4	304	E5.5	304	E5.6	304
Class 2	-	-	-	-	-	-	E5.4	304	E5.5	304	E5.6	304
Class 3	-	-	-	-	-	-	E5.4	304	E5.5	304	E5.6	304
Responder												
Class 1	-	-	-	-	-	-	-	-	-	-	E5.6	304
Class 2	-	-	-	-	-	-	-	-	-	-	E5.6	304
Class 3	-	-	-	-	-	-	-	-	-	-	E5.6	304
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	E5.1	303	-	-	-	-	E5.4	304	E5.5	304	E5.6	304
Class 3	E5.1	303	-	-	-	-	E5.4	304	E5.5	304	E5.6	304
Data Phase												
Class 1	E5.1	303	E5.2	303	-	-	E5.4	304	E5.5	304	E5.6	304
Class 2	E5.1	303	E5.2	303	-	-	E5.4	304	E5.5	304	E5.6	304
Class 3	E5.1	303	E5.2	303	-	-	E5.4	304	E5.5	304	E5.6	304
Responder Data Phase												
Class 1	E5.1	303	-	-	E5.3	304	-	-	-	-	E5.6	304
Class 2	E5.1	303	-	-	E5.3	304	-	-	-	-	E5.6	304
Class 3	E5.1	303	-	-	E5.3	304	-	-	-	-	E5.6	304
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	E5.1	303	-	-	E5.3	304	-	-	-	-	E5.6	304
Class 3	E5.1	303	-	-	E5.3	304	-	-	-	-	E5.6	304

Table 10.9 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 43 - 48

PTI 172703

Configurations	Error Conditions											
	49. Transmit Parity		50. Laser Fault		51. Out of Sync		52. Loss of Signal					
Fabric / Pt-to-Pt	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Class 2	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Class 3	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Responder												
Class 1	-	-	E5.8	305	E5.9	306	E5.10	306				
Class 2	-	-	E5.8	305	E5.9	306	E5.10	306				
Class 3	-	-	E5.8	305	E5.9	306	E5.10	306				
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported for Command Sequence											
Class 2	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Class 3	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Data Phase												
Class 1	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Class 2	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Class 3	E5.7	305	E5.8	305	E5.9	306	E5.10	306				
Responder Data Phase												
Class 1	-	-	E5.8	305	E5.9	306	E5.10	306				
Class 2	-	-	E5.8	305	E5.9	306	E5.10	306				
Class 3	-	-	E5.8	305	E5.9	306	E5.10	306				
Status Phase												
Class 1	Class 1 is not supported for Status Sequence											
Class 2	-	-	E5.8	305	E5.9	306	E5.10	306				
Class 3	-	-	E5.8	305	E5.9	306	E5.10	306				

Table 10.10 Tachyon Fabric / Pt-to-Pt Configuration Error Conditions 49 - 52

PTI 172704

Configurations	Error Conditions											
	1. No BB_Credit for Tx		2. Link Down		3. Any SOFx1		4. Init failure due to no response		5. Loop re-init		6. LIPf received	
Loop	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Class 3	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Responder												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Class 3	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Class 3	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Data Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Class 3	E6.1	307	E6.2	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Responder Data Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.1	307	E6.3	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Class 3	E6.1	307	E6.3	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Status Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.1	307	E6.3	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309
Class 3	E6.1	307	E6.3	307	E6.4	307	E6.5	308	E6.6	308	E6.7	309

Table 10.11 Tachyon Loop Configuration Error Conditions 1 - 6

PTI 172705

Configurations	Error Conditions											
	7. On Loop, Non-Partici-pating		8. Lost AL_PA during re-init		9. Failed to acquire desired AL_PA		10. State Timeout		11. OPN/ CLS		12. OPN returned, missing node	
Loop	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #	Err #	Pg #
Non-SCSI Initiator												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Class 3	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Responder												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Class 3	E6.8	309	E6.9	309	E6.7	309	E6.10	310	-	-	-	-
SCSI Initiator Cmd Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Class 3	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Data Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Class 3	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Responder Data Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Class 3	E6.8	309	E6.9	309	E6.7	309	E6.10	310	-	-	-	-
Status Phase												
Class 1	Class 1 is not supported on Loop											
Class 2	E6.8	309	E6.9	309	E6.7	309	E6.10	310	E6.11	310	E6.12	311
Class 3	E6.8	309	E6.9	309	E6.7	309	E6.10	310	-	-	-	-

Table 10.12 Tachyon Loop Configuration Error Conditions 7 - 12

PTI 172706

Configurations	Error Conditions					
	13. NOS/ LOS is received					
Loop	Err # Pg #					
Non-SCSI Initiator						
Class 1	Class 1 is not supported on Loop					
Class 2	E6.13 311					
Class 3	E6.13 311					
Responder						
Class 1	Class 1 is not supported on Loop					
Class 2	E6.13 311					
Class 3	E6.13 311					
SCSI Initiator Cmd Phase						
Class 1	Class 1 is not supported on Loop					
Class 2	E6.13 311					
Class 3	E6.13 311					
Data Phase						
Class 1	Class 1 is not supported on Loop					
Class 2	E6.13 311					
Class 3	E6.13 311					
Responder Data Phase						
Class 1	Class 1 is not supported on Loop					
Class 2	E6.13 311					
Class 3	E6.13 311					
Status Phase						
Class 1	Class 1 is not supported on Loop					
Class 2	E6.13 311					
Class 3	E6.13 311					

Table 10.13 Tachyon Loop Configuration Error Condition 13

PTI 172707

10.5 Descriptions of Error Conditions

The following sub-sections (referenced from the Error Condition Tables in the preceding section) describe the various error conditions.

Each description includes the following:

- The cause or reason for the error notification
- Tachyon and the host recovery steps to return Tachyon to a normal operating state
- Any additional recovery information for the host to prevent other errors or protocol violations from occurring due to the error being described

The host must perform all exchange management functions for any sequence with errors. This management process is outside the scope of Tachyon operation, and is not discussed in this user's manual.

10.5.1 Error Conditions E1.1-E1.12

E1.1 Timeout Due to Lost ACK or No EE_Credit

Tachyon runs a timer for sequences that it transmits. Tachyon sets the ACK Time Out bit in the outbound_completion message if the final ACK for the sequence is not returned within ED_TOV of sending the final frame of the sequence. Tachyon sets the ACK Time Out bit if a frame for the sequence is delayed because of a lack of EE_Credit for longer than ED_TOV.

Tachyon Operation

- Send an outbound_completion message to the host with the ACK Time Out bit set to one, the Class 1 Connection Open bit set to one or cleared to zero as appropriate, and Tachyon's remaining EE_Credit count.
- Freeze OSM.
- Send any subsequent ACKs received for the sequence to the host as described in error note "E2.1 Unexpected ACK".

Host Recovery

- Perform Abort Sequence (ABTS)/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E1.2 No BB Credit for Transmit

Tachyon has a timer that monitors when BB_Credit is zero while in N_Port ACTIVE state. The Credit Error bit in the Frame Manager Status register indicates that Tachyon is connected to a fabric (or an N_port) and the fabric (or the N_port) has not returned BB_Credit within ED_TOV. Refer to "E6.1 No BB Credit for Transmit on FC-AL" on page 307. Also, refer to "E5.4 Timeout Due to Outbound Transmission Delay" on page 304.

Tachyon Operation

If Tachyon has zero BB_Credit for ED_TOV while in the N_Port ACTIVE state, Tachyon:

- Sets the Credit Error bit in the Frame Manager Status register.
- Sends a frame_mgr_interrupt completion message to the host.
- Continues to set the Credit Error bit in the Frame Manager Status register every ED_TOV until BB_Credit is replenished.

Host Recovery

- Read the Frame Manager Status register to determine the cause of the Frame Manager interrupt.
- Perform a link reset by writing the Link Reset command to the Frame Manager Control register.
- After Tachyon completes Link Reset, perform Link Reset recovery. Refer to "E2.11 Link Down With Class 1 Connection" on page 285, "E2.12 Link Down During Class 2 or 3 Sequences" on page 286, and "E2.13 Link Down During SCSI Inbound Processing" on page 287.

PTI 172709

E1.3 Receipt of ACK With Abort Bits = 01, 10, or 11

Tachyon supports ACK with abort bits = 01. If an ACK is received with 10 or 11, Tachyon manages it as if it was set to 01.

Tachyon checks the ACK frame against the current sequence being sent by OSM. If the ACK does not match, Tachyon behaves as described in error note "E2.1 Unexpected ACK". If the ACK does match, Tachyon performs the processes as defined in the Tachyon Operation section below.

Tachyon Operation

Upon receipt of the first ACK with abort, Tachyon:

- Stops sending frames.
- Sends an outbound_completion message to the host with the Abort Requested bit set, the Class 1 Connection Open bit set or cleared as appropriate and Tachyon's remaining EE_Credit count.
- Freezes OSM.
- Sends any subsequent ACKs received for the sequence to the host described in error note "2.1 Unexpected ACK".

Host Recovery:

- Perform ABTS/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E1.4 ACK With History Bit = 0 received

An ACK with the history bit = 0 is a normal in order ACK and is treated under normal processing.

E1.5 ACK With History Bit = 1 received

An ACK with history bit = 1 indicates that at least one frame was received out of order on a remote node. Tachyon does not distinguish between in order ACKs or OOO ACKs. These ACKs are treated under normal processing.

Tachyon uses a method of simply counting the number of frames transmitted and the number of ACKs received, and when they match, credit is balanced.

E1.6 ACK_1 Received When Expecting ACK_0.

This condition is allowed only on error reporting. If a bad frame is received by the remote node, it should immediately ACK with the abort bits set to stop transmission of frames.

Tachyon checks the ACK frame against the current sequence being sent by OSM. If the ACK does not match, Tachyon behaves as described in error note "E2.1 Unexpected ACK". If the ACK does match, Tachyon performs the processes as defined in the Tachyon Operation section below.

Tachyon Operation

- If no abort bits are set in the ACK_1, Tachyon behaves as described in error note "E2.1 Unexpected ACK".
- If abort bits are set in the ACK_1, Tachyon behaves as described in error note "E1.3 Receipt of ACK With Abort Bits".

Host Recovery:

Refer to "E1.3 Receipt of ACK With Abort Bits = 01, 10, or 11" on page 277 and "E2.1 Unexpected ACK" on page 282.

E1.7 Receipt of P_RJT or F_RJT

Tachyon does not distinguish between P_RJT and F_RJT.

Tachyon checks the RJT frame against the current sequence being sent by OSM. If the RJT doesn't match, Tachyon behaves as described in error note "E2.1 Unexpected ACK". If the RJT does match, Tachyon performs the processes as defined in the Tachyon Operation section below.

Tachyon Operation

- Stop sending frames.
- Send the RJT frame to the host with an inbound_unknown_frame completion message so that the host can determine the cause of the reject.
- Send an outbound_completion message to the host with the Rejected bit set, the Class 1 Connection Open bit set to one or cleared to zero as appropriate, and Tachyon's remaining EE_Credit count.
- Freeze the OSM.
- Send any subsequent ACKs received for the sequence to the host aas described in error note "E2.1 Unexpected ACK".

Host Recovery

- The host should examine the rejected frames to determine the cause of the reject and determine the state of the sequence and the Class 1 connection. If no reject was received with EOFt, the sequence is indeterminate. If a reject was received with EOFt, the sequence is already terminated.
- Perform ABTS/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172711

E1.8 BSY

Tachyon does not distinguish between P_BSY and F_BSY.

Tachyon can handle BSYs for the first frame of the sequence differently than the other BSYs, depending on the programming of the Sequence Interlock bit in the ODB or SOFc1.

Tachyon checks the BSY frame against the current sequence being sent by the OSM. If the BSY does not match, Tachyon behaves as described in error note "E2.1 Unexpected ACK". If the BSY does match, Tachyon performs the processes as defined in the Tachyon Operation section below.

Tachyon Operation

- If the BSY is received for the first frame of a sequence, and the Sequence Interlock bit is set to one or the frame was an SOFc1, Tachyon re-sends the first frame of the sequence up to 16 times.
- If the BSY is for frames other than the first frame or the Sequence Interlock bit is not set to one, Tachyon sends the BSY to the host with an inbound_unknown_frame completion message. EE_Credit is not updated for BSYs sent to host..
- If a BSY is received, and the Retry Disable bit is set to one in the Tachyon Configuration register, Tachyon terminates the sequence without re-sending the frame, and behaves as in error note "E1.9 BSY Retries Exceeded".

Host Recovery

- If the host wants to re-transmit the busied frame, it can re-create the busied frame by copying the frame header and data to a single buffer and transmitting it via the HPCQ. If the frame re-transmission is successful, Tachyon receives the ACK and the sequence completes normally.
- If the host does not want to re-transmit, it simply waits for a timeout. Refer to "E1.1 Timeout Due to Lost ACK or No EE_Credit" on page 276.

E1.9 BSY Retries Exceeded

Tachyon retries sending the first frame of an X_ID interlock sequence or SOFc1 frame up to 16 times. If the 16th try is busied, an error condition is flagged.

Tachyon Operation

- Send an outbound_completion message to the host with the Retries Exceeded bit set and the Class 1 Connection bit set or cleared as appropriate.
- Freeze the OSM..

Host Recovery

Note that the sequence was not established in this case and therefore the host does not need to perform the ABTS/exchange error recovery protocol.

- Perform Sequence Recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172712

E1.10 Receipt of P_RJT or F_RJT on SCSI Outbound Operation

Tachyon does not distinguish between P_RJT and F_RJT.

Tachyon checks the RJT frame against the current sequence being sent by the OSM. If the RJT does not match, Tachyon behaves as described in error note "E2.1 Unexpected ACK". If the RJT does match, Tachyon performs the processes as defined in the Tachyon Operation section below.

Tachyon Operation

- Stop sending frames.
- Send the RJT frame to the host with an inbound_unknown_frame completion message so that the host can determine the cause of the reject.
- Sends an outbound_completion message to the host with the Rejected bit set, the Class 1 Connection Open bit set to one or cleared to zero as appropriate, Tachyon's remaining EE_Credit count and the Trans_ID of the Outbound SEST Entry that was being processed when the error occurred.
- Freeze the OSM.
- Send any subsequent ACKs received for the sequence to the host as described in error note "E2.1 Unexpected ACK".

Host Recovery

- The host should examine the rejected frames to determine the cause of the reject and to determine the state of the sequence and the Class 1 connection. If no reject was received with EOFt, the sequence is indeterminate. If a reject was received with EOFt, the sequence is already terminated.
- Perform ABTS/exchange error recovery processing and connection recovery as needed. The host should also clear the STE Valid bit in theOutbound SEST entry.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E1.11 BSY Retries Exceeded on SCSI Data

Tachyon retries sending the first frame of an X_ID interlock sequence or SOFc1 frame up to 16 times. If the 16th try is busied, an error condition is flagged.

Tachyon Operation

- Send an outbound_completion message to the host with the Retries Exceeded bit set and the Trans_ID of the outbound SEST entry that was being processed when the error occurred.
- Freeze the OSM.

Host Recovery

The sequence was not established, however the exchange was established by the CMD sequence, therefore the exchange must be terminated.

- The host must termintate the SCSI exchange and perform connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172713

E1.12 RJT or BSY on ACK Frames

The fabric can reject (RJT) or busy (BSY) an ACK frame. Tachyon sends theses frames to the host as inbound unknown frames. The host is expected to simply discard them and have the remote node time out.

Tachyon Operation

- Send the RJT or BSY frame to the host with an inbound_unknown_frame completion message.

Host Recovery

- The host should discard the frame.
- No recovery is required at Tachyon's level.

PTI 172714

10.5.2 Error Conditions E2.1-E2.26

E2.1 Unexpected ACK

An unexpected ACK is any Link Control frame (ACK, BSY or RJT) that does not match the current outbound sequence being sent by the OSM.

Tachyon Operation

- Tachyon sends an unexpected ACK, BSY, or RJT to the host with an inbound_unknown_frame completion message.

Host Recovery

- The host must check the Link Control frame against any possible abnormally terminated sequences it is tracking in exchange status blocks.
- No recovery is required at Tachyon's level.

E2.2 Unexpected EOFdt During Class 1 Outbound Connection

The only time an EOFdt is expected is with the final ACK for the current outbound Class 1 sequence. If a frame with an EOFdt is received while an outbound Class 1 connection is open and the EOFdt is not received on the final ACK for the sequence, the Class 1 connection is closed and an error condition is flagged.

Tachyon Operation

- If the frame is an Link Control frame that does not match the current outbound sequence, in addition to the following, Tachyon behaves as described in error note "E2.1 Unexpected ACK".
- Terminate the outbound connection.
- If any outbound sequence is in progress (regardless of class), Tachyon terminates the sequence, sends an outbound_completion message to the host with the Class 1 Error bit set, the Class 1 Connection Open bit cleared and the remaining EE_Credit count.
- Freeze the OSM.

Host Recovery

- Perform ABTS/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172715

E2.3 Unexpected EOFdt During Class 1 Inbound Connection

If a frame with an EOFdt is received while an inbound Class 1 connection is open, the Class 1 connection is closed and an error condition is flagged.

Tachyon Operation

- If the frame is a Link Control frame that does not match the current outbound sequence, in addition to the following, Tachyon behaves as described in error note "E2.1 Unexpected ACK".
- If an inbound MFS is in progress, it is terminated and an inbound_mgr_completion message is sent to the host with the Class 1 Error bit set.
- The inbound Class 1 connection is closed.

Host Recovery

- The host must check the Link Control frame against any possible abnormally terminated sequences it is tracking in exchange status blocks.
- No recovery is required at Tachyon's level.

E2.4 Unexpected ACK with EOFt During Transmission

An ACK with an EOFt indicates that at least one ACK frame was received out of order by Tachyon. Tachyon ignores the EOFt delimiter on ACKs.

Tachyon counts the number of frames transmitted and the number of ACKs received, and when they match, credit is balanced.

E2.5 Unexpected EOFt During Reception

Tachyon ignores the EOFt delimiter on Class 1 and 2 data frames. An EOFt on a Class 3 data frame is considered normal operation.

E2.6 EOFdti During Class 1 Outbound Connection

If a frame with an EOFdti is received while an outbound Class 1 connection is open, the Class 1 connection is closed and an error condition is flagged.

Tachyon Operation

- Send the frame to the host with an inbound_unknown_frame completion message.
- Terminate the outbound connection.
- If any outbound sequence is in progress (regardless of class), Tachyon terminates the sequence, sends an outbound_completion message to the host with the Class 1 Error bit set, the Class 1 Connection Open bit cleared and the remaining EE_Credit count.
- Increment the Generated EOFa error count in the Frame Manager Link Error Status Counters #2 register.
- Freeze the OSM.

Host Recovery

- Perform ABTS/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E2.7 EOFdti During Class 1 Inbound Connection

If a frame with an EOFdti is received while an inbound Class 1 connection is open, the Class 1 connection is closed and an error condition is flagged.

Tachyon Operation

- Send the frame to the host with an inbound_unknown_frame completion message.
- If an inbound MFS is in progress, it is terminated and an inbound_mfs_completion message is sent to the host with the Class 1 Error bit set.
- Increment the Generated EOFa error count in the Frame Manager Link Error Status Counters #2 register.
- The inbound Class 1 connection is closed.

Host Recovery

- No recovery is required at Tachyon's level.

E2.8 EOFdti Without Class 1 Connection

Since the frame is invalid, it does not affect any inbound Class 2 or 3 sequence processing. The exception is if an SOFc1 is sent to establish an outbound Class 1 connection, and an ACK with EOFdti is received in response, Tachyon sends an outbound_completion message to the host with the Class 1 Error status bit set and freezes the outbound channel.

Tachyon Operation

- Send the frame to the host with an inbound_unknown_frame completion message.
- Increments the Generated EOFa error count in the Frame Manager Link Error Status Counters #2 register.
- If an SOFc1 frame was sent to establish an outbound Class 1 connection, Tachyon terminates the sequence, sends an outbound_completion message to the host with the Class 1 Error bit set and the Class 1 Connection Open bit cleared, and finally freezes the OSM.

Host Recovery

- If an outbound Class 1 sequence was aborted, perform sequence recovery processing.
- If an outbound Class 1 sequence was aborted, return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172717

E2.9 EOFa

Since the frame is to be aborted, it does not affect any sequence processing and is simply discarded.

Tachyon Operation

- Discard the frame.
- Increment the Received EOFa error count in the Frame Manager Link Error Status Counters #2 register.

Host Recovery

- No recovery is required at Tachyon's level.

E2.10 EOFni

Since the frame is invalid, it does not affect any sequence processing and is simply discarded.

Tachyon Operation

- Discard the frame.
- Increment the Generated EOFa error count in the Frame Manager Link Error Status Counters #2 register.

Host Recovery

- No recovery is required at Tachyon's level.

E2.11 Link Down With Class 1 Connection

A Link Down condition causes the ISM and the OSM to complete the processing of any pending frames and then to flush all FIFOs. Any Class 1 connection is closed.

Tachyon Operation

- The Frame Manager detects the Link Down condition and sets the Link Down bit in the Frame Manager Status register and sends a frame_mgr_interrupt completion message to the host.
- If an inbound MFS was in progress, Tachyon terminates the sequence and sends an inbound_mfs_completion message to the host with the Link Down error bit set.
- If an outbound sequence was in progress, Tachyon terminates the sequence and sends an outbound_completion message to the host with the Link Down error bit set.
- All internal FIFOs are flushed.
- The Class 1 connection is closed.
- Transmit BB_Credit is re-initialized to the value programmed by the host in the Frame Manager Configuration register.
- Freeze the OSM.

PTI 172718

Host Recovery

- Read the Frame Manager Status register to determine the cause of the Frame Manager interrupt.
- Verify that the link is up by examining the Frame Manager Status register's Loop and/or Port Machine states. If the link is not up, wait for another frame_mfs_interrupt completion message.
- Perform ABTS/exchange error recovery processing, connection recovery and login recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E2.12 Link Down During Class 2 or 3 Sequences

A Link Down condition causes the ISM and OSM to complete the processing of any pending frames and then to flush all FIFOs.

While the Fibre Channel standard states that a Class 2 or 3 sequence is not necessarily affected by a link reset, it does require that BB_Credit be reset to LOGIN values. To do this requires Tachyon to flush buffers, which causes timeouts due to lost frames. To alleviate extra on-chip error handling, Tachyon terminates any sequence in both the outbound and inbound direction.

Tachyon Operation

- The Frame Manager detects the LR and sets the Link Down bit in the Frame Manager Status register and sends a frame_mgr_interrupt completion message to the host.
- If an inbound MFS was in progress, Tachyon terminates the sequence and sends an inbound_mfs_completion message to the host with the Link Down error bit set.
- If an outbound sequence was in progress, Tachyon terminates the sequence and sends an outbound_completion message to the host with the Link Down error bit set.
- All internal FIFOs are flushed.
- Transmit BB_Credit is re-initialized to the value programmed by the host in the Frame Manager Configuration register.
- Freeze the OSM.

Host Recovery

Tachyon's management of sequences was abnormally terminated by the Link Down, but from Fibre Channel's standpoint, any active Class 2 or 3 sequences have not been terminated and will require the ABTS protocol to be executed.

- Read the Frame Manager Status register to determine the cause of the Frame Manager interrupt.
- Verify that the link is up by examining the Frame Manager Status register's Loop and/or Port Machine states. If the link is not up, wait for another frame_mfs_interrupt completion message.
- Perform ABTS/exchange error recovery processing and login recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172719

E2.13 Link Down During SCSI Inbound Processing

A Link Down condition causes Tachyon's ISM and OSM to complete the processing of any pending frames and then to flush all FIFOs.

A Link Down has no affect on the operation of the inbound SCSI management, with the exception that the Link Down flushes the inbound FIFO, which may contain SCSI frames. These frames are lost, which causes a ULP timeout to occur. Also, if the data transfer was occurring on a Class 1 connection, the connection is lost and the sequence is abnormally terminated. Tachyon does not send a completion message to the host for the SCSI sequence and the ULP timeout will detect the error.

Tachyon Operation

- Tachyon behaves as described in error note "E2.11 Link Down With Class 1 Connection" if a Class 1 connection open and as described in error note "E2.12 Link Down During Class 2 or 3 Sequences" if no Class 1 connection open.

Host Recovery

- Refer to "E2.11 Link Down With Class 1 Connection" on page 285 and "E2.12 Link Down During Class 2 or 3 Sequences" on page 286.

E2.14 Link Failure With Class 1 Connection

The Link Failure condition differs from the Link Down in severity. Link Failure is indicated when the Frame Manager sets the Link Failure bit in the Frame Manager Status register.

Refer to "E5.8 Laser Fault" on page 305, "E5.9 Out of Sync" on page 306, and "E5.10 Loss of Signal" on page 306.

Tachyon Operation

- The Frame Manager detects the Link Fail condition and sets the Link Down bit in the Frame Manager Status register and sends a frame_mgr_interrupt completion message to the host.
- If an inbound MFS was in progress, Tachyon terminates the sequence and sends an inbound_mfs_completion message to the host with the Link Down error bit set.
- If an outbound sequence was in progress, Tachyon terminates the sequence and sends an outbound_completion message to the host with the Link Down error bit set.
- All internal FIFOs are flushed.
- The Class 1 connection is closed.
- Transmit BB_Credit in the Frame Manager Configuration register is cleared to zero.
- Freeze the OSM.

PTI 172720

Host Recovery

- Read the Frame Manager Status register to determine the cause of the Frame Manager interrupt.
- Verify that the link is up by examining the Frame Manager Status register's Loop and/or Port Machine states. If the link is not up, wait for another frame_mfs_interrupt completion message.
- Set the BB_Credit in the Frame Manager Configuation register to one and perform login recovery as needed.
- Perform ABTS/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E2.15 Link Failure During Class 2 or 3 Sequences

The Link Failure condition differs from the Link Down in severity. Link Failure is indicated when the Frame Manager sets the Link Failure bit in the Frame Manager Status register.

Refer to "E5.8 Laser Fault" on page 305, "E5.9 Out of Sync" on page 306, and "E5.10 Loss of Signal" on page 306.

Tachyon Operation

- The Frame Manager detects the Link Fail condition and sets the Link Down bit in the Frame Manager Status register and sends a frame_mgr_interrupt completion message to the host.
- If an inbound MFS was in progress, Tachyon terminates the sequence and sends an inbound_mfs_completion message to the host with the Link Down error bit set.
- If an outbound sequence was in progress, Tachyon terminates the sequence and sends an outbound_completion message to the host with the Link Down error bit set.
- All internal FIFOs are flushed.
- Transmit BB_Credit in the Frame Manager Configuration register is cleared to zero.
- Freeze the OSM.

Host Recovery

Tachyon's management of sequences was abnormally terminated by the Link Failure, but from Fibre Channel's standpoint, any active Class 2 or 3 sequences may have not been terminated and will require the ABTS protocol to be executed.

- Read the Frame Manager Status register to determine the cause of the Frame Manager interrupt.
- Verify that the link is up by examining the Frame Manager Status register's Loop and/or Port Machine states. If the link is not up, wait for another frame_mfs_interrupt completion message.
- Set the BB_Credit in the Frame Manager Configuation register to one and perform login recovery as needed.
- Perform ABTS/exchange error recovery processing as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

PTI 172721

E2.16 Link Failure During SCSI Inbound Operations

The Link Failure condition differs from the Link Down in severity. Link Failure is indicated when the Frame Manager sets the Link Failure bit in the Frame Manager Status register.

Refer to "E5.8 Laser Fault" on page 305, "E5.9 Out of Sync" on page 306, and "E5.10 Loss of Signal" on page 306.

A Link Failure has no effect on the operation of the inbound SCSI management, with the exception that the Link Down flushes the inbound FIFO, which may contain SCSI frames. These frames are lost, which causes a ULP timeout to occur. Also, if the data transfer was occurring on a Class 1 connection, the connection is lost and the sequence is abnormally terminated. Tachyon does not send a completion message for the SCSI sequence and the ULP timeout will detect the error.

Tachyon Operation

- Tachyon behaves as described in error note "E2.14 Link Failure With Class 1 Connection" if a Class 1 connection open and as described in error note "E2.15 Link Failure During Class 2 or 3 Sequences" if no Class 1 connection open.

Host Recovery

- Refer to "E2.14 Link Failure With Class 1 Connection" on page 287 and "E2.15 Link Failure During Class 2 or 3 Sequences" on page 288.

E2.17 ABTS Received During Sequence Transmission

The host cannot stop the OSM from processing data. Because of this, if the remote node is trying to abort the sequence, then the host must wait until it receives an outbound_completion message from Tachyon. At that time, the host can perform the ABTS protocol. Any ABTS frames received by Tachyon do not affect the state of outbound sequences, i.e., ABTS frames received for an outbound sequence do not stop the sequence from being sent.

Tachyon Operation

- Send the ABTS frame to the host with an inbound_SFS_completion message.
- Tachyon ACKs the ABTS frame.

Host Recovery

- No recovery is required at Tachyon's level.
- Perform ABTS/exchange error recovery processing as needed.

PTI 172722

E2.18 ABTS Received During Non-SCSI Sequence Reception

An ABTS frame immediately terminates an inbound non-SCSI sequence.

Tachyon Operation

- If an inbound MFS is in progress, the sequence reassembly is terminated and an inbound_MFS_completion message is sent to the host with the Abort Requested bit set.
- Send the ABTS frame to the host with an inbound_SFS_completion message.
- Tachyon ACKs the ABTS frame.

Host Recovery

- No recovery is required at Tachyon's level.
- Perform ABTS/exchange error recovery processing as needed.

E2.19 ABTS Received During SCSI Sequence Reception

An ABTS frame immediately terminates a SCSI sequence.

Tachyon Operation

- Tachyon terminates SCSI exchange management on the affected exchange and clears the STE Valid bit to zero in the SEST entry.
- Send the ABTS frame to the host with a bad_SCSI_frame_completion message.
- Tachyon ACKs the ABTS frame only if the Bad SCSI Auto ACK bit is set in the Tachyon Configuration register.

Host Recovery

- No recovery is required at Tachyon's level.
- Perform ABTS/exchange error recovery processing as needed.

E2.20 ABTX Received During Sequence Transmission

The ABTX is not the method of exchange termination preferred by the Fibre Channel profile. An ABTS with the last sequence of exchange bit set is the preferred method. The ABTS/LS method lets Tachyon detect the end of the sequence and exchange.

The host cannot stop the OSM from processing data. Because of this, if the remote node is trying to abort the sequence, then the host must wait until it receives an outbound_completion message from Tachyon. At that time, the host can perform the ABTX protocol.

Tachyon Operation

- Send the ABTX frame to the host with an inbound_SFS_completion message.
- Tachyon ACKs the ABTX frame.

Host Recovery

- No recovery is required at Tachyon's level.
- Perform ABTS/exchange error recovery processing as needed.

PTI 172723

E2.21 ABTX Received During Non-SCSI Sequence Reception

The ABTX is not the method of exchange termination preferred by the Fibre Channel profile. An ABTS with the last sequence of exchange bit set is the preferred method. The ABTS/LS method lets Tachyon detect the end of the sequence and exchange.

Tachyon provides no mechanism for the host to prematurely terminate a non-SCSI sequence reassembly. The sequence must timeout or complete. The non-SCSI exchange management is done exclusively by the host.

Tachyon Operation

- Send the ABTX frame to the host with an inbound_SFS_completion message.
- Tachyon ACKs the ABTX frame.

Host Recovery

- No recovery is required at Tachyon's level.
- Perform ABTS/exchange error recovery processing as needed.

E2.22 ABTX Received During SCSI Sequence Reception

The ABTX is not the method of exchange termination preferred by the Fibre Channel profile. An ABTS with the last sequence of exchange bit set is the preferred method. The ABTS/LS method lets Tachyon detect the end of the sequence and exchange.

The ABTX payload indicates which OX_ID is to be terminated. This payload is not interpreted by Tachyon. This frame is processed by Tachyon as an SFS and passed to the host. The host must process the frame and then start the exchange abort process.

Tachyon Operation

- Send the ABTX frame to the host with an inbound_SFS_completion message.
- Tachyon ACKs the ABTX frame.
- Upon detection of the write to the Tachyon Flush SEST Cache Entry register, if the SEST entry exists in Tachyon's cache, it will write back the cached SEST entry with the STE Valid bit cleared.
- Send any subsequent frames received for the sequence to the host with a bad_scsi_frame completion message.

Host Recovery

- No recovery is required at Tachyon's level.
- The host gets the OX_ID out of the ABTX frame and clears the STE Valid bit in the SEST entry for this OXID and then writes the OXID to the Tachyon Flush SEST Cache Entry register.
- Perform ABTS/exchange error recovery processing as needed.

PTI 172724

E2.23 RMC Received as Class 1 Connection Initiator

The RMC frame has the E_C bit set. Setting this bit immediately terminates the outbound connection and any in-process outbound sequences.

Tachyon Operation

- Sends the RMC frame to the host with an inbound_unknown_frame completion message. The RMC frame is not acknowledged.

Host Recovery

- Perform a link reset by writing the Link Reset command to Frame Manager Control register.
- After Tachyon completes Link Reset, perform Link Reset recovery. Refer to "E2.11 Link Down With Class 1 Connection" on page 285, "E2.12 Link Down During Class 2 or 3 Sequences" on page 286, and "E2.13 Link Down During SCSI Inbound Processing" on page 287
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E2.24 RMC Received as Class 1 Connection Recipient (non-SCSI)

The RMC frame has the E_C bit set and the ACK sent has an EOFdt delimiter. This immediately terminates the inbound connection and any in-process inbound sequence.

Tachyon Operation

- Transmits an ACK for the RMC frame with an EOFdt delimiter.
- If an inbound MFS is in progress, it is terminated and an inbound_mfs_completion message is sent to the host with the Class 1 Error bit set.
- The inbound Class 1 connection is closed.
- Sends the RMC frame to the host with an inbound_sfs_completion message.

Host Recovery

- No recovery is required at Tachyon's level.
- Perform ABTS/exchange error recovery processing as needed.

PTI 172725

E2.25 RMC Received by Class 1 Connection Recipient (SCSI)

The RMC frame has the E_C bit set and the ACK sent has an EOFdt delimiter. This immediately terminates the inbound connection and any in-process inbound sequence.

Tachyon Operation

- If the Bad SCSI Auto ACK bit is set in the Tachyon Configuration register, transmit an ACK for the RMC frame with an EOFdt delimitter.
- If an inbound MFS is in progress, it is terminated and an inbound_mfs_completion message is sent to the host with the Class 1 Error bit set.
- The inbound Class 1 connection is closed.
- If the OX_ID of the frame is valid, the STE Valid bit is cleared and the SEST entry is updated.
- Send the RMC frame to the host with an inbound_bad_scsi completion message.

Host Recovery

- Perform ABTS/exchange error recovery processing as needed.

The host should perform the following steps if the Bad SCSI Auto ACK bit is clear in the Tachyon Configuration register:

- Perform a link reset by writing the Link Reset command to Frame Manager Control register.
- After Tachyon completes Link Reset, perform Link Reset recovery. Refer to "E2.11 Link Down With Class 1 Connection" on page 285, "E2.12 Link Down During Class 2 or 3 Sequences" on page 286, and "E2.13 Link Down During SCSI Inbound Processing" on page 287.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E2.26 LCR

LCR is not directly supported by Tachyon and the FCSI profiles do not allow its use. Tachyon does not ACK an LCR frame.

Tachyon Operation

- Send LCR frame to the host with an inbound_unknown_frame completion message.

Host Recovery

- No recovery is required at Tachyon's level.

10.5.3 Error Conditions E3.1-E3.10

E3.1 SOFc1 During Established Class 1 Connection

Receiving a frame with an SOFc1 delimiter while a Class 1 connection is open is an error condition.

Tachyon Operation

- Send frame to the host with a inbound_unknown_frame completion message (the frame is not acknowledge).

Host Recovery

- No recovery is required at Tachyon's level. However, the host should perform a Link Reset because the state of the connection is in question.
- Perform a link reset by writing the Link Reset command to the Frame Manager Control register.
- After Tachyon completes the Link Reset, perform Link Reset recovery. Refer to "E2.11 Link Down With Class 1 Connection" on page 285, "E2.12 Link Down During Class 2 or 3 Sequences" on page 286, and "E2.13 Link Down During SCSI Inbound Processing" on page 287.

E3.2 SOFi1 or SOFn1 as Class 1 Connection Initiator

Class 1 data frames (SOFi1 and SOFn1) received during an inbound Class 1 connection (Tachyon is the Class 1 connection recipient) is normal operation, unless the frame does not match the S_ID and D_ID of the current connection.

Receiving a Class 1 data frame while there is an outbound Class 1 connection (Tachyon is the Class 1 connection initiator) indicates that bi-directional Class 1 communication is occurring. Tachyon does not support bi-directional Class 1 connections, and this is a Tachyon error condition. The remote node may be trying to perform bi-directional Class 1 communication, is confused about the state of the connection or is performing some sort of error recovery. The host may need to perform a Link Reset to clean up the connection.

Tachyon Operation

- Send frame to the host with a inbound_unknown_frame completion message (the frame is not acknowledged).

Host Recovery

- No recovery is required at Tachyon's level.

In the case where the received frames are error recovery frames, the host must respond appropriately. If the received frames are truly out of context, the host is expected to perform a Link Reset because the state of the connection is in question.

PTI 172727

E3.3 SOFi1 or SOFn1 Received Without an Established Class 1 Connection

This is an error case regardless of any other Tachyon operation.

Tachyon Operation

- Send the frame to the host with a inbound_unknown_frame completion message (the frame is not acknowledged).

Host Recovery

- No recovery is required at Tachyon's level.

The host is expected to send a P_RJT for the frame via the High Priority Channel.

E3.4 Unexpected Frame With E_C as Class 1 Connection Initiator

The E_C bit is only valid if received in a Class 1 data frame. The E_C bit is ignored in frames of other classes. Receiving a Class 1 data frame with the E_C bit set while an outbound Class 1 connection is open is an error condition.

Tachyon performs as described in error note "E3.2 SOFi1 or SOFn1 as Class 1 Connection Initiator".

E3.5 Unexpected Frame With E_C as Class 1 Connection Recipient

Unexpected inbound frames are those that do not match the S_ID of the Class 1 connection initiator or those that are non-data frames meant to prematurely terminate the connection. These include RMC, NOP, BA_ACC, or BA_RJT.

Tachyon Operation

- If the S_ID does not match the S_ID of the Class 1 connection initiator:
  - Send the frame to the host with an inbound_unknown_frame completion message.
  - The frame is not acknowledged and the connection is unaffected.
- If the S_ID matches the S_ID of the Class 1 connection initiator:
  - If a multiframe sequence reassembly operation is in progress, it is terminated and an inbound_mfs_completion message is sent to the host with the Class 1 Error bit set.
  - If the frame is a SFS, it is sent to the host with an inbound_sfs_completion message.
  - ACK the frame with an EOFdt and remove the inbound connection.

Host Recovery

No recovery is required at Tachyon's level.

E3.6 Bi-directional SOFc1 as a Connection Initiator

As a connection initiator, the host must never establish a bi-directional connection. Tachyon does not support Class 1 bi-directional connections.

PTI 172728

E3.7 Bi-directional SOFc1 as a Connection Recipient

A bi-directional connect request is treated by Tachyon as a normal unidirectional connect request. The difference is that the host could use the connection to send data back, but this is not supported by Tachyon and it is expected that the host will not try to take advantage of the connection. Tachyon does not process Class 1 OCQ commands while the inbound Class 1 connection exists, so only the High Priority Channel could be used to send frames against the connection.

E3.8 NOP Frame Received

An NOP frame is treated as a normal single frame data sequence.

E3.9 Other Basic Link Services Frames Received

NOP, ABTS, and RMC are discussed in separate explanitory notes. Any other BLS frames (BA_ACC, BA_RJT, etc) are single frame data sequences and are processed as such.

E3.10 Extended Link Services Frames Received

All ELS frames are single frame data sequences and are processed as such.

10.5.4 Error Conditions E4.1-E4.17

E4.1 Receiver Timeout

Receiver timeouts do not affect outbound sequence management.

Tachyon Operation

- The multiframe sequence reassembly operation is terminated.
- An inbound_mfs_completion message is sent to the host with the Time Out bit set.

Host Recovery

- No recovery is required at Tachyon's level.

E4.2 Inbound Class 1 Inactivity Timeout

The Class 1 inactivity timer monitors when an inbound Class 1 connection is open (Tachyon is the Class 1 connection recipient) but no sequences are being received. When the Class 1 connection is open and no sequences have been received for ED_TOV time, Tachyon sends an inbound_C1_timeout completion message to the host to inform it of the idle Class 1 connection. The timeout does not affect the state of the connection. The completion message is simply a warning and only occurs once.

Tachyon Operation

- Send an inbound_C1_timeout_completion message to the host.

Host Recovery

- It is not considered a violation of the FCSI profiles to leave a connection open without transferring data for an ED_TOV time. The host may respond to the inbound_C1_timeout completion message by performing a Link Reset. The Link Reset would (abnormally) terminate the connection and enable Class 1 connections to other nodes.
- No recovery is required at Tachyon's level.

E4.3 Inbound BB Credit Violation

An inbound BB_Credit violation may occur if the Inbound Data FIFO overflows. If a credit violation results in a FIFO overflow, the FIFO simply discards the frame that caused the overflow and no notification is given to the host. The problem is that it may not be possible to store the address of the offender, therefore notifying the host is minimally useful.

Tachyon Operation

- If a frame is received that does not fit into the Inbound Data FIFO, it is discarded.
- There is no indication sent to the host that this has occurred.

Host Recovery

- No recovery is required at Tachyon's level.

PTI 172730

E4.4 Received Frame Too Long

The maximum receive frame length is 2076 (2048 payload plus FC header). This value is hard-coded in the Frame Manager and checked as frames arrive. The host must login with the 2048 byte maximum payload specified. When a node sends a longer frame, it is a protocol violation.

Tachyon Operation

- If a frame is received that is too long, it is discarded.
- Increment the Generated EOFa error count in the Frame Manager Link Error Status Counters #2 register.
- There is no indication sent to the host that this has occurred.

Host Recovery

- No recovery is required at Tachyon's level.

E4.5 Frame Received With Expiration / Security Header

Tachyon has no way to validate the security or timer fields of the header, so any frame received with this header is passed to the host to be processed. No ACK or other response frame is sent.

Tachyon Operation

- Send the frame to the host with a inbound_unknown_frame completion message (the frame is not acknowledged).

Host Recovery

- The host should send a response frame via the High Priority Channel.

No recovery is required at Tachyon's level.

E4.6 Frame Received for a Second MFS

Tachyon has resources to perform only one multiframe sequence reassembly at a time.

Tachyon Operation

If the Disable P_BSY bit is set in the Tachyon Configuration register, refer to "3.7.5 Multiframe Sequence, Deferred P_BSY Mode" on page 60.

If the Disable P_BSY bit is not set in the Tachyon Configuration register, or if the incoming frame is an SOFc1, Tachyon:

- Automatically sends a P_BSY response to the frame.
- Discards the frame with no notification to the host.
- Increments the P_BSYs Sent counter in the Tachyon Receive Frame Error Counter Register.

Host Recovery

If the Disable P_BSY bit is set in the Tachyon Configuration register, refer to "3.7.5 Multiframe Sequence, Deferred P_BSY Mode" on page 60.

If the Disable P_BSY bit is not set in the Tachyon Configuration register

- No recovery is required at Tachyon's level.

PTI 172731

E4.7 Class 2/3 MFS Received as Class 1 Connection Initiator

This is a normal operation.

E4.8 Class 2/3 MFS Received as Class 1 Connection Recipient (non-SCSI)

Tachyon can perform only one sequence reassembly at a time. Since an established Class 1 connection implies that a sequence can arrive at any time, Tachyon must set aside its reassembly resource to handle a Class 1 multiframe sequence. Therefore, any Class 2 or 3 mutliframe seqeunce frames that are received during an inbound Class 1 connection are handled as described in explanitory note "4.6 Frame Received for Second MFS".

E4.9 OOO Reception of a Class 1 Frame During Non-SCSI Operation

Reception of an OOO Class 1 frame is an error. This indicates that a frame has been lost in the fabric and will never be delivered.

Tachyon Operation

- Assemble the sequence as described in "3.7.4 Multiframe Sequence, Out of Order (OOO) Reception" on page 59.
- Send the ACK for the received frame and all subsequent frames received for the sequence to the sequence initiator with the abort bit set. Note that the last data frame of the sequence may have the E_C bit set, but Tachyon does not close the connection to allow for error recovery. This means Tachyon continues to send ACKs with EOFn until an ABTS is received, which is acknowledged with an EOFdt.

Host Recovery

The multiframe sequence terminates when either an ABTS frame is received or a timeout occurs due to the lost frame.

- No recovery is required at Tachyon's level.

E4.10 OOO Reception of a Class 2 or Class 3 Frame During Non-SCSI Operation

Tachyon cannot handle the reassembly of sequences whose frames arrive of out of order. Instead, it places the frames in host memory and provides the host with enough information to reassemble the sequence. This is described in "3.7.4 Multiframe Sequence, Out of Order (OOO) Reception" on page 59.

Tachyon Operation

- Assemble the frames for the sequence as described in "Multiframe Sequence, Out of Order (OOO) Reception".
- Set the history bit in the ACK. This bit is set in any ACK sent for subsequent frames received for the sequence, until all the frames are received. The last ACK sent for the sequence has the history bit cleared.

Host Recovery

No host recovery is required, but the host must scan the various completions to allow it to reassemble the sequence segments.

- No recovery is required at Tachyon's level.

PTI 172732

E4.11 OOO Reception of a Class 2 or Class 3 Frame During SCSI Operation

If SCSI assists are enabled and the OOO Reassembly Disable bit is not set in the Tachyon Configuration register, Tachyon reassembles out of order SCSI frames without host assistance.

Tachyon Operation

- Automatically places the frame payload in host memory where it belongs.
- Set the history bit in the ACK. This bit is set in any ACK sent for subsequent frames received for the sequence, until all the frames are received. The last ACK sent for the sequence has the history bit cleared.

Host Recovery

- No recovery is required at Tachyon's level.

E4.12 OOO Reception of a Class 2 or Class 3 Frame During SCSI In Order Reassembly Mode

If SCSI assists are enabled and the OOO Reassembly Disable bit is set in the Tachyon Configuration register, and a frame is received out of order, it is assumed that an error has occurred.

Tachyon Operation

- Send frame to the host with a bad_scsi_frame completion message.
- If the Bad Scsi Auto ACK bit is set in the Tachyon Configuration register, an ACK is sent with the history bit set.
- The STE Valid bit in the SEST entry for this exchange is cleared and the SEST entry is updated.

Host Recovery

- No recovery is required at Tachyon's level.

E4.13 Unknown Frame Received (non-SCSI)

Tachyon sends an inbound_unknown_frame completion message to the host when it receives a frame that it cannot interpret.

Inbound_unknown_frame completion messages are sent to the host when the following conditions occur:

1. If the ACK Disable bit in the Tachyon Configuration register is set.
2. If a frame with an EOFdti (invalid delimiter) is received.
3. If the SOF type of the frame is not recognized.
4. If the link control frame is not recognized (i.e., R_CTL [7:4] = link control, R_CTL [3:0] >= 0111, or DF_CTL [6] is set to EXP_SEC_HDR, or the R_CTL field is not DEV_DATA, VIDEO_DATA, BSC_LNK_SERVICE, or XTD_LNK_SERVICE, or LINK_CNTL)
5. If a Class 1 data frame is received, and no inbound Class 1 connection is open.
6. If the S_ID or D_ID of a Class 1 frame does not match those of the node the Class 1 connection is established with.
7. If a Class 1 data frame is received and an outbound Class 1 connection is open.
8. If an SOFc1 is received and a Class 1 connection is already open.
9. If an SOFc1 is received and Tachyon is in loop mode.
10. If a LNK_CNTL frame with incorrect Sequence Context (FCTL [22]) is received.
11. If a BSY is received for an outbound sequence which is not a BSY for the first frame of the sequence.
12. If a RJT for an outbound frame is received.
13. If an ACK with EOFdt is received which does not match the outbound sequence, arrives when OSM is in an unexpected state, or an inbound Class 1 connection is open (with no MFS reassembly in progress).

Tachyon Operation

- Send the frame to the host with a inbound_unknown_frame completion message (the frame is not acknowledged).

Host Recovery

- The host should interpret the frame's contents and respond with the appropriate response frame via the High Priority Channel.
- In general, when Tachyon receives an unknown frame, Tachyon does not enter an error state, and therefore does not need error recovery. A notable exception is when Tachyon receives an unknown frame that contains an EOFdt or EOFdti terminator which forces the termination of an open Class 1 connection. The host may need to perform a Link Reset or the connection recovery procedure.
- No recovery is required at Tachyon's level.

PTI 172734

E4.14 Bad SCSI Frame Received

SCSI frames are those frames that have the Non-SCSI Bit clear in the OX_ID (Refer to "Choosing X_ID Values" on page 145.). A bad_scsi_frame completion message is sent to the host when a SCSI data frame is received and any of the following conditions are true:

1. The OX_ID is outside the valid range of the SEST.
2. The SEST entry indicated by the OX_ID has the STE Valid bit clear. (This may occur if the exchange was terminated before Tachyon receives all of the frames of the exchange)
3. The Inbound Bit of the OX_ID does not match the SCSI Direction bit in the SEST entry.

Tachyon Operation

- Send the frame to the host with a bad_scsi_frame completion message.
- If the Bad Scsi Auto ACK bit is set in the Tachyon Configuration register, an ACK will be sent with the history bit set.

Host Recovery

No recovery is required at Tachyon's level.

E4.15 Frame From Unlogged-In Remote Node

Tachyon does not keep login parameters. Therefore, Tachyon does not discriminate between frames from logged in nodes and other nodes. Frames and sequences from a node that hasn't logged in are treated as though the node did log in.

E4.16 Frame Received for an Invalid Non-SCSI Exchange

Tachyon does not keep exchange information for non-SCSI exchanges. Therefore, Tachyon does not validate frames or sequences based on exchange state.

E4.17 Frame Received for an Invalid SCSI Exchange

Tachyon uses the SCSI Exchange State Table to track SCSI exchanges. If a frame for an invalid exchange is received, it is passed to the host for processing. Refer to "E4.14 Bad SCSI Frame Received" on page 302.

PTI 172735

10.5.5 Error Conditions E5.1-E5.10

E5.1 SCSI: Upper Layer Timeout

When the host experiences a ULP timeout for SCSI, the host must ensure that Tachyon is no longer processing frames for that exchange. This is accomplished by performing the following:

1. Clear the STE Valid bit in the SEST entry.
2. Read the Tachyon Flush SEST Cache Entry register and ensure the Update bit is clear.
3. Write the OX_ID and set the Update bit in the Tachyon Flush SEST Cache Entry register, forcing Tachyon to discard any internal status for the exchange.
4. Read the Tachyon Flush SEST Cache Entry register and ensure the Update bit is cleared (indicating the operation is complete) before using the SEST entry.

Tachyon Operation

Upon receipt of the Flush OX_ID operation:

- If Tachyon has the SEST entry cached on the chip, it writes back stored status and clears the STE Valid bit in the SEST entry.
- If Tachyon does not have the SEST entry stored internally, no action is taken on the SEST entry.
- Tachyon clears the Update bit in the Tachyon Flush SEST Cache Entry register.

Host Recovery

- Before writing the Flush SEST Cache Entry register, the host must clear the STE valid bit in the SEST entry and ensure the Update bit in the Flush SEST Cache Entry register is clear.
- No Host recovery of Tachyon is required after the Flush OX_ID.
- The host should poll the Update bit in the Flush SEST Cache Entry register to be sure the operation has been completed before using the SEST entry.

E5.2 SCSI Outbound: FCP_XFER_RDY Queued in Host When Exchange Abnormally Terminated

When Tachyon gets behind in processing outbound FCP_XFER_RDYs, it queues them in host memory using the SEST. If the exchange is terminated prematurely, the FCP_XFER_RDY request cannot be removed from the queue by the host. The host must invalidate the exchange by clearing the STE Valid bit in the SEST entry. When Tachyon processes the FCP_XFER_RDY, it sees the STE Valid bit is clear and does not process the FCP_XFER_RDY.

Tachyon Operation

- Read the SEST entry for the queued FCP_XFER_RDY and notes the cleared STE Valid bit. Program the OSM with an invalid transfer request.
- The OSM aborts the transfer request and sends an outbound_completion_message to the host with the Host Programming Error bit set and the Trans_ID of the SEST entry that was being processed.
- Freeze the OSM.

PTI 172736

Host Recovery

- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E5.3 SCSI Frame Received Does Not Match SEST Direction bit

When a SCSI frame is received by Tachyon, the Direction Bit in the SEST is checked against the "I" bit of the frame's OX_ID to be sure the direction (inbound or outbound) matches (Refer to "Choosing X_ID Values" on page 145.). If a match does not occur, Tachyon manages it like an invalid SEST entry. Refer to "E4.14 Bad SCSI Frame Received" on page 302.

E5.4 Timeout Due to Outbound Transmission Delay

Tachyon keeps track of the time between when a frame is DMA'd from host memory to when it is actually transmitted on the link. If this time exceeds ED_TOV, a timer within Tachyon expires. Tachyon sends an outbound_completion message to the host with the Frame Time Out bit set. The Frame Time Out bit indicates that more frames were supposed to be sent but the current sequence exceeded the ED_TOV time for the transmission of the frame.

Tachyon Operation

- Send an outbound_completion message to the host with the Frame Time Out bit set.
- Freeze the OSM.

Host Recovery

- Perform ABTS/exchange error recovery processing and connection recovery as needed.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E5.5 Host Programming Error

The Host Programming Error bit in the outbound_completion message indicates that Tachyon was given conflicting parameters within an outbound request. This situation should only occur as described in error note "E5.2 SCSI Outbound: Data Descriptor Queued in Host When Exchange Abnormally Terminated". If it occurs under any other condition, the host should examine the ODB and EDB it sent to Tachyon to determine where the conflict is.

E5.6 Elastic Store Error

The Elastic Store Error bit in the Frame Manager Status register indicates that Tachyon has experienced an overflow or underflow in its Inbound Data FIFO. This condition occurs if the remote or local clock frequency is out of tolerance as specified by the Fibre Channel Specification or if a remote node is defective and is continually transmitting data (since primitive sequences are required to keep the elastic store centered). This error requires that hardware must be replaced on the remote node or the local node.

Tachyon Operation

- No recovery needed.

Host Recovery

- Replace the bad hardware.

PTI 172737

E5.7 Transmit Parity Error

The Transmit Parity Error in the Frame Manager Status register is a fatal error. It indicates that the transmit data path has experienced an internal corruption. This is not a recoverable error and Tachyon halts data transfer operation to prevent the transmission of bad data.

Tachyon Operation

- Tachyon asserts the ERROR_L signal on the TSI bus.
- Tachyon halts DMA operation, but responds to slave accesses of its status registers to inform the host of what has happened.

Host Recovery

- The host must either assert the RESET_L signal to Tachyon or issue a Software Reset command to the Tachyon Control register.
- Re-initialize Tachyon.

E5.8 Laser Fault

The Laser Fault bit in the Frame Manager Status register indicates that the PLM has experienced a failure with its laser operation. This bit is a qualifier to further define the cause of a Link Failure. Refer to "E2.14 Link Failure With Class 1 Connection" on page 287, "E2.15 Link Failure During Class 2 or 3 Sequences" on page 288, and "E2.16 Link Failure During SCSI Inbound Operations" on page 289.

Tachyon Operation

- Sets the Laser Fault and Link Failure bits in the Frame Manager Status register
- Sends a frame_mgr_interrupt completion message to the host.

Host Recovery

- The host sets the Enable External Loopback bit in the Frame Manager Control register to turn on the E_WRAP signal for at least 5 ms to force the PLM to execute its laser circuitry reset. This may or may not clear the laser fault. If it does, then normal operation can resume, if not, the interface must be replaced.
- After the host clears the Enable External Loopback bit in the Frame Manager Control register, it may take up to 10.3 seconds (according to the GLM specification) for the Laser Fault condition to clear.
- The host may want to perform Link Failure Management.
- If Tachyon is an N_Port, the host must write a Clear LF command to the Frame Manager Control register to enable Tachyon to recover from the failure when the Laser Fault condition has been cleared.

E5.9 Out of Sync

The Out of Synchronization error bit indicates that Tachyon cannot acquire symbol lock with the incoming data stream for greater than RT_TOV time. Loss of synchronization can be due to a degraded input signal or a loss of the signal for a time shorter than that required to trigger the loss of signal condition. This bit is a qualifier to further define the cause of a Link Failure. Refer to "E2.14 Link Failure With Class 1 Connection" on page 287, "E2.15 Link Failure During Class 2 or 3 Sequences" on page 288, and "E2.16 Link Failure During SCSI Inbound Operations" on page 289.

Tachyon Operation

- Set the Out of Synchronization and Link Failure bits in the Frame Manager Status register.
- Send a frame_mgr_interrupt completion message to the host.
- Tachyon attempts to re-establish synchronization without host intervention.

Host Recovery

- The host may want to perform Link Failure Management.
- If Tachyon is an N_Port, the host must write a Clear LF command to the Frame Manager Control register to enable Tachyon to recover from the failure when synchronization has been re-established. If Tachyon is an L_Port, the host waits for Tachyon to re-establish synchronization and waits for the link to come up.

E5.10 Loss of Signal

The Loss of Signal bit in the Frame Manager Status register indicates that the input signal has been lost for an extended period of time. This bit is a qualifier to further define the cause of a Link Failure. Refer to "E2.14 Link Failure With Class 1 Connection" on page 287, "E2.15 Link Failure During Class 2 or 3 Sequences" on page 288, and "E2.16 Link Failure During SCSI Inbound Operations" on page 289.

Tachyon Operation

- Set the Loss of Signal and Link Failure bits in the Frame Manager Status register.
- Send a frame_mgr_interrupt completion message to the host.

Host Recovery

- The host may want to perform Link Failure Management.
- If Tachyon is an N_Port, the host must write a Clear LF command to the Frame Manager Control register to enable Tachyon to recover from the failure when the input signal has been restored.

10.5.6 Error Conditions E6.1-E6.13

E6.1 No BB Credit for Transmit on FC-AL

The Credit Error and Loop bits in the Frame Manager Status register indicates that Tachyon is connected to a remote node on the loop and is attempting to send a frame and has not been able to acquire BB_Credit within ED_TOV.

Tachyon Operation

If the timer expires while waiting to transmit a frame, Tachyon:

- Sets the Credit Error bit in the Frame Manager Status register.
- Sends a frame_mgr_interrupt completion message to the host.

Tachyon operation from the reinitialization is described in error note "E6.6 Loop reinitialization".

Host Recovery

- Read the Frame Manager Status register. The host can either wait for the frame to be sent or re-initialize the the loop in an attempt to clear the problem.
- Write the Initialize command to the Frame Manager Control register to reinitialize the loop. This clears a terminal congestion problem and allows forward progress.

If the host reinitializes the loop, the sequence transmission aborts and requires its own recovery.

Host recovery from the reinitialization is described in error note "E6.6 Loop reinitialization".

E6.2 Link Down During Class 2 or Class 3 Sequences on FC-AL

The Link Down error bit indicates that the sequence was abnormally terminated because the Frame Manager experienced a link failure, a link reset or a loop reinitialization. Refer to "E2.12 Link Down During Class 2 or 3 Sequences" on page 286 and "E2.15 Link Failure During Class 2 or 3 Sequences" on page 288.

E6.3 Link Down During SCSI Inbound Operations on FC-AL

A link down has no affect on the operation of the inbound SCSI management, with the exception that it flushes the inbound FIFO that could contain SCSI frames. These frames are lost, which in turn causes a ULP timeout to occur. Refer to "E2.13 Link Down During SCSI Inbound Processing" on page 287 and "E2.16 Link Failure During SCSI Inbound Operations" on page 289.

E6.4 Receipt of Any Class 1 (SOFx1) Frames on FC-AL

Tachyon does not support Class 1 operation on an arbitrated loop.

Tachyon Operation

- Send the frame to the host with an inbound_unknown_frame completion message.

Host Recovery

- No recovery is required at Tachyon's level.
- The host should send a RJT response frame via the High Priority Channel.

PTI 172740

E6.5 Initialization Failure, No Proper Response

Once the host has programmed Tachyon to go online, it verifies that a failure condition (loss of signal, loss of sync or laser fault) does not exist prior to attemting to initialize. Once the receiver is up Tachyon enters the Initialize state and attempts to connect to a loop . It waits for a LIP in the initialize state for twice RT_TOV. If a LIP is not received, Tachyon enters Old Port state and attempts to initialize as an N_Port.

Tachyon Operation

If no meaningful order set is received for either loop or an FC link once initialization is started, Tachyon:

- Sets the Link Failure bit in the Frame Manager Status register.
- Sends a frame_mgr_interrupt completion message to the host.

Host Recovery

- Read the Frame Manager Status register. The Frame Manager Status indicates link failure.
- Write the Clear LF command to the Frame Manager Control register to allow Tachyon to resume initialization.

E6.6 Loop Reinitialization

Each time the loop re-initializes, Tachyon's loop address could change. The host must check that this has not happened. The following operation assumes that Loop Initialization successfully completed once.

Tachyon Operation

- When the loop has reinitialized, the Link Down bit in the Frame Manager Status register is set.
- Send a frame_mgr_interrupt completion message to the host.
- When initialization is complete, the Link Up bit in the Frame Manager Status register is set.
- Send a frame_mgr_interrupt completion message to the host (unless the host has not read the Frame Manager Status since the previous frame_mgr_interrupt completion message was sent).

Host Recovery

- Read the Frame Manager Status register, note that the link is down and wait for it to come back up.
- If the link is not up, wait for the second completion message, read the Frame Manager Status register, and note that the link is up.
- Read the Frame Manager Configuration register and check that the AL_PA is still the same.
- Perform login recovery and sequence recover as needed.

PTI 172741

E6.7 LIP Failure Received

When a loop node experiences a failure, it restarts initialization. To do this it sends out a LIP ordered set with the address portion of the LIP set to its AL_PA. This can be used by other nodes to determine where the failure is taking place.

Tachyon Operation

- When the LIP Failure is received, the LIPf and Link Down bits in the Frame Manager Status register are set.
- Sends a frame_mgr_interrupt completion message to the host.
- The AL_PA from the LIP Failure is copied to the Frame Manager AL_PA Received register.

Host Recovery

- Read the Frame Manager Status register, note that the link is down and that a LIPf was received.
- If the host wishes to monitor who is experiencing problems, it reads the Frame Manager AL_PA Received register. With this AL_PA information, the host may find the location of the fault. Note that if a valid AL_PA has not be acquired by the device sourcing the LIP Failure, a 0xF7 is read.
- At this point the host either waits until the loop comes back up, or it tries to do other loop recovery based on the LIPf information.

E6.8 Loop Initialized, Node is Non-Participating

When a loop initializes with more than 128 nodes, some of the nodes are not allowed to acquire an address. These nodes go non- participating. As each reinitialization occurs, they will again attempt to become participating.

Tachyon Operation

- Once initialization is complete the Non-Participating and Link Up bits in the Frame Manager Status register are set.
- Sends a frame_mgr_interrupt completion message to the host.

Host Recovery

- Read the Frame Manager Status register, note that the link is up but that the node is not participating. The host cannot send or receive data on this loop. There is no recovery from this condition, other than waiting for another node to drop out of the loop which frees up an AL_PA.

E6.9 Loop Re-initialized, Previously Acquired AL_PA Lost

When the loop reinitializes, it is possible that a fabric has been reassigning addresses and has taken the address that was assigned to this node. When this happens, the node reinitializes and acquires a new address. The host must check whether its address has changed. If the new address is acceptable, normal operation continues. If the host cannot tolerate its address changing, it may go non-participating. It does this by clearing all the address mode bits in the Frame Manager Configuration register and forcing another loop initialization by writing an Initialize Command to the Frame Manager Control register.

E6.10 Timeout on FC-AL

When Tachyon exceeds ED_TOV while in one of the following states: ARB, OPEN, OPENED, XMIT CLS, or RX CLS, Tachyon sends a frame_mgr_interrupt completion message to the host.

Tachyon Operation

- When Tachyon enters one of the following loop states: ARB, OPEN, OPENED, XMIT CLS, or RX CLS, it starts a timer.
- When ED_TOV time expires, Tachyon sets the Loop State Timeout bit in the Frame Manager Status register.
- Send a frame_mgr_interrupt completion message to the host.
- Tachyon clears and re-starts the timer. It will continue to set the Loop State Timeout bit in the Frame Manager Status register every ED_TOV as long as the loop stays in the state that started the timer.

Host Recovery

- Read the Frame Manager Status register to recognize that the Loop State Timeout bit has been set.
- At this point, the host may wait for another event to occur, or the host may re-initialize the loop if it thinks the loop is in a bad state.

E6.11 Loop Opened and Immediately Closed by Remote Node

Once Tachyon has won access to the loop, it sends an OPN to the remote node. If the remote node responds to the OPN with an immediate CLS, it is signaling that it cannot accept the frame at this time. Tachyon releases the loop and then attempts to reacquire it and OPN again. This cycle is repeated indefinitely or until a timeout occurs that causes the frame to be discarded. The host is unaware that this cycling is occurring.

PTI 172743

E6.12 OPN Sent (to a Bad AL_PA) and Immediately Received Back

When Tachyon sends an OPNfr (open broadcast replicate), a return is expected. However, if Tachyon sends an OPNxy that is returned, then the addressed node no longer exists on the loop.

Tachyon Operation

- The OSM freezes. Copy the AL_PA to the Frame Manager Received AL_PA register.
- Set the Bad AL_PA bit in the Frame Manager Status register.
- Send a frame_mgr_interrupt completion message to the host.
- If LCr equals zero, the frame remains in the Outbound Frame FIFO while the OSM is frozen. While the frame is in the Outbound Frame FIFO, the host cannot send frames via the High Priority Channel. The OSM unfreezes when the host sets the Error Release bit. Once the OSM unfreezes, discard the frame.

  If LCr is greater than zero, then the frame was already transmitted. When the Frame Manager receives the frame that it transmitted (since the destination AL_PA does not exist on the loop), discard this frame. No inbound completion message is generated for this received frame.
- If an outbound sequence is in progress, send an outbound_completion message with the Frame Time Out bit set to one.

Host Recovery

- Read the Frame Manager Status register and note that a Bad AL_PA was detected.
- Read the Frame Manager AL_PA Received Rgister to obtain the bad AL_PA and log it as no longer being valid.
- Return Tachyon to normal operation by setting the Error Release bit in the Tachyon Control register.

E6.13 NOS/OLS is Received

If a Tachyon initializes as an L_Port, and another Tachyon initializes as an N_Port, the L_Port Tachyon recognizes the OLSs and sets the NOS/OLS Received bit in its Frame Manager Status register. Setting the NOS/OLS Received bit indicates that a NOS and/or OLS was received after a loop circuit was established. The L_Port Tachyon then generates a frame_mgr_interrupt completion message to its host. The L_Port host can recognize this condition and, if no other ports exist on the loop, it should re-initialize the Frame Manager as an N_Port.

Appendix A. Application Notes

A.1 PCB Layout Suggestions

Tachyon Application Note #01
Revision 5.0
May 13, 1996

Introduction

This application note provides suggestions on the physical layout of a Printed Circuit Board (PCB) containing Tachyon and a Gigabit Link Module (GLM).

PCB Layers and Signals

The following PCB layers are suggested:

Layers	Description
Layer 1	Signal (Side A)
Layer 2	Ground
Layer 3	Signal
Layer 4	Signal
Layer 5	Power
Layer 6	Signal (Side B)

Table A.1.1 Layer Descriptions

Signal Routing

- The signals from Tachyon to the GLM should be routed on Layers 1 and 6.
- All transmit signals (i.e. TBC, TX [20..0]) should be routed on Layer 1. All receive signals (i.e. RBC, RX [20..0]) should be routed on Layer 6. Trace widths of these signals should be 6 mils with an 8 mil gap between them. Routing this way reduces the risk of crosstalk between transmit and receive.
- The Tachyon System Interface (TSI) side of Tachyon may be routed on all signal layers. Trace widths should be 8 mils with a gap of 8 mils.
- There is only a single plane for ground (Layer 2). All the Tachyon Vss signals connect to this plane, including the PLLVss signal.

Adapter Board Requirements

Delay Specifications

1. The maximum delay of the traces is 78.74 ps/cm (200 ps/in)
2. The minimum delay of the traces is 55.12 ps/cm(140 ps/in)

Receive Trace Specifications

1. RBC trace maximum length = [shortest RX trace + 3.81 cm]
2. RBC trace minimum length = [longest RX trace - 1.27 cm]

Transmit Trace Specifications

1. TX trace maximum length = 8.26 cm (3.25 in)
2. TX trace minimum length = 3.81 cm (1.50 in)
3. On all TX outputs, specifically TX [19..0], LCKREF_L, and EWRAP, Tachyon requires a minimum capacitance of 10 pF (GLM specification assumes a 10 pF load is present on Tachyon) and a maximum capacitance of 25 pF.

Clock Generator Trace Specifications

1. Clock Generator trace to Tachyon < [Clock Generator trace to GLM + 7.62 cm]
2. Clock Generator Rise Time = 2.8 ns

SCLK Routing to Tachyon and to the GLM

SCLK should be routed on Layer 1 with an 8 mil trace width. This trace must be as short as possible and have Tachyon as its only load. The trace should have guard traces along it.

Tachyon's GLM interface receives its transmit clock from an on-board oscillator. The output of this oscillator has two loads, Tachyon's TBC pin and the GLM interface's TBC pin. Thus the oscillator output must be split with one route through a 10 Ω resistor going to Tachyon, and the other route through another 10 Ω resistor going to the GLM pin. The route length to Tachyon and to the GLM transceivers should be equal. Also, the 10 Ω resistors should be near the oscillator output. These traces should be on Layer 1 and have guard traces along their length.

The 10 Ω resistor may have to be increased or reduced depending on signal quality at the load.

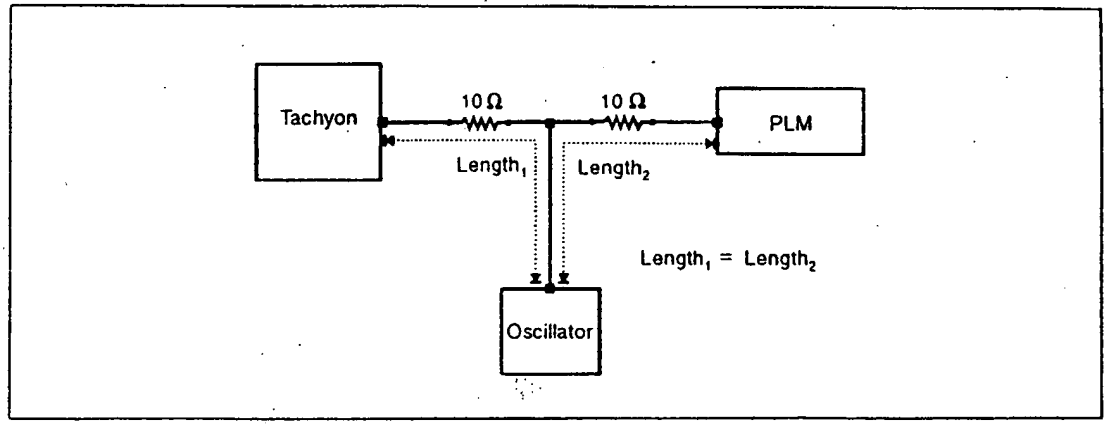

Figure A.1.1 SCLK Routing to Tachyon and to the GLM

PTI 172746

Calculation of Pull-up Resistor on the Tri-Stated I/O of Tachyon

Calculation for maximum pull-up resistor that can be used on the tri-stated I/O of Tachyon:

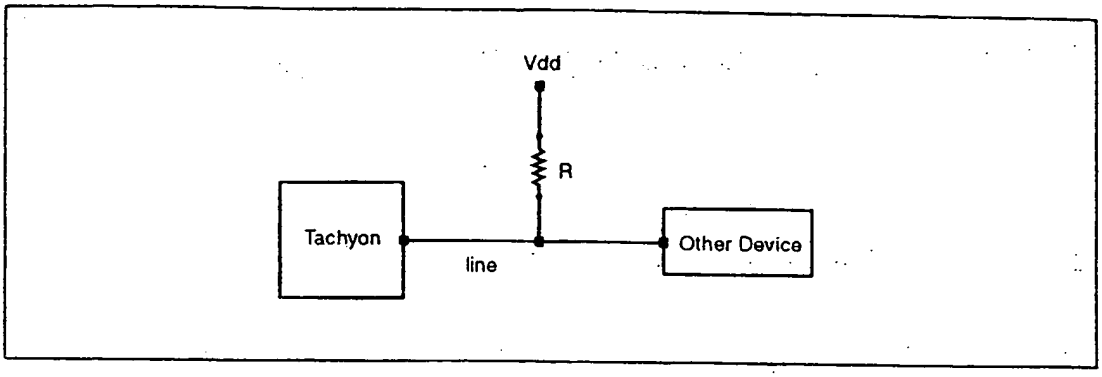

Figure A.1.2 Pull-Up Resistor on the Tri-Stated I/O of Tachyon

Values:

Ioz for Tachyon	=	10	uA
Ioz for Other Device	=	10	uA
Vdd minimum low	=	3.00	V
Vih minimum	=	2.00	V

Using the values above, the maximum value of R when the line is tri-stated to keep Vih at the minimum level is:

R = (3 - 2)Volts / 20 uA = 50 kΩ (this includes resistor tolerance)

GLM Signal Strappings Requirements

- TX_SI must be tied low with a maximum resistor value of 1 kΩ
- EN_CDET must be tied high with a maximum resistor value of 1kΩ

GLM Placement Recommendation

For optimal results, place Tachyon on the opposite side of the printed circuit board and behind the GLM Samtec connector. This GLM and Tachyon locality optimizes trace length and routing.

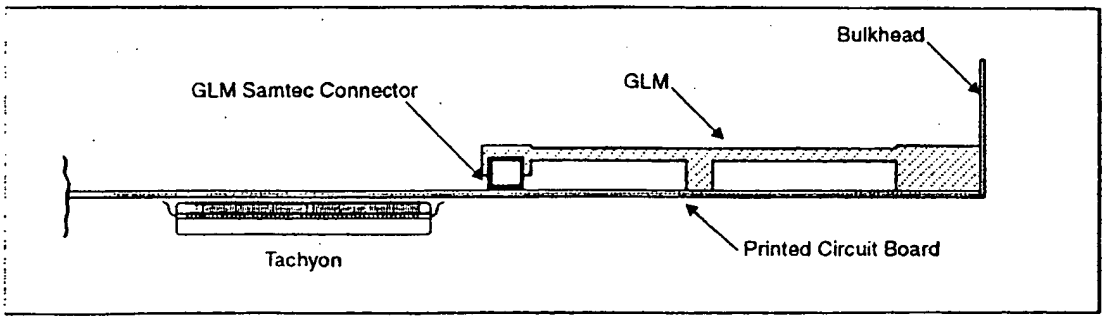

Figure A.1.3 GLM and Tachyon Locality for Trace Length and Routing Optimization

Power Connections

A continuous ground plane (Vss) should be provided for the entire PCB. The power plane (Vdd_Other) should be split into different regions to reduce noise and increase isolation between Tachyon and other components. Three methods are suggested:

1. **Ferrite Beads Method**

   This was the original recommendation for power connections and is adequate for existing designs.

   VddTAC is the region for Tachyon power. This region receives its power through a ferrite bead from Vdd_Other and requires decoupling capacitors. All of the Tachyon Vdd and Vdd2 pins connect directly to VddTAC, with the exception of PLLVdd, which connects to VddTAC through another ferrite bead. A 0.01 uf decoupling capacitor should be placed near the PLLVdd pin to provide decoupling to the ground plane. The ferrite beads should have a minimum saturation current of 2 Amps. A 25 mil trace (minimum, 100 mil preferred) should be used to connect PLLVdd through the ferrite bead to VddTAC. Tachyon Vss and PLLVss are both connected to the PCB ground plane.

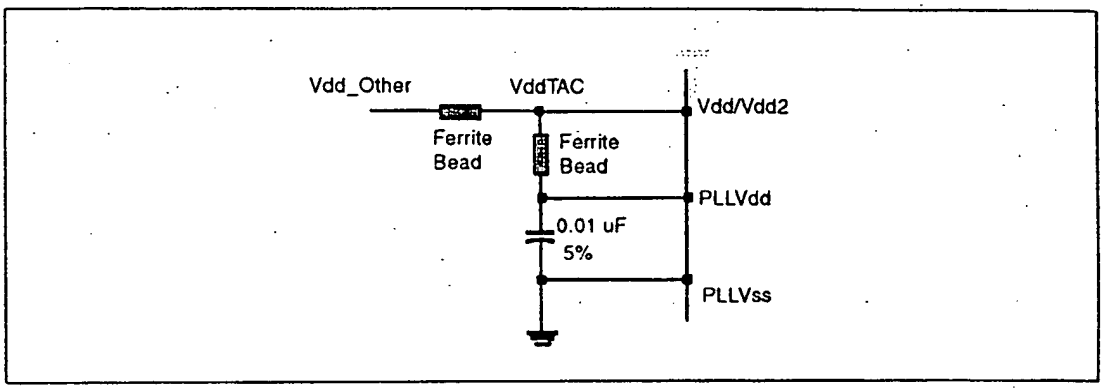

Figure 4 Ferrite Beads Method

The VddTAC plane resides under the Tachyon part as follows:

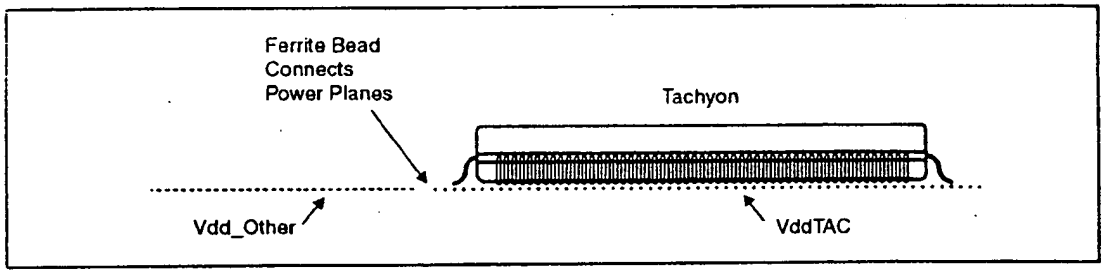

Figure A.1.5 Power Plane Regions

Power (Vdd and Vdd2) should be routed with a 100 mil trace width through fuses or ferrite beads prior to connecting into the plane through vias. Multiple vias are needed to meet current requirements for bringing power onto the plane.

PTI 172748

2. **PLL Ferrite Bead Alternative Method**

This is the current recommendation due to its optimal noise immunity to Tachyon's internal PLL circuitry.

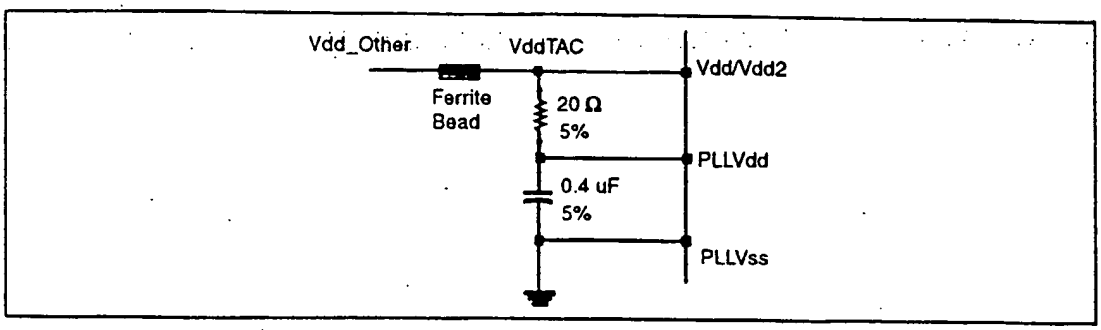

Figure A.1.6 PLL Ferrite Bead Alternative Method

3. **Multiple Ferrite Beads for Maximum Isolation Method**

This is recommended for systems that measure a large amount of noise on their circuits when using method 2 (PLL Ferrite Bead Alternative Method).

Tachyon has been designed with multiple internal power planes with separate input pins to isolate the output buffer/drivers, central core functions, and PLL oscillator. These power inputs are referenced as Vdd, Vdd2, and PLLVdd. In a typical situation, Vdd and Vdd2 are connected at the board level, with additional filtering for PLLVdd. Some applications may require additional isolation and improved noise immunity between Vdd and Vdd2. This alternative may be obtained by using multiple ferrite beads connecting each of the power regions (Vdd, Vdd2, PLLVdd) to Vdd_Others and additional decoupling capacitors for each region.

Separate low-inductance power regions for both Vdd and Vdd2 need to be provided on the PCB under the Tachyon package, separated from each other with multiple layers. The preferred arrangement is to have the Vdd and Vdd2 power regions on opposite sides of the PCB ground plane. Because PLLVdd is a single input into Tachyon, only a short, low-inductance trace needs to be provided from the ferrite bead to the PLLVdd pin. Decoupling capacitors for all three power regions should be located close to the Tachyon power input pins.

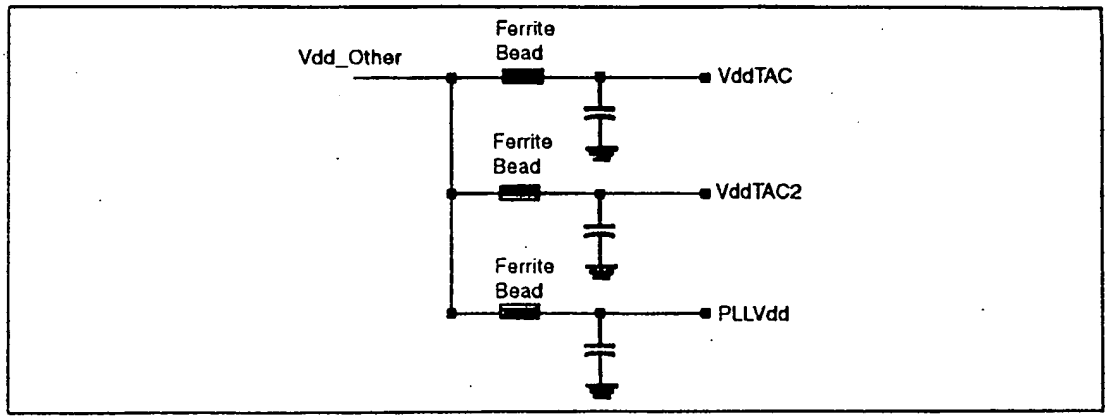

Figure A.1.7 Multiple Ferrite Beads for Maximum Isolation Method

PTI 172749

A.2 Implementing Tachyon with Little Endian System

Tachyon Application Note #02
Revision 3.0
May 13, 1996

To use Tachyon with a little endian system, implement the following steps:

Step 1:

Connect:

Tachyon's TAD pin 31 to bit 31 of the little endian system
Tachyon's TAD pin 30 to bit 30 of the little endian system
...
Tachyon's TAD pin 0 to bit 0 of the little endian system

For a 32-bit word, bit 31 is the most significant bit and bit 0 is the least significant bit.

Big Endian:

Most Significant Byte — Least Significant Byte

31			0
Byte 0	Byte 1	Byte 2	Byte 3

Little Endian:

Most Significant Byte — Least Significant Byte

31			0
Byte 3	Byte 2	Byte 1	Byte 0

By connecting Tachyon's TAD 31 to Little Endian's bit 31, the most significant byte of the Tachyon bus gets tied to the most significant byte of the Little Endian bus and the least significant byte of the Tachyon bus gets tied to the least significant byte of the Little Endian bus in a 32-bit transfer. Tying these buses together allows Tachyon to place any address that it wants to access on the correct byte lane (MSB to MSB, etc.). This causes a byte swap on all byte-wide transfers. 32-bit transfers are not swapped.

PTI 172750

Step 2:

Implement a Hardware swap of all data for DMA transfers. This should only be done on DMA transfers where Tachyon masters the bus and either reads or writes data to host memory.

Slave reads and writes should not be swapped and addresses should not be swapped.

All data transferred serially is transferred byte by byte. Therefore, data going out on the link is transferred byte-wide. For byte-wide transfers, the data was byte-swapped in Step 1 and must be swapped back.

Slave accesess are not byte-wide transfers and a byte swap is not needed. Slave accesses are 32-bit transfers and the byte lanes match.

DMA transfers may be byte-wide (data to be sent out on the link) or 32-bit (Tachyon control structure) transfers. However, the type of DMA transfer is indistinguishable and therefore all DMA transfers must be considered byte-wide. For byte-wide transfers, the data was byte-swapped in Step 1 and must be swapped back.

Step 3:

Host software must do a byte swap on all of Tachyon's DMA'd control structures. Anything Tachyon DMA's to or from the host must be swapped in software so Tachyon interprets it correctly. Data that is sent out on the link does not have to be swapped in software.

In Step 2, DMA link data and control structures transfers were byte-swapped. Either all the link data or all the control structures need to be swapped in software. Because there are fewer control structures than link data, it is more efficient to swap all the control structures rather than all of the data. To ensure that Tachyon can interpret its DMAed control structures properly, all the control structures must be swapped in software.

Step 4:

Since Slave reads and writes are not swapped by the Hardware swap (refer to 'Step 2:' above), the information written to Tachyon in Slave reads and writes does not have to be software swapped to be interpreted correctly by Tachyon.

Refer to Step 2 explanation on why Slave reads and writes do not need to be swapped.

Step 5:

Swapping is defined as follows:

Byte 0 is swapped with Byte 3

Byte 1 is swapped with Byte 2

This example attempts to explain the reasons the changes are necessary for Tachyon to be used with a Little Endian System.

There are three different types of data which move across the backplane both to and from Tachyon:

1. Slave reads and writes / 32-bit wide transfer
2. Link data (data that is sent over the link) / byte-wide transfer
3. Tachyon control structures / 32-bit wide transfer

PTI 172751

A.3 Limited Airflow Applications

Tachyon Application Note #03
Revision 3.0
May 13, 1996

When using Tachyon in applications where airflow is less than 1.5 meters/sec, careful attention must be given to ensure that Tachyon junction temperatures remain below 110 ° C. If junction temperatures exceed 100 ° C then the functional lifetime of the device may be reduced and/or potential data loss may result.

Junction Temperature Calculation

$T_{junction} = T_{ambient} + (P_{maximum} * \theta_{ja})$

$T_{junction}$ = The calculated junction temperature in ° C.

$T_{ambient}$ = The ambient air temperature around device in ° C.

$P_{maximum}$ = The maximum power dissipation of device in Watts. For Tachyon this is 3.5 Watts.

θ_{ja} = The package thermal resistance in ° C/Watt.

Package Thermal Resistance

The package thermal resistance, with respect to ambient air temperature, is highly dependent upon the airflow. The following table shows how the thermal resistance of the Tachyon package changes with respect to airflow. The table includes values with and without a heat sink attached.

Parameter	Conditions				Units
Airflow	0.0	1.0	2.0	3.0	meters/sec
208 MQuad	15.0	12.0	11.0	10.0	θ_{ja} – ° C/Watts
208 MQuad w/ Heat Sink	13.0	9.0	7.5	6.5	θ_{ja} – ° C/Watts

Table A.3.1 Package Thermal Resistance for 208 MQuad

Example Heat Sink:

Thermaloy P/N 18455B
Black Anodize Aluminum Omnidirectional Post Type
Dimensions : width 18 mm, length 21 mm, height 10 mm
Posts : 42 @ 1.2 mm x 1.2 mm x 8 mm
Approximate Effective Surface Area: 1990 mm sq.

The heat sink mentioned above is for reference only and is not available from Hewlett-Packard Company.

PTI 172752

A.4 Performance Suggestions

Tachyon Application Note #04
Revision 3.0
May 13, 1996

Tachyon's performance, i.e., its ability to move data to and from the host,can be improved by following the suggestions below.

1. Set Write and Read stream size fields in the Tachyon Configuration register to the highest value allowed by the application in which Tachyon is being used.
2. Always place data and headers on memory boundaries of either the read stream size value or data size, whichever is smaller. For example, program a Read Stream Size of four into the Tachyon Configuration register by writing 0x01 to the rs bits. The header length is 32 bytes and the EDB length is 2048 bytes. The header may be word aligned, but for performance reasons place it on a 32 byte boundary. If the header size increases above 32 bytes, place it on a sizeof( ) boundary, up to the streaming size. Since Tachyon holds the bus for only 128 bytes of data at a read stream of four, no increase in performance is obtained by placing the EDB pointers on boundaries greater than 128 bytes.
3. Choose values for the MFS Buffer Length, SFS Buffer length, and SCSI Buffer Length registers that allow the data to be stored in the fewest buffers possible. The buffers should be at least as big as the Write Stream Size.

Examples

With an airflow of 1.5 meters/sec and an ambient air temperature of 70 ° C, the junction temperature is calculated:

$T_{junction}$ = 70 ° C + (3.5 * 11.5) = 110.25 ° C

If the ambient air temperature is greater than 70 ° C or airflow is less than 1.5 meters/sec then a heat sink may be needed. The dimensions and surface area of the heat sink can be tailored to the specific application requirements.

An example of an application which requires special attention is where airflow is 0 meters/sec. There are two solutions to this problem:

1. Put an upper limit on the ambient air temperature: 110 ° C = $T_{ambient}$ + (3.5 * 15), solving for $T_{ambient}$ yields: $T_{ambient}$ = 110 ° C - (3.5 * 15) = 57.5° C.
2. Reduce θ_{ja} by using a heat sink which provides the required thermal resistance: 110 ° C = 70 ° C + (3.5 * θ_{ja}), solving for θ_{ja} yields: θ_{ja} = (110° - 70°) / 3.5 = 8.5° C per Watt.

PTI 172754

A.5 SCSI Command Forwarding

Tachyon Application Note #05
Revision 2.0
May 13, 1996

The SCSI Command Forwarding feature allows a controller to receive a SCSI command from a remote node and then forward it to one of its I/O devices. Tachyon can support this feature.

SCSI Command Forwarding can be used in an application where many I/O devices are controlled by a common controller. When Tachyon sends a command to a controller, the controller may forward the command to the appropriate device. Tachyon can then communicate directly with that device without needing the controller as an intermediary.

Initiator Read with Forwarding Controller Example

When Tachyon sends a FCP_CMND to the controller, the controller may determine that the command is for a specific device and may decide to forward the FCP_CMD to that device. The device returns the FCP_DATA and FCP_RSP directly back to Tachyon, bypassing the controller.

Initiator Write with Forwarding Controller Example

When Tachyon sends a FCP_CMND to the controller, the controller may determine that the command is for a specific device and may decide to forward the FCP_CMD to that device. The device sends a FCP_XFER_RDY directly back to Tachyon indicating that it is ready to receive the data. Tachyon uses the S_ID of that device to send the FCP_DATA and FCP_RSP directly to that device, bypassing the controller.

Note	SCSI Command Forwarding cannot be used on a Loop topology.

PTI 172755

Appendix B. Host Driver Notes

1. If the host has parity enabled, then the driver should properly configure parity within Tachyon via the Tachyon Configuration register. If parity is not properly configured before Tachyon drives signals on the address/data bus (TAD[31..0]), such as during a slave read, a parity error may result.
2. If the host must invalidate a SCSI Read Transaction, it must first clear the STE Valid bit in the SEST entry and then issue the Flush OX_ID command. Both steps must be done to guarantee that the entry is invalidated, since Tachyon may have the OX_ID cached internally.
3. For Class 1 operation, Tachyon only supports unidirectional connections. Transfer of initiative with data returning on the same connection is not allowed. The only exception is with error recovery sequences. For error recovery sequences, initiative must transfer in order to execute recovery procedures.
4. Tachyon does not support streamed sequences.
5. Do not change the uni-directional bit to bi-directional at any time during a Class 1 connection.
6. If Tachyon is connected point-to-point, Tachyon requires intermix support from the fabric or the remote node.
7. Send Networking and SCSI Link Level protocols on non-SCSI-Assisted OX_IDs only to prevent them from being interpreted as requiring FCP hardware assists.
8. For a non-SCSI-assisted transaction, set bit 15 in the OX_ID field in the Tachyon Header Structure. For SCSI-assisted transaction, clear bit 15 to zero.
9. Tachyon passes all unrecognized link control frames, e.g., LCR, RJTs, unprocessed BSYs, unrecognized ACKs, to the host without generating ACKs.
10. Tachyon accepts and acknowledges received frames, regardless of the D_ID. The host must check the D_ID to determine which service or alias was the destination of the frame.
11. Tachyon accepts and acknowledges FC-4 device data, video data, basic link services, and extended link services frames. Tachyon manages all other frames as unknown frames and passes them to the host without generating ACKs.
12. Tachyon writes a Relative Offset (RO) value into every transmitted frame except High Priority frames. The host must set the F_CTL bit, which indicates whether or not the destination should pay attention to the RO value.
13. If an error occurs on Class 1 sequences, the OSM suspends transmission of the current sequence and waits for error recovery. The host must reset the OCQ, unfreeze the OSM, and re-establish the queue entries if any Class 1 sequences in the queue expect an established connection. The host must modify the header of the first sequence on the queue to force it to re-establish the connection. It must do this before releasing the OSM lock.
14. All addresses given to Tachyon must be word-aligned and all lengths must be a multiple of 4 bytes. Sequences that require odd lengths must set the Fill bytes field in the ODB to account for extra pads bytes. The total length must be rounded up to the next 4-byte boundary to account for the odd bytes.
15. In the ODB, the Hdr_Len field is 32 bytes or more.
16. Zero length buffers are not allowed in the EDB entries.
17. The Tachyon header length programmed in the ODB must be less than or equal to the frame length. Tachyon does not divide or segment the Tachyon header A/L.

PTI 172756

18. It is recommended that the High Priority Channel only be used for error recovery or sending Link Control frames. It should not be used to establish or to destroy Class 1 connections. If a Class 1 connection must be closed abnormally, the host should not send a frame via the HPCQ with an EOFdt to close the connection. To close the Class 1 connection, the host should issue a link reset with to the Frame Manager. The High Priority Channel can be used for Class 3 small payloads for N_Ports. Refer to "3.6.8 Small Data Payload Transmit Processes" on page 47.

19. The OCQ, HPCQ, and SCSI Exchange Manager can each send sequence information to the OSM (Refer to "Figure 3.1 Tachyon Internal Block Diagram" on page 17.) If the OSM receives information from these three sources simultaneously, it prioritizes the information. The OSM transmits the HPCQ sequence information first, it then arbitrates fairly between the OCQ and SCSI Exchange Manager, and sends the associated network or SCSI information. However if a Class 1 connection is already opened, the OSM only accepts network ODBs.

20. When Tachyon is not in a point-to-point configuration and logs into a remote node that only supports one sequence, e.g., a remote node using Tachyon, the Sequence Interlock bit in the ODB should be set. Setting the Sequence Interlock bit enables Tachyon to re-send frames when the remote node's inbound channel is busy. This is especially applicable in a homogeneous network. Refer to "3.7.5 Multiframe Sequence, Deferred P_BSY Mode" on page 60.

21. When a write (outbound) requires more than one FCP_XFER_RDY for SCSI, the host must process all FCP_XFER_RDYs that are not processed by the card. The host must guarantee that the SEQ_ID it uses is different from the one that Tachyon uses for the first sequence.

22. Tachyon does not have a timer to monitor how long it takes the host to transmit data on a Class 1 or loop connection. It is up to the host to monitor itself not to send more data than is allowed by the FCSI Profile time limit.

23. At the start of a new exchange, the host should queue only the first sequence and wait for the completion message. In the completion message, a RX_ID that was assigned by the remote node for use in subsequent sequences is returned.

24. The host must set the BB_Credit parameter in the N_Port LOGIN frames to zero when connected to a loop to prevent over-running the inbound FIFO.

25. Tachyon does not have access to login parameters and therefore cannot validate frame data associated with login parameters.

26. The host should process completion messages quickly in order to replenish the IMQ and avoid blocking the inbound path.

27. The host should process all completion messages in the IMQ before writing consumer index entries to Tachyon. If the host writes a consumer index during message processing, Tachyon generates additional interrupts because the message queue does not appear empty.

28. Class 1 services are not supported on loop.

29. In order to send the Fabric Login frame when not connected to an arbitrated loop, the driver must set the fabric BB_Credit parameter in the Frame Manager Configuration register to one.

30. If a Tachyon initializes as an L-Port, and another Tachyon initializes as an N_Port, the L_Port Tachyon recognizes the OLSs and sets the NOS/OLS Received bit in its Frame Manager Status register. Setting the NOS/OLS Received bit indicates that a NOS and/or OLS was received after a loop circuit was established. The L_Port Tachyon then generates a frame_mgr_interrupt completion message to its host. The L_Port host can recognize this condition and, if no other ports exist on the loop, it should re-initialize the Frame Manager as an N_Port.

PTI 172757

31. The driver must manage sequence IDs across all aliases.

32. The host should not attempt to do a Flush OX_ID command for SCSI Write Transactions since Tachyon does not cache write entries on the chip.

33. When an ERROR_L is asserted by Tachyon, the host queries Tachyon's registers to determine the problem. During the query, if another error occurs, the host may hang. To avoid hanging, the host should implement a time out during the query.

34. For an outbound transaction, if an ACK-return-timeout occurs (ED_TOV expires), and more ACKs are received for the same transaction, the following may occur:

    a. An ACK timeout may occur but the ACK could arrive before the outbound _completion is generated. In this case, the completion may indicate a balanced credit and the timeout error.
    b. The host could possibly receive a completion for the unexpected ACK before it receives the completion for the ACK timeout. There is no guarantee of ordering of these two completions.

35. If the Disable AUTO P_BSY bit in the Tachyon Configuration register is set and the host initiates a link reset, the host should set the Clear Deferred P_BSY bit after successfully achieving link reset and emptying the Deferred P_BSY Queue. Refer to "3.7.5 Multiframe Sequence, Deferred P_BSY Mode" on page 60.

36. If an outbound Class 1 connection is opened and an unexpected frame with EOFdt is received, the OSM closes the connection, sets the Class 1 Error bit in the outbound_completion message, and freezes. If the host sends a Class 2 or Class 3 sequence while this Class 1 connection is opened, the sequence is aborted if the Class 1 connection closes due to an unexpected frame with EOFdt.

37. Reset the OCQ using the OCQ Reset bit in the Tachyon Control register only when the OSM is frozen.

38. An OCQ or HPCQ entry can be retrieved by Tachyon and buffered for the OSM even if the OSM is in the frozen state.

39. When the host receives an SOFc1 marked as a bad SCSI frame and the Bad SCSI Auto ACK bit in the Tachyon Configuration register is cleared to zero (which disables the ACKs), the host must initiate a link reset. Do not send an ACK to open the connection.

40. Tachyon does not guarantee the FC-PH maximum inter-frame delay (of ED_TOV) requirement for all frames of sequences sent using multiple ODBs. Tachyon adheres to the FC-PH Standard and guarantees to send all frames of an ODB within ED_TOV amount of delay between frames of the sequence. However, if the Continue Sequence bit in the ODB is set, Tachyon does not guarantee that the first frame sent from the subsequent ODB is transmitted before ED_TOV is exceeded.

41. If the Continue Sequence bit (Word 2/bit 20 in ODB) and ACK_0 bit (Word 2/bit 24 in ODB) are both set, a sequence timeout occurs. A sequence timeout occurs because the expected ACK can not be received in the last frame if the Continue Sequence bit has been set.

42. With an Out of Order Fabric using Class 3 FCP-Initiator assists, Tachyon may receive a Status frame before the data sequence completes. In this case, Tachyon prematurely notifies the host that the data sequence has completed even though the data sequence is incomplete. The remaining frames of the data sequence are sent to the SFSBQ as inbound bad SCSI frames. The host should continue to wait for the remaining frames until RA_TOV expires.

    If Tachyon receives these "missing" frames before RA_TOV (indicated by the Byte Count value in the Inbound SEST Entry plus the newly arrived frames equaling the expected number of bytes of the sequence), the frames may be reassembled with the rest of the data sequence. If Tachyon does not receive all the frames by RA_TOV (or the host does not want to manually reassemble the sequence), the host should start upper level error recovery.

PTI 172758

43. If Tachyon receives an ACK with the Abort Requested bit set or a RJT during an active outbound sequence, Tachyon immediately terminates the sequence, freezes its OSM, and sends an outbound_completion message to the host indicating the error. Tachyon does not wait for BB_Credit to balance, as stated in the Tachyon User's Manual.

44. If a Link Down occurs during the transmission of an outbound sequence before a successful outbound_completion message is sent to the host, the outbound_completion message will indicate a Link Down (LD) even though the previous sequence may have been transferred successfully. For example, if all ACKs were received for a sequence, but then the link goes down before the outbound_completion message is sent, the outbound_completion message will indicate a Link Down. When this situation occurs, the host cannot determine from the completion message whether the sequence was transferred successfully. The host must determine through Exchange Management if the sequence was transferred successfully. The host may choose to re-send the sequence.

45. If Tachyon runs out of SFS, MFS, or IMQ buffers, it will not initialize from link down until the buffers are replenished. In loop mode, this situation would keep the loop down.

46. Only multiframe sequences can be sent when the SOFc1 and E_C bits are both set to one.

PTI 172759

Appendix C. Building Data Structures

C.1 Tachyon Header Structure

Building the Tachyon Header Structure

The host should provide all the information to build the Tachyon Header Structure. The algorithm for building the Tachyon Header Structure can be simplified by always filling in the AL_PA field. Tachyon ignores this AL_PA field if it is not on a loop.

Algorithm for Building the Tachyon Header Structure

1. Set the SOF and EOF bits based upon the class type. For almost all sequences, these should be SOFI2, SOFI3, or SOFC1, and EOFN.
2. Write the AL_PA regardless of whether Tachyon is on the loop. This step is not required, but it is helpful in case it is decided later that Tachyon will be on a loop.
3. If Timestamps are being used, then a Tachyon Header Structure containing a valid Timestamp should have been received. A copy of that Timestamp should be used to send the frame. The Timestamp is used to prevent the frame from being sent if more than ED_TOV time has elapsed since it was received.
4. Write the R_CTL and the D_ID (the D_ID must match the D_ID field in the ODB).
5. Write the S_ID.
6. Write the F_CTL and Type fields (F_CTL should not have the End_Sequence bit set to one).
7. Set the SEQ_CNT to zero.
8. Select and write a unique Sequence ID for this exchange.
9. Write the RX_ID and OX_ID.
10. Write the RO and DF_CTL.

Assumptions for Building the Tachyon Header Structure

1. This algorithm is only applicable to normal I/O's processed by the OCQ.
2. Class 1 sequences always are multiframe sequences and are terminated at the end of a sequence, therefore SOFc1 (0x3) can always be written to the Tachyon Header Structure's SOF field.

C.2 ACKs

Algorithm for Building ACKs

Under normal conditions Tachyon generates ACKs. Under certain error conditions the driver must generate them. Use the following steps as a guide to build ACKs from Tachyon Header Structures for received frames.

1. Swap the D_ID and S_ID fields.
2. Change the R_CTL field to indicate an ACK frame.
3. In the F_CTL field of the Tachyon Header Structure
   a. Invert the Exchange Context (bit 23) and Sequence Context (bit 22) bits.
   b. Clear the Sequence Retransmitted bit (bit 9) to zero.
   c. Clear the Abort Sequence Condition bits (bits 5..4) to zero.
   d. For Class 1, set the Unidirectional Transmission bit (bit 8) to one.
   e. For Class 1 or Class 2, set the EOF field as appropriate.

PTI 172761

C.3 ODB

Building the ODB

Each ODB contains a Transaction Identifier (Trans_ID), which is used to match completion messages to requests. Tachyon returns this value on outbound completion messages. It is suggested that the Trans_ID be set to the ODB's entry number in the OCQ or a pointer to a driver managed structure. This enables matching completion messages to OCQ entries easier. It also enables error recovery and memory management to be easier.

Because the HPCQ and the OCQ have their own interrupts, their Trans_IDs do not have to be unique. However, if SCSI assists are used, the SEST uses the Outbound Message Channel to send its payload. This means that a SEST outbound completion and an OCQ outbound completion are indistinguishable from one another if they have the same Trans_ID.

Algorithm for Building the ODB.

1. Select the next free OCQ entry.
2. Build the Tachyon Header Structure.
3. Build the EDB.
4. Write the Fill bits (if any) along with the other control bits to the CNTL field of the ODB.
5. Write the SEQ_ID from the Tachyon Header Structure into the ODB.
6. Write the D_ID from the Tachyon Header Structure into the ODB.
7. Write the RX_ID from the Tachyon Header Structure into the ODB.
8. Write the Trans_ID field.
9. Write the Header Address Pointer (Hdr_Addr) into the ODB.
10. Write the EDB's address (EDB_Addr) into the ODB.
11. Write the Total Sequence Length (Tot_Len) into the ODB.
12. Write the Header Buffer Length (Hdr_Len) into the ODB.
13. Write the Maximum Frame Length (Frame_Len) from the Login parameters into the ODB. Use the Fiber Channel default parameters if performaing login.
14. Write the other fields as required.

C.4 EDB

Algorithm for Building the EDB.

1. Determine the number of EDB A/L pairs required.
2. Allocate memory for the EDB A/L pairs in multiples of four pairs (32 bytes).
3. Write the first Buf_Addr with the data's address.
4. Write the data length into the Buf_Len field.
5. If this is the last A/L pair, set the End Bit (E).
6. If it is not the last A/L pair, write the start of the next buffer into the next A/L pair.
7. Repeat steps 4. through 6. until all the data is described by the A/L pairs.
8. Write the Header Bit (H) and the Frame Boundary Bit (F), if required.

Assumptions for building the EDB

1. Data is contiguous for each A/L pair.
2. The data must be a multiple of four bytes, except for the last A/L pair. If the last A/L pair is not a multiple of four bytes, use pad bytes to fill the data to a 4-byte boundary.

C.5 High Priority Frame Structure

Algorithm for Building the High Priority Frame Structure

Refer to " Algorithm for Building the HPDB" on page 336.

PTI 172764

C.6 HPDB

Algorithm for Building the HPDB

Building the HPDB includes building the High Priority Frame Structure (HPFS). The algorithm for this process follows.

1. Select the next free HPCQ entry. This is the HPDB.
2. Allocate contiguous memory for the HPFS. This includes 32 bytes for the Tachyon Header Structure, a maximum of 2048 bytes for the payload, and 4 bytes for the EOF word at the end of the structure. The payload includes the data and all optional headers. The data must be a multiple of 4 bytes in length. Fill bytes are used to pad out data to multiples of 4 bytes.
3. Build the Tachyon Header Structure fields in the HPFS.
4. Write the EOF field to the last 4 bytes of the HPFS.
5. Write the Loop Close (CL) bit of the HPFS, if desired.
6. Write the Trans_ID.
7. Write the Header Address Pointer (Hdr_Addr) to the HPDB.
8. Write the Frame_Length (Frame_Len) to the HPDB. This is the amount of memory allocated for the HPFS in 2. above.
9. Write the Completion Message Disable (C) and Completion Message Interrupt Disable (I) bits of the HPDB, if required

C.7 SEST Entries

Algorithm for Building the Inbound SEST Entry

1. Acquire an empty SEST entry.
2. Build the SDB and write its location to the SEST entry's SCSI Inbound Descriptor Block Address (SDB_Addr) field.
3. If In Order Reassembly mode is being used, write the RO to the SEST entry's Exp_RO field, otherwise clear this field to zero.
4. Write the Buffer Offset (Offset), if required. This field must be cleared to zero if In Order Reassembly is used.
5. Assert the STE Valid (V) bit and the SCSI Direction (D) bit.

Algorithm for Building an Outbound SEST Entry

Building an Outbound SEST Entry is similar to building an ODB. However, some of the fields are in a different location and others have a different function. The main rules for building an Outbound SEST Entry are that the OX_ID in the Tachyon Header Structure must be set to the SEST entry's index in the SEST and the STE Valid (V) bit must be the last bit set. All fields in the SEST entry must be set to zero at initialization.

1. Acquire an empty SEST entry.
2. Build the Tachyon Header Structure and write its location to the Outbound SEST Entry's Header Address Pointer (Hdr_Addr) field.
3. Build the EDB and write its location to the Outbound SEST Entry's EDB_addr.
4. Write the SEQ_ID from the Tachyon Header Structure to the Outbound SEST Entry.
5. Write the D_ID from the Tachyon Header Structure to the Outbound SEST Entry.
6. Write the RX_ID from the Tachyon Header Structure to the SEST Entry.
7. Write the Trans_ID.
8. Write the remainder of the fields in the Outbound SEST Entry.
9. Assert the STE Valid (V) bit.

Outbound SEST entries need their STE Valid (V) bit cleared to prevent the OSM from sending the SCSI data if a FCP_XFER_RDY arrives after an outbound SCSI transaction is aborted.

C.8 SDB

Algorithm for Building the SDB, In Order Reassembly

1. Calculate how many buffers of SCSI Buffer Length are required to hold the incoming data.
2. Allocate memory for the SCSI buffers in multiples of SCSI Buffer Length and write the addresses of the SCSI buffers to the SDB's Buffer Address (Buf_Addr) fields.
3. Set the End of SDB (Es) bit of the End Delimiter field to one. All other bits of the End Delimiter field should be cleared to zero.

Algorithm for Building the SDB, OOO Reassemby

1. Calculate how many buffers of SCSI Buffer Length are required to hold the incoming data.
2. Allocate exactly 256 bytes of memory for the SDB, 256-byte aligned.
3. Allocate memory for the SCSI buffers in multiples of SCSI Buffer Length and write the addresses of the SCSI buffers to the SDB's Buffer Address (Buf_Addr) fields.
4. For any unused buffers, write the address of a safe location in memory to the Buf_Addr fields.
5. The End Delimiter is ignored, so do not set the End of SDB (Es) bit in the SDB.

Appendix D. 8B/10B Encoding/Decoding

Internal 8-bit data bytes are transmitted over the Fibre Channel link as 10-bit encoded characters. Data is encoded primarily to improve transmission quality by maintaining a DC balance, which is a balance of 1s and 0s on the link. Encoding also allows for sufficient transitions in the serial bit stream to make clock recovery possible at the receiver, provides unique control characters, and enables error checking at the character level by detecting invalid encoded transmission characters. The 8B/10B encode/decode method used by Fibre Channel is patented by IBM.

Encoding and decoding are performed by the FC-1 Encode/Decode layer. Encoding works by dividing each internal 8-bit data byte into a 5-bit sub-block comprising the low-order five bits and a 3-bit sub-block comprising the high-order three bits. For example, an 8-bit byte whose bits are labeled in FC-1 notation as A (for the low-order bit) through H

```
7 6 5 4 3 2 1 0
H G F E D C B A
```

is subdivided into two sub-blocks.

```
7 6 5       4 3 2 1 0
H G F       E D C B A
```

The sub-blocks are swapped and their bit order is reversed. Then, by looking up the 8-bit byte in a FC-1 encoding table an additional bit is added to each sub-block, creating the encoded 10-bit transmission character whose bits are labeled a through j. The transmission character consists of a 6-bit sub-block and a 4-bit sub-block.

```
0 1 2 3 4       5 6 7
A B C D E       F G H
a b c d e i     f g h j
```

Bit a is the first bit transmitted onto the link. Bits are transmitted serially.

Decoding is the converse of encoding.

Using 10 bits for transmission characters supplies 1024 unique bit patterns rather than the limited 256 combinations for 8-bit characters. 1024 bit patterns allow two encodings for each internal character plus additional bit patterns for special control characters. 444 of the 1024 possible 10-bit patterns are used by Fibre Channel. The values used provide the best transition density by limiting the run of consecutive 1s and 0s on the link. The remaining 580 combinations are invalid and result in code violations if they occur.

8B/10B encoded characters can contain three types of bit patterns, six bits set to 1 and four bits set to 0, six bits set to 0 and four bits set to 1, or five bits set to 1 and five bits set to 0. The relative number of 1s to 0s is called disparity. Sub-blocks containing more 1s than 0s have positive disparity. Sub-blocks containing more 0s than 1s have negative disparity. Disparity is neutral when the number of 1s and 0s are equal. Disparity is independently determined for each 6-bit and 4-bit sub-block. Disparity at the end of each 10-bit character is the disparity at the end of its 4-bit sub-block. 122 of the 256 byte-wide data characters have both a positive disparity encoding and a negative disparity encoding. 72 characters are neutral with one encoding. The remaining 62 characters are neutral with two encodings. Disparity starts out negative when the link comes up.

PTI 172768

Running disparity is the cumulative disparity of all previous characters. Running disparity alternates after each positive or negative disparity character, but does not change after neutral disparity characters. Running disparity is used to determine the encoding and decoding of characters. To maintain the DC balance, if the running disparity is positive, a negative disparity encoding is used for the next character. If the running disparity is negative, a positive disparity encoding is used for the next character. Running disparity is never neutral. A neutral disparity character does not change the running disparity, for example, if the running disparity is positive, it remains positive at the end of the neutral disparity character.

A disparity error occurs if the running disparity is positive and a positive disparity character is received, or the running disparity is negative and a negative disparity character is received. Disparity errors may not be detected on the character in error. Neutral disparity characters following the error delay its detection.

8B/10B encoding uses a special notation of the form Zxx.y to identify characters. Each of the 256 bit patterns possible in a byte is assigned a symbol. For example, D21.1 is the symbol for a byte containing the hexadecimal value 35, which has the bit pattern 00110101. The symbol's Z digit is a control variable that identifies the character type. The letter D, as the Z digit, designates a "valid data byte." Fibre Channel control characters are special symbols and are designated using a K. The data symbol's xx digits are the decimal value of the byte's 5-bit sub-block, which in the example is decimal 21 for binary 10101. The data symbol's y digit is the decimal value of the byte's 3-bit sub-block, which in the example is decimal 1 for binary 001.

Twelve special control symbols are defined by the 8B/10B encoding scheme, but only one, the K28.5, is used by Fibre Channel. The K28.5 character contains a "comma," which is unique in that its encoding consists of either two 0s followed by five 1s (binary 0011111) or two 1s followed by five 0s (binary 1100000). No other character contains these bit patterns. These special patterns allow the comma to be used for byte synchronization.

10-bit encoded characters are never transmitted individually. They always are part of a 40-bit transmission word. Transmission words are either data words or ordered sets. Data words contain data symbols exclusively and transmit encoded data. Ordered sets are groups of four characters, consisting of a K28.5 as the first symbol followed by three data symbols. The second symbol is usually a D21.4 or D21.5. The third and fourth symbols normally are repeated for error checking. For example, the SOF Connect Class 1 (SOFc1) ordered set consists of the symbols K28.5, D21.5, D23.0, and D23.0. Ordered sets are used for frame delimiters, primitive signals, and primitive sequences.

Frame delimiter ordered sets consist of SOFs and EOFs, which identify frame boundaries. EOF ordered sets have two encodings but always force a negative running disparity. As a result, running disparity between transmitted frames is always negative. SOFs require a negative running disparity.

Primitive signal ordered sets consist of IDLE, Receiver Ready (R_RDY), and the Arbitrated Loop primitive signals Arbitrate (ARB), Open (OPN), Close (CLS), and Mark (MRK). Primitive signals fill the gaps between transmitted frames in order to maintain proper link transmission characteristics during periods of inactivity. Transmitters must send at least six primitive signals between frames. IDLEs may be inserted or deleted singly by switches, repeaters, or test tools for clock elasticity purposes between adjacent nodes. A receiver expects at least two IDLEs between frames. IDLE and R_RDY primitive signals maintain negative disparity on the link.

Primitive sequence ordered sets are used for communication at a level more primitive than frames. They are used for link initialization and error recovery. Primitive sequence ordered sets include Offline (OLS), Link_Reset (LR), Loop Initialization (LIP), Not_Operational (NOS), and Link_Reset_Response (LRR). A minimum of three occurrences of a primitive sequence must be received before it has a meaning.

For more information, refer to "FC-PH Revision 4.3+, Chapter FC- 1 8B/10B Transmission Code, p. 63 - 71."

PTI 172769

Fibre Channel Frame Quick Reference: FC-PH 4.3

HEWLETT PACKARD

	SOF	FRAME HEADER	DATA FIELD: optional headers: Expiration_Security	Network	Association	Device	PAYLOAD	fill	CRC	EOF	"normal" Link Primitives
bytes	4	24	16	16	32	0 | 16 | 32 | 64	≤ 2112	0-3	4	4	IDLE | R_RDY ... ≥ 24tx, 8rx

SOF delimiters
c1 = Class 1 connect request
ix = initiate sequence on Class x
nx = sequence continuation on Class x
f = fabric

EOF delimiters
t = terminate (sequence) *last ACK_x or Class 3 Data frame*
dt = disconnect-terminate
a = abort
n = normal
dti = dt + invalid
ni = n + invalid

N_port or F_port address identifiers
FF FF FF = reserved for broadcast
FF FF FE = Fabric F_Port
FF FF FD = Fabric Controller
FF FF FC = Name Server = Directory Service
FF FF FB = Time Server
FF FF FA = Management Server
FF FF F9 .. FF FF F0 = reserved
Fx xx xx = N_port alias
Ex xx xx = N_port alias
all others assigned by fabric at LOGIn
00 00 00 xx = AL_PA on "closed loop"
00 00 00 00 = unidentified

bits 23-14: Exchange/Sequence Control
bit 23: Exchange context: originator \| responder
bit 22: Sequence context: initiator \| recipient
bit 21: First_Sequence of exchange
bit 20: Last_Sequence of exchange
bit 19: End_Sequence = last Data frame of Sequence
bit 18: E_C = End of Connection Pending (Class 1)
bit 17: C_S = Chained_Sequence Active (Class 1)
bit 16: Sequence initiative: transfer
bit 15: X_ID reassigned
bit 14: Invalidate X_ID
bit 13-12: ACK_Form (FC-PH2)
no assistance provided \| ACK_1 \| ACK_0/ ACK_N
bit 11: Reserved
bit 10: Reserved
bit 9: Sequence retransmission
bit 8: Unidirectional transmission
bits 7-6: Continue Sequence Condition
next sequence to follow: No information \| immediately \| soon \| delayed
bits 5-4: Abort Sequence Condition
policy requested on 1st Data frame of Exchange:
Abt, Discard multi seqs \| Abt, Discard \| seq \| Process - inf. bufs \| Discard multi seqs w/retrans
for ACK frame (from Sequence Recipient):
Sequence: Continue \| Abort, Perform ABTS \| Stop \| Immediate Sequence retransmission
bit 3: Relative Offset present
bit 2: (Reserved for) Exchange reassembly
bits 1-0: Fill Data Bytes
no fill \| last 1 \| last 2 \| last 3 ... bytes of Data Field

Class 4
VC_ID: virtual circuit
Class Dependent Field **Frame Control**

bit#	Word 0 (31–27 / 23–0)	Word 1 (31–23 / 23–0)	Word 2 (31–23 / 23–0)	Word 3 (31–23 / 23–15 / 15–0)	Word 4 (31–15 / 15–0)	Word 5 (31–23 / 23–16 / 15–7 / 7–0)	
	Routing Bits | Info. Category	D_ID	CS_CTL (FC-PH2) | S_ID	TYPE	F_CTL	SEQ_ID | DF_CTL | SEQ_CNT	OX_ID | RX_ID	H | N	a)
	R_CTL						X_ID	Action | Reason | res. | V.U	b)
								Relative Offset (RO)	c)
bytes	1 | 3	1 | 3	1 | 3	1 | 1 | 2	2 | 2	4		

Routing Control

Data (FT-1) frame types (2112byte max. payload):
0000:FC-4 Device_Data frame, and
0100:Video_Data
0 = Uncategorized info.
1 = Solicited Data
2 = Unsolicited Control
3 = Solicited Control
4 = Unsolicited Data
5 = Data Descriptor. Payload: 4byte: Offset, Length, Res., FC-4 Optional information
6 = Unsolicited Command. Payload: 8 byte FC-4 Entity Address, FC-4 Command information
7 = Command Status 4byte: Command status, Res., FC-4 Optional status
> 7 = Unspecified

0010:Extended Link Service, and
0011:FC-4 Link_Data frame
2 = Unsolicited Control
3 = Solicited Control

1000: Basic Link Service: 8byte max. payload
0:NOP = No Operation
1:ABTS = Abort Sequence
2:RMC = Remove Connection
3:reserved
4:BA_ACC = Basic_Accept
5:BA_RJT = Basic_Reject

Link_Control (FT-0) frame types (0byte payload):
1100: Link_Control frame: 0byte payload
Acknowledges:
0: ACK_1
1:ACK_N or ACK_0
Link_Responses:
2:P_RJT = N_Port Reject
3:F_RJT = Fabric Reject
4:P_BSY = N_Port Busy
5:F_BSY(DF) = Fabric Busy to Data frame
6:F_BSY(LC) = Fabric Busy to Link_Control frame
Link_Commands:
7:LCR = Link Credit Reset

Data structure type
FC-4 Device&Link_Data Types
4 = IS8802 - 2 LLC (in order)
5 = IS8802-2 LLC/SNAP
8 = SCSI - FCP
9 = SCSI - GPP
10. to 15. = Reserved - SCSI
16. = Reserved - IPI-3
17. = IPI-3 Master
18. = IPI-3 Slave
19. = IPI-3 Peer
20. = Reserved for IPI-3
21. = CP IPI-3 Master
22. = CP IPI-3 Slave
23. = CP IPI-3 Peer
24. = Reserved - SBCCS
25. = SBCCS (SSA) Channel
26. = SBCCS (SSA) Control Unit
27. to 31. = Reserved - SBCCS
32. = Fibre Channel Services
33. = FC-FG
34. = FC-XS
35. = FC-AL
36. = SNMP
37. to 39. = Reserved - Fabric Services
40. to 47. = Reserved - Futurebus
64. = HIPPI - FP
65. to 67. = Reserved - HIPPI
224. to 255. = Vendor Unique

Fast_Path_Video TYPEs
208 to 255 = Vendor Unique

Basic&Extended Link Service Types:
0 = Basic Link Service
1 = Extended Link Service
208 to 255 = Vendor Unique

F_BSY Link_Control frame TYPE field (CCCC = Link_Control code = Info. Category field)
0001CCCC = Fabric Busy
0011CCCC = N_Port busy

ALL OTHERS reserved

Data_Field Control **Used for:**
ACK_N (N=N), ACK_1 (N=1), ACK_0 (N=0): a)
H = History bit: all previous ACKs transmitted \| not P_BSY and x_RJT: b)
Data Frame and F_BSY(DF): c)
bits 17-16: Device Header length: 0 \| 16 \| 32 \| 64 bytes
bits 19-18: reserved
bit 20: Association_Header: 32 bytes
bit 21: Network_Header: 16 bytes
bit 22: Expiration_Security Header: 16 bytes
bit 23: Reserved for Extended Frame_Header

PTI 172770

Terms and Abbreviations

A/L Pair	Address / Length Pair
ABT	Abort
ABTS	Abort Sequence
ABTX	Abort Exchange
ACK	Acknowledge
AL_PA	Arbitrated Loop Physical Address
AL_PD	Arbitrated Loop Physical Destination Address
AL_PS	Arbitrated Loop Physical Source Address
AL_TIME	Arbitrated Loop Time Out Value
ANSI	American National Standards Institute
ARB	Arbitrate primitive signal
ARB(F0)	ARB primitive with AL_PA of F0 is used for access fairness and during loop initialization
ARB(x)	L_Port with AL_PA of x is arbitrating for control of the loop
AVCS	Address Valid Chip Select
BB	Buffer-to-Buffer
BSY	Busy
C	The C programming language
CLS	Close open port primitive signal
CNTL	Control
COM_DET	Comma Detect
CRC	Cyclic Redundancy Check
D_ID	Destination Identifier
DIO	Direct I/O

PTI 172771

DMA	Direct Memory Access
E_C	End Connection
ED_TOV	Error Detect Time Out Value
EDB	Extended Descriptor Block
EE	End-to-End
EISA	Extended Industry Standard Architecture
EOF	End of Frame
EOFa	End of Frame Abort
EOFdt	End of Frame Disconnect Terminate
EOFdti	End of Frame Disconnect Terminate Invalid
EOFn	End of Frame Normal
EOFni	End of Frame Normal Invalid
EOFt	End of Frame Terminate
EOS	End of Sequence
EWRAP	Electrical Wrap Enable
F_CTL	Frame Control
F_Port	Fabric Port
FC	Fibre Channel
FC-0	Fibre Channel Layer 0 Physical layer
FC-1	Fibre Channel Layer 1 Encode/Decode layer
FC-2	Fibre Channel Layer 2 Framing Protocol layer
FC-3	Fibre Channel Layer 3 Common Services layer
FC-4	Fibre Channel Layer 4 Protocol Mappings layer
FC-AL	Fibre Channel Arbitrated Loop topology

PTI 172772

FC-EP	Fibre Channel Enhanced Physical
FC-PH	Fibre Channel Physical and Signaling Interface
FCP	Fibre Channel Protocol for SCSI
FCSI	Fibre Channel System Initiative
FIFO	First In First Out
FL_Port	Fabric Loop Port
FLOGI	Fabric Login
FM	Frame Manager
GLM	GigaBaud Link Module Family; Gigabit Link Module
HP	Hewlett-Packard Company
HPCQ	High Priority Command Queue
HPDB	High Priority Descriptor Block
HPFS	High Priority Frame Structure
I/O	Input / Output
ID	Identifier
IEEE	Institute of Electrical and Electronics Engineers
IMQ	Inbound Message Queue
INT	Interrupt
IP	Internet Protocol
IPI	Intelligent Peripheral Interface
ISA	Industry Standard Architecture
ISM	Inbound Sequence Manager
ISR	Interrupt Service Routine
JTAG	Joint Test Action Group

KByte	Kilo Bytes; 1028 bytes
L_Port	Loop Port, either a FL_Port or NL_Port
LCKREF	Lock to Reference
LF	Link Failure
LIFA	Loop Initialization Fabric Assigned address - loop initialization sequence
LIHA	Loop Initialization Hard Assigned address - loop initialization sequence
LILP	Loop Initialization Loop Position - loop initialization sequence
LIP	Loop Initialization Primitive sequence
LIPA	Loop Initialization Previously Assigned address - loop initialization sequence
LIRP	Loop Initialization Report Position - loop initialization sequence
LISA	Loop Initialization Soft Assigned address - loop initialization sequence
LISM	Loop Initialization Select Master - loop initialization sequence
LPB	Loop Port Bypass primitive sequence
LPE	Loop Port Enable primitive sequence
LR	Link Reset
LRR	Link Reset Response
MBaud	Mega Baud
MByte	Mega Bytes; 1,048,576 bytes
MFS	Multiframe Sequence
MFSBQ	Multiframe Sequence Buffer Queue
MIB	Management Information Base
MUX	Multiplexer
N_Port	Node Port
NFS	Network File System

PTI 172774

NL_Port	Node Loop Port
NLOGI	Node Login
NOP	No Operation
NOS	Not Operational State
NWH	Network Header
OCQ	Outbound Command Queue
ODB	Outbound Descriptor Block
OLS	Offline State
OOO	Out of Order
OPN	Open primitive signal
OPNfr	Open all ports, broadcast replicate mode primitive, with characters 3 (f) and 4 (r) of the OS = 0xFF
OPNr	Open Replicate primitive, either OPNfr or OPNyr
OPNy	Open primitive, either OPNyx or OPNyy, where y is the AL_PD
OPNyr	Open Port(y), selective replicate mode primitive, with character 3 (y) of the OS = AL_PD and character 4 (r) = 0xFF
OPNyx	Open Port(y), full-duplex from Port(x) primitive, where y is the AL_PD and x is the AL_PS
OPNyy	Open Port(y), half-duplex primitive, with both characters 3 and 4 of the OS = AL_PD
OS	Ordered Set
OSM	Outbound Sequence Manager
OX_ID	Originator Exchange Identifier
P_BSY	N_Port Busy
P_RJT	N_Port Reject
PCI	Peripheral Component Interconnect
PLL	Phase Locked Loop
PLM	Physical Link Module

PTI 172775

PLOGI	Port Login
R_CTL	Routing Control
R_RDY	Receiver Ready
RBC	Receive Byte Clock
RDY	Ready
RJT	Reject
RO	Relative Offset
RRO	Random Relative Offset
RT_TOV	Receiver-Transmitter Time Out Value
RX	Receive
RX_ID	Responder Exchange Identifier
S_ID	Source Identifier
SCLK	System Clock
SCSI	Small Computer System Interface, either SCSI-2 or SCSI-3
SCSI-2	The SCSI architecture specified by ANS X3.131-1994
SCSI-3	The SCSI architecture specified by ANS X3.270-199X
SDB	SCSI Descriptor Block
SEQ_CNT	Sequence Count
SEQ_ID	Sequence Identifier
SEST	SCSI Exchange State Table
SFS	Single Frame Sequence
SFSBQ	Single Frame Sequence Buffer Queue
SNAP	Standard Network Access Protocol
SOF	Start of Frame primitve

PTI 172776

SOFc1	Start of Frame Connect Class 1 primitive
SOFf	Start of Frame Fabric primitive
SOFi1	Start of Frame Initiate Class 1 primitive
SOFi2	Start of Frame Initiate Class 2 primitive
SOFi3	Start of Frame Initiate Class 3 primitive
SOFn1	Start of Frame Normal Class 1 primitive
SOFn2	Start of Frame Normal Class 2 primitive
SOFn3	Start of Frame Normal Class 3 primitive
SRO	SCSI Relative Offset
STE	State Table Entry
TAD	TSI Address Data
TAP	Test Access Port
TBC	Transmit Byte Clock
TBG	Tachyon Bus Grant
TBR	Tachyon Bus Request
TCK	Test Clock
TCP	Transmission Control Protocol
TDI	Test Data In
TDO	Test Data Out
THS	Tachyon Header Structure
TMS	Test Mode Select
TRST	Tachyon Reset
TSI	Tachyon System Interface
TX	Transmit

UDP	User Datagram Protocol
UF	Unfair Access
ULP	Upper Level Protocol
WWN	World Wide Name
X_ID	Exchange Identifier
XMIT	Transmit
XFR_RDY	Transfer Ready Sequence

Glossary

8B/10B Encoding

A method used to convert 8-bit bytes into encoded 10-bit transmission characters. The encoding scheme maintains an equal number of 1s and 0s on the link, provides control characers and error detection, and a special notation to identify characters. To encode a character, the byte is divided into a 5-bit and a 3-bit sub-block. The 5-bit sub-block is encoded into six bits and the 3-bit sub-block is encoded into four bits. Most characters have a positive disparity encode and a negative disparity encode. The encode selected is determined by the running disparity at the end of the previous character.

10B/20B Demux (Demultiplexer)

Responsible for receiving incoming encoded data, either 10-bits wide or 20-bits wide, from the Physical Link Module (PLM) and packing it into 20-bits for decoding. A 20-bit data width is used with a 100 Mbytes/sec link speed. All other link speeds transmit 10-bit wide data. The data width is specified in the Parallel ID field of the PLM Interface.

16B/20B Encoder

An encoder that converts outbound 16-bit wide data from the elastic store into two 8-bit pieces, each of which is encoded into a 10-bit transmission character by using the 8B/10B encoding algorithm.

20B/10B Mux (Multiplexer)

Responsible for selecting the proper width, either 10 bits or 20 bits, of the data path for the specific type of Physical Link Module (PLM) being used. The data width is specified in the Parallel ID field of the PLM Interface. A 20-bit data width is used with a 100 Mbytes/sec link speed. All other link speeds transmit 10-bit wide data.

20B/16B Decoder

Responsible for converting 20-bit wide data received from the 10B/20B Demultiplexor from the transmission codes used on the link into 8-bit data bytes.

ACK FIFO

A FIFO used to hold ACK frames until thay can be sent out. The ACK FIFO holds up to eight ACK frames.

Abort Exchange (ABTX)

A Fibre Channel extended link service that is sent in a separate exchange and is used to abort a specified exchange. The payload identifies the target exchange, which is identified by OX_ID, RX_ID, and S_ID. An ABTX can be requested only by the originator or the responder of the exchange.

Abort Sequence (ABTS)

A Fibre Channel basic link service that is sent in a single frame and is used to abort a specified sequence. The sequence is identified by OX_ID, RX_ID, and SEQ_ID.

Acknowledgement (ACK)

A link control frame used for flow control to indicate a frame has been received. Fibre Channel defines three acknowledgment models, ACK_1 in which one ACK frame is sent for each frame received, ACK_N in which one ACK is sent for a group of N frames, and ACK_0 in which one ACK is sent for all frames within a sequence. ACK_1 is the required default.

Arbitrated Loop

One of three existing Fibre Channel topologies, in which from two to 126 devices are interconnected serially in a single loop circuit without hubs and switches. Communication is managed using an arbitration process in which the lowest port address has the highest priority for communicating. The arbitrated loop topology supports all classes of service and guarantees in order delivery of frames when the source and destination are on the same loop.

Arbitrated Loop Physical Address (AL_PA)

A unique 1-byte address used to identify each port on an arbitrated loop. An AL_PA for each L_Port is assigned during loop initialization.

Arbitrated Loop Time Out Value (AL_TIME)

The number of milliseconds the port state machine waits for responses during arbitrated loop initialization. The default value is 15 milliseconds.

Bad CRC Count

One of four error counters in the Frame Manager Link Error Status #2 Register. It indicates the number of frames received with a bad CRC.

Bad SCSI Frame

An unrecognized, unknown, or invalid SCSI frame, such as an inbound SCSI frame having an invalid OX_ID, or an out of order SCSI frame received during an in order reassembly.

Bad SCSI Frame Completion Message

A 32-byte Inbound Message Queue (IMQ) entry indicating an inbound SCSI frame containing an invalid OX_ID.

Bad Tx Character Count

One of four error counters in the Frame Manager Link Error Status #1 Register. It indicates the number of times the 8B/10B decode failed to detect a valid 10-bit code.

Baud

The encoded bit rate per second.

PTI 172780

Buffer-to-Buffer Credit (BB_Credit)

Buffer-to-Buffer credit manages the flow of frames between two ports at the ends of a single link. BB_Credit specifies the number of frames which can be received by the next node connected to Tachyon (either an F_Port or an N_Port) before that receiver's buffers fill up. The BB_Credit of the receiver is discovered during Login, and is supplied to Tachyon in the Frame Manager Configuration Register. Each time a frame is sent, BB_Credit is decremented, and each time an R_RDY link primitive is received, BB_Credit is incremented.

Buffer Warning Completion Message

A 32-byte completion message sent to warn the host when the MFSBQ or SFSBQ are one entry from being empty or the IMQ is two entries from being full.

Busy (BSY)

A link control frame used for flow control to indicate that a frame could not be delivered. BSYs are returned in response to a transmitted frame. BSYs can indicate that either the fabric is busy (F_BSY) or the N_Port is busy (P_BSY).

Channel

A direct or cross-point switched connection between communicating devices, usually a processor and an I/O peripheral device, which transports data at the highest speeds and with the least delay possible.

Circular Queue

A contiguous area in host memory that can be thought of as an array of queue elements having no "first" or "last" element. Queues are used by Tachyon to pass messages and memory descriptors between itself and its host system. The queues contain entries that describe things that have been done, things to do, and lists of resources Tachyon can use. Each queue is defined by two parameters that must be provided by the host, the base address of the queue list, and the length of the queue, which is the number (zero based) of queue entries. All elements in all of the queues are 32 bytes in length.

Circular Queue Index (Q_Index)

A 32-bit register value used to reference the buffer queue entry from where the last address was taken.

Class 1 Service

A service which establishes a dedicated hard or circuit-switched connection between communicating N_Ports and confirms delivery of frames, or notification of non-delivery. The connection may span multiple exchanges. Class 1 frames are always delivered in order.

Class 2 Service

A service which provides a connectionless, frame switched link, with acknowldgement confirmed or notification of non-delivery. Frames are individually routed and may be delivered out of order if multiple paths exist through the fabric. Buffer-to-buffer flow control is used on every frame.

PTI 172781

Class 3 Service

A service which provides a connectionless, frame switched link, with delivery not guaranteed and with no acknowldgement of delivery or non-delivery. Frames are individually routed and may be delivered out of order if multiple paths exist through the fabric.

Comma

A 7-bit sequence, either 0011111 or 1100000, in the 8B/10B encoding scheme, which is used for byte synchronization.

Consumer Index

A 32-bit register used by the consumer of circular queue elements to indicate the index of the next queue element to be processed. The host is the consumer of the Inbound Message Queue. Tachyon is the consumer of all other queues. The consumer index exists in producer space, which eliminates the need for any read operation of this index across the backplane interface and would reduce host CPU performance.

Credit

Credit is a flow control mechanism that throttles link traffic by limiting the number of frames each N_Port or F_Port can send. A credit limit, which represents the maximum number of outstanding frames that can be transmitted by an N_Port or F_Port without causing a buffer overrun condition at the receiver, is established at login. Available credit is decremented for each frame that is sent and incremented for each R_RDY response received. Credit cannot be incremented above the credit limit nor can it be decremented below zero. Two types of credit can be used, buffer-to-buffer credit (BB_Credit) and end-to-end credit (EE_Credit).

Cyclic Redundancy Check (CRC)

A 4-byte error correcting code used to verify the accuracy of a frame's contents.

Data Field Control (DF_CTL)

A field in the Fibre Channel frame header that indicates whether any optional headers are present. FCP does not require any optional headers.

Decoding

Validity checking of received transmission characters and generation of valid data bytes and special codes from those characers.

Deferred P_BSY Mode

A mode of operation in which Tachyon does not automatically send a P_BSY to the remote node. When Tachyon receives a new MFS while in this mode, Tachyon sends the MFS to the host via the SFSBQ and generates an inbound_busied_frame completion message. The host stores this frame in a Deferred P_BSY queue. The host sets Deferred P_BSY mode by setting the Disable AUTO P_BSY bit in the Tachyon Configuration register.

PTI 172782

Destination Identifier (D_ID)

The address identifier used to indicate the targeted destination of the transmitted frame.

Disparity

Disparity is the difference between the number of 1s and 0s in an encoded transmission character. Disparity is positive if the encoded character contains more 1s than 0s. Disparity is negative if the character contains more 0s than 1s. Neutral disparity means the number of 0s and 1s is equal.

Disparity Error

An error detected if the running disparity is positive and a positive disparity encoded character is received, or the running disparity is negative and a negative disparity encoded character is received. Disparity errors cannot indicate which character is in error as neutral disparity can delay the detection of the error.

Elastic Store

Responsible for retiming received data from the link clock (RX_CLK) to the internal clock (SCLK). The elastic store holds 16-bit wide data.

Encapsulation

The method for mapping ULP objects into Fibre Channel frames.

Encoding

Generation of transmission characters from valid data bytes and special codes.

End of Frame (EOF)

A 4-byte value delimiting the end of a frame.

End-to-End Credit (EE_Credit)

Specifies the number of frames that Tachyon can send to a remote D_ID without receiving an ACK. EE_Credit is decremented each time a frame is sent and incremented each time an ACK link control frame is received. This does not apply to Class 3 service.

Error Detect Time Out Value (ED_TOV)

The Fibre Channel Error Detect Time Out Value in milliseconds. The default is 500 milliseconds.

Exchange

The basic mechanism used for managing an operation. An exchange identifies information transfers consisting of one or more related non-concurrent sequences which may flow in the same or opposite directions, but always in half duplex mode. An exchange may span multiple Class 1 dedicated connections. An exchange is identified by an OX_ID and a RX_ID.

Extended Descriptor Block (EDB)

A data structure that defines the data payload of an outbound Fibre Channel sequence. Each entry in the EDB consists of an 8-byte Address/Length pair that defines a data buffer in host memory containing data to be transmitted. The address to an EDB is contained in the sequence's Outbound Descriptor Block and must be 32-byte aligned.

Fabric

A generic term used to describe a crosspoint switched topology, which is one of the three existing FC topologies. The fabric interconnects the various N_Ports attached to it. A fabric consists of one or more fabric elements, which are the switches responsible for frame routing. Frames are routed by using the D_ID field in the FC-2 frame header. The fabric structure is transparent to the N_Ports connected to it. It also relieves the ports from the responsibility for station management. A maximum of 224 devices can be interconnected in a fabric topology.

Fabric Busy (F_BSY)

A link control frame used for flow control to indicate that a frame could not be delivered by the fabric because the fabric was busy or the destination N_Port was busy.

Fabric Port (F_Port)

The access point of the fabric for physically connecting an N_Port.

FC-0 Physical Layer

The lowest of three FC-PH layers, FC-0 defines the characteristics of the physical interface, media, and data rates, including connectors, cables, speeds, transmitters and receivers. The Physical Link Module (PLM) performs the FC-0 functions.

FC-1 Encode/Decode Layer

The middle of three FC-PH layers, FC-1 defines the encoding, decoding, byte synchronization, and error control required to transmit and receive data.

FC-2 Framing Protocol Layer

The highest of three FC-PH layers, FC-2 defines the signaling protocol that specifies the rules and provides the mechanisms, such as the frame structure and byte sequences, needed to transfer blocks of data end-to-end. FC-2 defines a suite of functions and facilities available for use by an FC-4.

FC-3 Common Services Layer

FC-3 provides a set of services which are common across multiple N_Ports of a FC node. No FC-3 functions are formally defined.

PTI 172784

FC-4 Protocol Mappings Layer

FC-4 provides a mapping function from ULP objects to Fibre Channel information units, which identify the protocol's constructs to the lower layers. Example SCSI constructs include Transfer Ready, Status, and Command. Each ULP has its own FC-4 mapping.

Fibre

A general term used to cover all transmission media types specified in FC-PH, such as optical fiber, twisted pair, and coaxial cable.

Fibre Channel (FC)

Logically, Fibre Channel is a bidirectional, full-duplex, point-to-point, serial data channel structured for high performance capability. Physically, the Fibre Channel is an interconnection of multiple communication ports, called N_Ports, interconnected by a switching network, called a fabric, a point-to-point link, or an arbitrated loop. Fibre Channel is a generalized transport mechanism that has no protocol of its own or native I/O command set, but can transport any existing Upper Level Protocols (ULPs) such as SCSI and IPI. Fibre Channel operates at speeds of 100 Mbytes/sec (full speed), 50 Mbytes/sec (half speed), 25 Mbytes/sec (quarter speed), or 12.5 Mbytes/sec (eighth speed), over distances of up to 100m over copper media or up to 10km over optical links. Fibre Channel can interconnect two devices in a point-to-point topology, from two to 126 devices in an arbitrated loop, and up to 224 devices in a fabric switched topology.

Fibre Channel Aribrated Loop (FC-AL)

The Fibre Channel standard, X3.272-199x, X3T11/Project 960D/Rev 4.5, that defines the arbitrated loop topology.

Fibre Channel Physical and Signaling Interface (FC-PH)

The architecture specified by the Fibre Channel standard, X3.230-199x, X3T11/Project 755D/Rev 4.3. FC-PH performs the functions required to transfer data from one N_Port to another. It defines three classes of service and is subdivided into three levels, FC-0, FC-1, and FC-2.

Fibre Channel Protocol for SCSI (FCP)

FCP defines a Fibre Channel mapping layer (FC-4) that uses FC-PH services to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target. Using FCP enables transmission and receipt of SCSI commands, data and status, across the Fibre Channel using the standard Fibre Channel frame and sequence formats. FCP operates with FC Classes of Service 1, 2, and 3 across fabric and arbitrated loop topologies.

Fibre Channel System Initiative (FCSI)

FCSI is a profile jointly developed by Hewlett-Packard, IBM and Sun for interoperability of their Fibre Channel products.

FIFO (First In First Out)

A type of queue in which data is written to the tail (rear) end and removed from the head (front) end, so that the data element written first is removed first.

Flow Control

A mechanism that uses credits to control the flow of frames between an originator and responder. The responder extends credit to the originator, giving advance permission to send a specified number of frames.

PTI 172785

Frame

The smallest, indivisible unit of information transfer used by FC-2. Frames are used for transferring data associated with a sequence. Frame size is dependent on the hardware implementation and independent of the ULP or the application software. Frames begin with a 4-byte SOF, end with a 4-byte EOF, include a 24-byte frame header and a 4-byte CRC, and can carry a variable data payload from 0 to 2112 bytes, the first 64 of which can be used for optional headers. Tachyon supports a maximum data payload of 2048 bytes, including optional headers. The maximum frame size including the SOF and EOF is 2148 bytes. Each frame within a sequence is identified with the SEQ_ID and a unique SEQ_CNT.

Frame Control (F_CTL)

A field in the FC-2 frame header containing control information about the frame content.

Frame Header

A 24-byte field following the SOF delimiter in the Fibre Channel frame, which contains identification information including the D_ID, S_ID, SEQ_ID, OX_ID, RX_ID and other operating parameters.

Frame Manager

The Tachyon front end responsible for the FC-1 functions of transmitting and receiving Fibre Channel frames. It is capable of generating interrupts to the host when certain link configuration changes occur that the host must respond to. The interrupt process occurs as part of the Loop initialization process and any time the link has been broken. The Frame Manager responds to a reset condition by initializing all of its registers to their default values.

Frame Manager Configuration Register

A 32-bit read/write register that contains Frame Manager configuration information.

Frame Manager Control Register

A 32-bit register written to by the host to release any pending state changes that need to be completed when the Frame Manager generates an interrupt.

Frame Manager Interrupt Completion Message

A 32-byte message indicating that a Frame Manager interrupt condition has been detected.

Frame Manager Link Error Status Counters #1 Register

A 32-bit read-only register containing four error counters that are incremented each time the associated error occurs. The counters are Loss of Signal Count, Bad Tx Character Count, Loss of Sync Count, and Link Fail Count. The counter values remain at 0xff when the maximum count is reached.

PTI 172786

Frame Manager Link Error Status Counters #2 Register

A 32-bit read-only register containing four error counters that are incremented each time the associated error occurs.. The counters are Received EOFa, Generated EOFa, Bad CRC Count, and Protocol Error Count. The counter values remain at 0xff when the maximum count is reached

Frame Manager Primitive Register

A 32-bit register, which is write only by the host and read only by Tachyon, containing the lower three bytes of the primitive that Tachyon should transmit in the host control state. The host must always set the upper byte to an encoded K28.5 character.

Frame Manager Received AL_PA Register

A 32-bit read-only register containing the 8-bit AL_PA which was received on the last LIPf or as a bad AL_PA.

Frame Manager RT_TOV/AL_TIME & ED_TOV Register

A 32-bit write-only register containing the Receiver Transmitter Time Out Value (RT_TOV), Arbitrated Loop Time Out (AL_TIME), and the Error Detect Time Out Value (ED_TOV).

Frame Manager Status Register

A 32-bit register read by the host in response to receiving a completetion interrupt from the Frame Manager to determine the current status and state of the Frame Manager.

Frame Manager World Wide Name (Hi/Lo) Register

Two 32-bit write-only registers containing the unique 8-byte World Wide Name (WWN) to be used during initialization. The most significant four bytes of the WWN are in the WWN Hi register while the least significant four bytes are in the WWN Lo register.

Full Duplex

A full duplex channel can transmit and receive data simultaneously.

Generated EOFa

One of four error counters in the Frame Manager Link Error Status #2 Register. It indicates the number of frames received with problems that caused the Frame Manager to attach an EOFa delimiter.

Half Duplex

A half duplex channel can transmit and receive, but not simultaneously. It is similar to traffic flow on a one-lane bridge.

PTI 172787

Hardware Assists

Much of the FCP protocol is managed by Tachyon in hardware, resulting in significant performance advantages over a software implementation. Hardware assists, in which Tachyon's SCSI Exchange Manager shares transaction management responsibilities with the host through use of the SEST, enable the host to reassemble many concurrent inbound I/O sequences.

High Priority Command Queue (HPCQ)

A host based data structure used by the host to instruct Tachyon to transmit single frames using high priority, preferential transmission, bypassing Outbound Command Queue traffic. This functionality is used to allow special Fibre Channel error recovery frames to be sent even when the OCQ channel is blocked due to an error It is also used when the host manually ACKs a frame. The host is the producer and Tachyon is the consumer of this queue.

High Priority Command Queue (HPCQ) Base Register

A 32-bit register indicating the physical address in host memory where the start of the HPCQ is located. This register must not be written after initialization.

High Priority Command Queue (HPCQ) Consumer Index Address Register

A 32-bit register indicating the host address where Tachyon should maintain its HPCQ consumer index. This register is maintained in host memory to allow the host fast access to the index.

High Priority Command Queue (HPCQ) Length Register

A 32-bit register that defines the length of the HPCQ by indicating the number (zero based) of 32-byte entries in the HPCQ. This register must not be written after initialization.

High Priority Command Queue (HPCQ) Producer Index Register

A 32-bit register used by the host to indicate to Tachyon that new commands to process are in the HPCQ. After the host fills in a queue entry it writes the index of the next empty queue entry to this register. When Tachyon determines that the producer index differs from its internally maintained consumer index, it processes the posted commands.

High Priority Descriptor Block (HPDB)

One element of the High Priority Command Queue. Since messages going out the High Priority channel are treated differently from normal outbound messages, the amount of information included in the ODB for a High Priority message is a subset of a normal ODB. The structure of the OBD for either type of message is the same, but the host does not need to fill in the OBD fields unused by the High Priority channel.

Idle (IDLE)

One of two FC-PH primitive signals, IDLEs are inserted and deleted between frames to maintain the link's negative running disparity and to regulate clock differences between nodes. A receiver expects a minimum of two IDLEs between frames.

PTI 172788

Information Unit (IU)

An organized collection of data specified by FC-4 to be transferred as a single sequence by FC-2. Each ULP mapped to Fibre Channel identifies constructs necessary to that protocol, for example, SCSI commands and control functions, or IP and ARP packets. These ULP constructs are mapped by FC-4 to IUs. IUs are independent of the transport mechanism.

Inbound Block Mover

A functional block within Tachyon responsible for DMAing inbound data into buffers specified by the MFSBQ, SFSBQ, or the SCSI Buffer Manager.

Inbound Buffer Channel MFS

A functional block within Tachyon responsible for managing the Multiframe Sequence Buffer Queue (MFSBQ). This block supplies addresses of empty MFS buffers to the Inbound Data Manager and generates a low buffer warning when the supply of MFS buffers runs low.

Inbound Buffer Channel SFS

A functional block within Tachyon responsible for managing the Single Frame Sequence Buffer Queue (SFSBQ). This block supplies addresses of empty SFS buffers to the Inbound Data Manager and generates a low buffer warning when the supply of SFS buffers runs low.

Inbound C1 Timeout Completion Message

A 32-byte message indicating the Inbound Class 1 inactivity timer has expired. A warning is given, but no other action is taken.

Inbound Data FIFO

A FIFO used to buffer incoming frames while their CRC is being verified. It is also used as high availability, temporary storage to facilitate the Fibre Channel flow control mechanisms. This FIFO is sized to hold a maximum of four 2KByte frames (including headers), though different classes of service may affect the BB_Credit that is used.

Inbound Data Manager

Maintains the host buffer structures for sequence reassembly of incoming data frames.

Inbound Message Channel

A functional block within Tachyon responsible for maintaining the host based IMQ. This includes supplying the Inbound Data Manager with the address of the next available entry in the IMQ and generating a warning to the host, via a completion message, when the number of available entries in the IMQ is down to two.

Inbound Message Queue (IMQ)

Used by Tachyon for communicating with the host. Information about new inbound sequences, inbound command completions, and various status messages are passed through the IMQ. Tachyon is the producer and the host is the consumer of this queue. The consumer index is used by the host to return ownership of processed completion entries to Tachyon. An IMQ entry is 32 bytes long.

Inbound Message Queue (IMQ) Base Register

A 32-bit register containing the base address in host memory of the IMQ. This register must not be written after initialization.

Inbound Message Queue (IMQ) Consumer Index Register

A 32-bit register written by the host to indicate which completion entries it has processed and are available to Tachyon for posting new completions. This value, along with the host's copy of the IMQ consumer index, gives the host an indication of how many completion messages have been posted since the last interrupt.

Inbound Message Queue (IMQ) Length Register

A 32-bit register that defines the length of the IMQ by indicating the number (zero based) of 32-byte elements in the IMQ. This register must not be written after initialization.

Inbound Message Queue (IMQ) Producer Index Address Register

A 32-bit register used to indicate the host address of Tachyon's IMQ producer index. This register should not be written after initialization and should be the last IMQ register initialized.

Inbound Multiframe Sequence (MFS) Completion Message

A 32-byte data structure indicating that a MFS has been received and processed by Tachyon.

Inbound Out of Order (OOO) Completion Message

A 32-byte data structure indicating an OOO sequence or frame has been received by the chip and processed.

Inbound SCSI Command Completion Message

When Tachyon receives an unsolicited command while SFS and SCSI assists are enabled, Tachyon sends this 32-byte completion data structure with an interrupt to the host.

Inbound SCSI Data Completion Message

When Tachyon is operating as a SCSI target and all frames for the SCSI sequence have been received, Tachyon sends this 32-byte completion message to the host.

Inbound SCSI Status Completion Message

A 32-byte data structure indicating a SCSI inbound status frame has been received and processed.

Inbound Sequence Manager (ISM)

Responsible for receiving all inbound frames, reassembling multiframe sequences, and establishing sequence completions. Data frames are received by the ISM in either single frame sequence, in order multiframe sequence, or out of order multiframe sequence. For each data frame received the ISM generates and sends an ACK or BSY response to the Outbound Frame FIFO. For each link control frame received the ISM notifies the OSM of the type of response.

Inbound Single Frame Completion Message

A 32-byte data structure written into the IMQ by Tachyon when a single frame that generates a host interrupt has been received. It contains an Interrupt Type, which defines the interrupt as an inbound single frame sequence completion, a Q_Index and Offset, which are used to place the single frame in host memory, and the total length of the data transferred into host memory.

Inbound Unknown Frame Completion Message

A 32-byte data structure indicating that an invalid frame, an unknown frame, or a reject response has been received. No response frame is sent.

Initialization

The logical procedure used by each L_Port on an arbitrated loop to determine its environment and acquire an AL_PA prior to transmitting or receiving data.

Interrupt Service Routine (ISR)

When an interrupt occurs, current processing is suspended and a branch is taken to a higher priority routine (an ISR) that performs the specific actions appropriate for the interrupt. At the conclusion of the ISR, normal processing resumes at the point at which the interrupt occured.

Link

A connection between two nodes, each having at least one N_Port, interconnected by a pair of optical or copper links, one inbound and one outbound.

Link Encapsulation

A method for mapping the IEEE 802.2 Link Level Control (LLC) protocol onto Fibre Channel. Link encapsulation is covered under Fibre Channel Link Encapsulation (FC-LE) X3T11/Project 955, under development.

Link Services

Fibre Channel's built in protocol used to manage link operations, not to transport ULP information. Examples of link services are the Abort Sequence (ABTS) and the Abort Exchange (ABTX).

Link Fail Count

One of four error counters in the Frame Manager Link Error Status #1 Register. It indicates the number of times the Frame Manager detected a NOS or other failure of the initialization protocol which caused a transition into the Link Failure state.

Login

The initial procedure all communicating nodes go through to establish service parameters and a common operating environment.

Loop Port (L_Port)

An N_Port or F_Port that supports arbitrated loop functions associated with arbitrated loop topology.

Loopback

A mode of FC-1 operation in which information passed to the FC-1 transmitter is shunted directly to the FC-1 receiver, overriding any signal detected by the receiver on its attached fibre.

Loss of Signal Count

One of four error counters in the Frame Manager Link Error Status #1 Register. It indicates how many times the Frame Manager detected a low to high transition on the lnk_unuse signal.

Loss of Sync Count

One of four error counters in the Frame Manager Link Error Status #1 Register. It indicates the number of times that loss of sync is greater than RT_TOV.

Multiframe Sequence (MFS)

If a sequence is larger than the frame size, the sending hardware disassembles the sequence into one or more frames and the receiving hardware reassembles the frames back into the sequence. This type of sequence is a multiframe sequence. Each frame in the multiframe sequence has the same SEQ_ID, but a unique SEQ_CNT, which identifies the frame's position within the sequence.

Multiframe Sequence Buffer Length Register

A 32-bit register that defines the length of the multiframe sequence receive buffers. The contents of this register should not be modified after initialization.

PTI 172792

Multiframe Sequence Buffer Queue (MFSBQ)

Used by the host to tell Tachyon the location in host memory of buffers that are used to receive and reassemble incoming data that has been split into an arbitrarily large number of frames. Each 32-byte queue entry contains eight pointers to buffers into which the data is placed. The size of these buffers is programmable and determined by the mfs_buff_len register. All buffers must be the same size, and must be a power of two bytes in length. The host is the producer and Tachyon is the consumer of this queue.

Multiframe Sequence Buffer Queue (MFSBQ) Base Register

A 32-bit register indicating the base address in host memory of Tachyon's MFSBQ. This register must not be written after initialization.

Multiframe Sequence Buffer Queue (MFSBQ) Consumer Index Register

A 32-bit register indicating the MFSBQ entry that Tachyon is currently using to process a multiframe sequence.

Multiframe Sequence Buffer Queue (MFSBQ) Length Register

A 32-bit register that defines the length of the MFSBQ by indicating the number (zero based) of 32-byte elements in the MFSBQ. This register must not be written after initialization.

Multiframe Sequence Buffer Queue (MFSBQ) Producer Index Register

A 32-bit register used by the host to pass buffers to Tachyon for use in reassembling incoming multiframe sequences in host memory.

Node

A node is an entity, system, or device that processes the ULP, FC-4, FC-3, and some of the FC-2 functions, and contains one or more N_Ports.

Node Port (N_Port)

A hardware entity within a node that supports the FC-PH. An N_Port processes FC-0, FC-1 and some FC-2 functions. It may act as an originator, a responder, or both.

Node Port Busy (P_BSY)

A link control frame used for flow control to indicate that a frame could not be delivered because the destination N_Port was busy.

Ordered Set

An ordered set consists of four 10-bit characters, consisting of a combination of data and special characters, which are used to provide very low-level link functions, such as frame demarcation and signaling between two ends of a link.

PTI 172793

Ordered Set/CRC Generator

Encapsulates data to be transmitted into FC-2 frames, and calculates and writes a 32-bit CRC into the frame.

Ordered Set Processor/CRC Checker

Responsible for detecting incoming frame boundaries, verifying the CRC, and passing the data to the inbound FIFO.

Originator

The logical function associated with an N_Port responsible for originating an exchange. A Fibre channel originator is comparable to a SCSI initiator.

Originator Exchange Identifier (OX_ID)

A 16-bit identifying value assigned by the originator of an exchange. The high-order bit (bit 15) is set if the OX_ID is for a non-SCSI assisted transaction. Bit 14 indicates the intended direction of the SEST entry for the OX_ID. This bit is set for an inbound SCSI sequence, and cleared for an outbound sequence. The remaining 14 bits are the value of the OX_ID for the exchange. For a SCSI exchange, the value must be the index of the SEST entry.

Out of Order (OOO)

Multiframe sequences that are received in nonsequential order based on the SEQ_CNT value.

Outbound ACK FIFO

A FIFO that is used to hold acknowledgement (ACK) frames until they can be sent out. The Outbound ACK FIFO holds up to eight ACK frames.

Outbound Block Mover

A functional block within Tachyon responsible for DMAing outbound data from host memory to Tachyon's Outbound Sequence Manager.

Outbound Command Queue (OCQ)

A host based data structure used by the host to issue commands that instruct Tachyon to transmit sequences with normal priority. The host is the producer and Tachyon is the consumer of this queue. The OCQ must be aligned on a sizeof(queue) boundary. OCQ entries are Outbound Descriptor Blocks (ODBs).

Outbound Command Queue (OCQ) Base Register

A 32-bit register containing the OCQ base address in host memory. This register must not be written after initialization.

Outbound Command Queue (OCQ) Consumer Index Address Register

A 32-bit register containing the address in host memory where Tachyon maintains its consumer index for the OCQ. This register is maintained in host memory to allow the host fast access to the index.

Outbound Command Queue (OCQ) Length Register

A 32-bit register defining the length of the OCQ, which is the number (zero based) of 32-byte entries in the OCQ. This register must not be written after initialization.

Outbound Command Queue (OCQ) Producer Index Register

A 32-bit register used by the host to indicate to Tachyon that new commands to process are in the OCQ. After the host fills in a queue entry it writes the index of the next empty entry in the queue to this register. When Tachyon determines that the producer index differs from its internally maintained consumer index, it processes the posted commands.

Outbound Completion Message

A 32-byte data structure indicating an outbound sequence from the OCQ has been processed or interrupted by the occurrence of an error condition.

Outbound Descriptor Block (ODB)

One 32-byte element of the OCQ. ODBs describe and define the fields in the Tachyon header structure of an outbound Fibre Channel sequence, including the SEQ_ID, D_ID, RX_ID, checksumming information, and a pointer to the EDB containing addresses of the data buffers, etc.

Outbound Frame FIFO

A FIFO used to buffer data before transmission to prevent under-run. This FIFO is sized to hold one frame. As the current frame is sent out, the Outbound Frame FIFO is simultaneously filled with the next frame, thus maximizing outbound performance.

Outbound High Priority Completion Message

A 32-byte data structure indicating an outbound sequence from the Outbound HPCQ has been processed.

Outbound Message Channel

The two functional blocks within Tachyon responsible for managing the host based OCQ and HPCQ.

Outbound Sequence Manager (OSM)

The functional block that manages the Fibre Channel protocol that sends an entire sequence of outbound data, including processing all the ACK frames and error handling. The OSM programs the Outbound DMA channel to retrieve a data sequence from host memory and segment it into individual frames for transmission. The OSM fairly arbitrates between processing ODBs from the OCQ and SCSI hardware assists.

Physical Link Module (PLM)

A physical component that manages the functions of the FC-0 layer. Also referred to as a GLM (Gigabit Link Module). Tachyon supports GLMs that conform to the FCSI-301-Revision 1.0 GLM Family.

Point-to-Point

One of three existing Fibre Channel topologies, in which two ports are directly connected by a link with no fabric, loop, or switching elements present.

Port

The hardware entity within a node that performs data communications over the Fibre Channel link.

Primitive Signal

Special encoded characters that fill the space between transmitted frames. Two primitive signals are defined for FC-PH, IDLE and Receiver Ready (R_RDY). Additional primitives are defined for FC-AL. Transmitters must send a minimum of six primitive signals between frames.

Producer Index

A 32-bit register used by the producer of circular queue elements to indicate the index of the next memory location in which a queue element will be placed. Tachyon is the producer of the Inbound Message Queue. The host is the producer of all other queues. The producer index exists in consumer space in order to eliminate the need for any read operation across the backplane interface, which would reduce performance of the host CPU.

Profile

Documents that define a subset of the feature-rich Fibre Channel standard. Implementations of profiles do not need to support all options defined in the Fibre Channel standard.

Protocol Error Count

One of four error counters in the Frame Manager Link Error Status #2 Register. It indicates the number of protocol errors the Frame Manager detected.

Queue Entry Offset (Offset)

References the address position within a buffer queue entry pointed to by the Q_Index.

Read Status Completion Message

A 32-byte data structure indicating that read status has been processed.

Reassembly

The process of recombining data that was segmented into multiple frames into contiguous data. Reassembly is handled by the FC-2 layer.

PTI 172796

Received EOFa

One of four error counters in the Frame Manager Link Error Status #2 Register. It indicates the number of frames received with an EOFa delimiter.

Receiver Ready (R_RDY)

One of two FC-PH primitive signals, R_RDYs are used to manage the flow of frames on the link. Receivers send R_RDYs to indicate a frame was received and a frame buffer is available for receipt of another frame.

Receiver Transmitter Time Out Value (RT_TOV)

The number of milliseconds the port state machine waits for responses during point-to-point initialization. The default value is 15 milliseconds. The host should change this value to 100 milliseconds after it determines that no arbitrated loop exists.

Reject (RJT)

A link control frame used for flow control to indicate that a frame could not be processed, either by the fabric (F_RJT) or the N_Port (P_RJT).

Relative Offset (RO)

The displacement, expressed in bytes, of the first byte of a payload related to an upper level-defined-origin for a given information category.

Responder

The logical function in an N_Port responsible for supporting the exchange initiated by the originator in another N_Port. A Fibre Channel responder is comparable to a SCSI target.

Responder Exchange Identifier (RX_ID)

A 16-bit identifying value assigned by the responder of an exchange.

Routing Control (R_CTL)

A field in the frame header that identifies the frame as part of a FCP operation and identifies the information category. The 8-bit R_CTL field consists of two subfields. Bits 31:28 are routing bits, which provide routing information for the receiving hardware. Bits 27:24 represent the information category, which identifies the type of information in the frame payload.

Running Disparity (RD)

A binary parameter indicating the cumulative disparity, either positive or negative, of all previously transmitted characters. RD is used in encoding and decoding transmission characters so that an equal number of 1s and 0s are transmitted on the link. RD alternates after each positive or negative disparity character. That is, if the RD is positive, the negative encoding is used for the next character. If the RD is negative, the positive encoding is used. A neutral disparity character does not change the RD. RD is negative between frames of a sequence.

PTI 172797

SCSI Buffer Length Register

A 32-bit register that defines the length of the SCSI receive buffers. The value in this register should be the length (zero based) of the SCSI receive buffer. The contents of this register should not be modified after initialization.

SCSI Buffer Manager

A functional block within Tachyon responsible for supplying the Inbound Data Manager with addresses of buffers to be used for inbound SCSI data frames.

SCSI Descriptor Block (SDB)

A data structure that defines the logically contiguous data buffers into which inbound SCSI data is to be received. Each word in the SDB is an address which points to a buffer in host memory. The buffer size is defined in the scsi_buff_len register.

SCSI Exchange Manager

In conjunction with the SEST the SCSI Exchange Manager provides Tachyon with the hardware assists for SCSI I/O operations. On outbound transactions the SCSI Exchange Manager generates a command to the OSM to transfer the data, just as if the host had requested the data transfer. On inbound transactions the SCSI Exchange Manager reassembles incoming data in host memory, as specified in the SEST entry.

SCSI Exchange State Table (SEST)

A shared access, host based data structure used by Tachyon and the host. It consists of an array of 32-byte entries that contain information about where inbound data is to be placed in system memory (Inbound SEST), or where outbound data resides in host memory and what parameters to use to transmit the data (Outbound SEST). The location of the table is programmable and defined by the SEST Base Register. Tachyon uses the OX_ID as an index into the SEST when processing a SCSI transaction.

SCSI Exchange State Table (SEST) Base Register

A 32-bit register indicating the base address of the SEST in host memory.

SCSI Exchange State Table (SEST) Length Register

A 32-bit register that defines the length of the SEST by indicating the number (zero based) of 32-byte entries in the SEST.

Segmentation

The process of breaking up sequences that are too large to fit in a single frame into multiple frames. Segmentation is done by the FC-2 layer.

PTI 172798

Sequence

A set of one or more data frames having a common SEQ_ID, transmitted unidirectionally from one N_Port to another N_Port, with a corresponding response, if applicable, transmitted in response to each data frame. Sequences are used for managing phases within an operation. Sequences are the recovery boundary in Fibre Channel and the smallest Fibre Channel construct known to the ULP.

Sequence Count (SEQ_CNT)

A unique value in a frame that identifies the frame's order within a multiframe sequence.

Sequence Identifier (SEQ_ID)

An 8-bit identifier used to identify a sequence in a particular exchange. All frames of a sequence have the same SEQ_ID. SEQ_IDs do not need to be consecutive and may be reused after the sequence has been acknowledged.

Single Frame Sequence

A sequence in which the data is contained in a single frame.

Single Frame Sequence Buffer Length Register

A 32-bit register that defines the length of the SFS receive buffers, which should be the value (zero based) of the SFSBQ buffer length. The contents of this register should not be modified after initialization.

Single Frame Sequence Buffer Queue (SFSBQ)

Used by the host to inform Tachyon of the location in host memory of buffers that are used to receive sequences contained within a single frame. Each 32-byte queue entry contains eight pointers to receive buffers. The size of these buffers is programmable and is determined by the sfs_buff_len register. All buffers must be the identical size, and must be a power of two bytes in length. The host is the producer and Tachyon is the consumer of this queue.

Single Frame Sequence Buffer Queue (SFSBQ) Base Register

A 32-bit register containing the SFBQ base address in host memory. This register must not be written after initialization.

Single Frame Sequence Buffer Queue (SFSBQ) Consumer Index Register

A 32-bit register indicating the SFBQ entry that Tachyon is currently using to process single frame sequences.

Single Frame Sequence Buffer Queue (SFSBQ) Length Register

A 32-bit register defining the length of the SFBQ by indicating the number (zero based) of 32-byte entries in the SFBQ. This register must not be written after initialization.

Single Frame Sequence Buffer Queue (SFSBQ) Producer Index Register

A 32-bit register used by the host to pass to Tachyon the buffer to use for receiving incoming single frame sequences in host memory.

Source Identifier (S_ID)

The address identifier used to indicate the source port of the transmitted frame.

Start of Frame (SOF)

A 4-byte value delimiting the start of a frame and indicating the class of service.

Tachyon

Hewlett-Packard's high performance Fibre Channel interface controller chip that supports both networking and mass storage over a single host connection. Tachyon provides a high performance controller at an inexpensive cost due to its significant number of design features including, support of 1062, 531, 266, and 133 MBaud link speeds; support of Fibre Channel Class 1, 2, and 3 services; compatibility with FCSI profiles for FC-IP and FCP; compliance to industry standard MIB-II network management; support of up to 2 KByte frame payload for all classes of service; and many other features.

Tachyon BB_Credit Zero Timer Register

A 32-bit register that serves as a timer that runs whenever the BB_Credit for a sequence goes to zero. This information gives the host an indication of the amount of time Tachyon has waited to transmit because it had no BB_Credit. The timer count is incremented every 10 microseconds.

Tachyon Configuration Register

A 32-bit register used by the host to initialize and configure Tachyon's operating modes.

Tachyon Control Register

A 32-bit register used by the host to reset all or portions of the Tachyon chip, and to request a dump of all internal Tachyon status.

Tachyon EE_Credit Zero Timer Register

A 32-bit register that serves as a timer that runs whenever the EE_Credit for a sequence goes to zero. This information gives the host an indication of the amount of time Tachyon has waited to transmit because it had no EE_Credit. The timer count is incremented every 10 microseconds.

Tachyon Flush OX_ID Cache Entry Register

A 32-bit register used by the host to force Tachyon to write back any internal status it has for the indicated Read SCSI OX_ID and to invalidate the SEST entry by clearing the valid bit. This operation is performed by the host when it determines that there is an error with the exchange and it wants to terminate it abnormally.

Tachyon Receive Frame Error Counter Register

A 32-bit register that allows the host to monitor the number of frames transmitted by Tachyon whose response was a P_BSY.

Tachyon Status Register

A 32-bit register used by the host to read error information and operational status of the Tachyon chip.

Tachyon System Interface (TSI)

The interface between Tachyon and the host system, accomplished using a basic transaction protocol that has two major variants, Writes and Reads. Every transaction has a master and a responder. If the host is the master of a transaction, then Tachyon is the responder. If Tachyon is the master of a transaction, then the host is the responder.

Topology

The physical or logical layout of nodes on a network.

Transmission Character

Any valid or invalid encoded character, consisting of data and special characters, transmitted across a physical interface specified by FC-0. 8-bit data is encoded into 10-bit transmission characters using the 8B/10B encoding algorithm.

Upper Level Protocol (ULP)

The protocol user of FC-4, for example, SCSI or IPI.

World Wide Name (WWN)

A 64-bit identifier consisting of fields identifying the naming convention, the node, and its ports. The WWN must be unique within a Fibre Channel domain.

X_ID Interlock

This is used to suspend transmission of a sequence until the responder's X_ID value is sent to the originator, who then includes the RX_ID value in subsequent frames.

Index

Numerics

A

B

PTI 172802

C

PTI 172803

D

PTI 172804

E

PTI 172805

PTI 172806

F

PTI 172807

G

H

PTI 172808

PTI 172809

J

L

PTI 172810

PTI 172811

P

PTI 172812

Q

R

PTI 172814

S

PTI 172815

T

PTI 172816

U

V

W

X

PTI 172817

Reader Comment Sheet

TACHYON User's Manual

HP Manual Part No. 5965-1247E

We welcome your evaluation of this manual. Your comments and suggestions will help us improve our publications. Remove this page, mail or FAX to (916) 785-2875. Use and attach additional pages if necessary.

	Agree				Disagree	N/A
The manual is well organized.	○	○	○	○	○	○
The information is technically accurate.	○	○	○	○	○	○
Information is easy to find.	○	○	○	○	○	○
Step-by-step procedures are easy to perform.	○	○	○	○	○	○
There are enough examples and pictures.	○	○	○	○	○	○
The examples and pictures are useful.	○	○	○	○	○	○

Comments: _____

_____

_____

_____

_____

Name: _____ Date:

Title:

Company:

Address:

City & State: _____ Zip:

Country:

Phone:

Hewlett-Packard has the right to use submitted suggestions without obligation, with all such ideas becoming property of Hewlett-Packard.

__ Check here if you would like a reply

NCD/SIL Learning Products

PTI 172818

NO POSTAGE
NECESSARY
IF MAILED
IN THE UNITED
STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 256 ROSEVILLE, CA

- POSTAGE WILL BE PAID BY ADDRESSEE -

Attention: SIL Learning Products (M.S. 5601)

Hewlett-Packard Company
Networked Computing Division
8000 Foothills Boulevard
Roseville, CA 95747-9987

Fold Here

PTI 172819

Tape Please do not staple Tape

HEWLETT PACKARD

For technical assistance or the location of your nearest Hewlett-Packard sales office, distributor or representative call:

Americas/Canada: 1-800-235-0312 or (408) 654-8675

Far East/Australasia: (65) 290-6305

Japan: (81 3) 3331-6111

Europe: Call your local HP sales office listed in your telephone directory. Ask for a Components representative.

Data Subject to Change

Printed in U.S.A. 5965-1247E (5/96)

PTI 172820

TACHYON
Errata

HEWLETT® PACKARD

HP Errata Part No. 5965-1433E
Printed in USA May 13, 1996

Revision 5.0

PTI 172821

Notice

Copyright and Trademark Information

TACHYON is a trademark of the Hewlett-Packard Company.

Patents are pending for technical information presented in this document.

Warranty

The information contained in this document is subject to change without notice. Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Print History

Revision 1.0, June 2, 1995
Revision 2.0, July 31, 1995
Revision 2.1, August 3, 1995
Revision 2.2, August 3, 1995
Revision 2.3, August 4, 1995
Revision 2.4, August 15, 1995
Revision 2.5, August 18, 1995
Revision 2.6, September 18, 1995
Revision 3.0, October 2, 1995
Revision 3.1, October 5, 1995
Revision 3.2, December 4, 1995
Revision 4.0, January 25, 1996
Revision 4.1, January 25, 1996
Revision 4.2, February 14, 1996
Revision 4.3, February 20, 1996
Revision 5.0, May 13, 1996

Modification Information

A vertical line in the left margin (known as a change bar) indicates new or revised text from Revision 4.3 to Revision 5.0.

PTI 172822

Contents

PTI 172823

1. General Information

1.1 Purpose of Document

The TACHYON Errata is an additional document to the TACHYON User's Manual. This TACHYON Errata describes:

- Revision 2 Enchancements, which are new functionality features that were not available in Revision 1.
- User's Manual updates, which are corrections and additions to the Tachyon User's Manual, First Edition.
- Revision 2 discrepancies, which are differences between Tachyon's performance versus the information that is documented in the Tachyon User's Manual, First Edition.
- Revision 1 discrepancies, which have been fixed and verified in Tachyon Revision 2.

1.2 Revision 2 Enhancements

New functionality has been added to Tachyon Revision 2 that was not available in Tachyon Revision 1.

1.2.1 OLS/NOS bit Enhancement

When the NOS/OLS Received bit in the Frame Manager Status register is set to one, it indicates that "Tachyon received a NOS or OLS.". Previously, this bit indicated that "Tachyon received a NOS or OLS after a loop circuit was established.". Tachyon now sets this bit in N_Port mode also. Refer to "Section 5.10.3 Frame Manager Status register" in the Tachyon User's Manual, First Edition, p.144.

1.2.2 New SCSI Freeze Function/Invalidating Outbound SEST Entries

To invalidate and re-use Outbound and Inbound SEST Entries, refer to p. 64-65 in the TACHYON User's Manual, First Edition.

1.2.3 TSI Timing Enhancement

Tachyon Revision 1 TSI timing has been improved. Refer to sections "7.7 TSI Functional Waveforms" (p. 219) and "7.8 TSI Timing Requirements" (p.244) in the Tachyon User's Manual, First Edition.

PTI 172824

2. User's Manual Updates

This section provides corrections or additional information to the TACHYON User's Manual, First Edition. The First Edition will be available by the end of May. A01 through A07 have been updated in the TACHYON User's Manual, First Edition, or earlier drafts.

2.1 Fabric Login Required bit is Set in Internal Loopback Mode (A08/B15)

When Tachyon is initialized in internal loopback mode, the Fabric Login Required bit (in the Frame Manager Status register) may be is set to one after initialization completes.

2.2 Class 1 Point-to-Point Standard/Profile Interpretation of SOFc1 (A09/B04)

In Class 1 point-to-point operation, Tachyon may send an SOFc1 after an inbound connection has been established. This situation only occurs on host systems that have a very high memory read reply latency. All of the following conditions must exist for this situation to occur:

- Tachyon sends an SOFc1
- Tachyon starts the read of the first block of EDBs for the ODB
- Tachyon receives a BSY and an SOFc1 from the remote node before the read of the EDB block has completed
- Tachyon processes the SOFc1 at the same time that the read of the EDB is received

This situation never occurs if the Class 1 connection is opened with an ODB that contains only the SOFc1 header and does not contain any payload (EDBs). Refer to Class 1 Connections (B03).

If Tachyon is the remote node (i.e., if it is the node that receives the SOFc1 after the connection is opened), it sends the frame to the host as an unknown frame. When this occurs, the host should perform link reset.

This situation does not affect fabric connections because non-stacked fabrics discard the SOFc1 and stacked fabrics "stack" the SOFc1.

Note — A09/B04 is not a protocol violation or a bug, but rather a Standard/Profile interpretation.

2.3 SCLK Requirements/PCB Layout Suggestion (A10)

Because SCLK is the reference signal to a PLL within Tachyon, it is very important to minimize noise and jitter on this signal, as well as to minimize noise on the power and ground to the PLL. The following requirements help to ensure a clean and stable operational environment.

- Do not drive the SCLK input with the output of a PLL.
- Use good analog power and ground isolation techniques.
- Use wide PCB traces for analog Vdd and Vss connections to the PLL, separate from the Tachyon Vdd and Vss supplies.
- Use proper Vdd and Vss de-coupling.
- Use good power and ground sources on the PCB.
- Because jitter is affected by the noise frequency in the analog Vdd and Vss, keep the noise level on the PLLVdd less than 10 MHz.

Note	If these requirements are not followed, internal timing of the chip may be compromised.

2.4 Recommended Operationg Conditions (A11)

Section 7.2 Recommended Operating Conditions, p. 237 and Section 8.2 Recommended Operating Conditions. p. 249 of the First Edition has been modified (where the change bars are located):

Label	Parameter	Minimum	Typical	Maximum	Units
Vdd	Recommended Supply Voltage	3.0	3.3	3.6	Volts
f_{sclk}	TSI Clock Frequency, SCLK	24	-	40	MHz
dc_{sclk}	Duty Cycle, SCLK (Refer to Note 1 below)	40	-	60	%
$t_{r/f}$	Rise/Fall Time, SCLK (Refer to Note 2 below)	-	2.5	3.5	ns
j_{sclk}	Maximum Jitter of SCLK (peak-to-peak)	50	-	50	ps
TBC_r	TBC Rise Time	0.5	-	2.8	ns
RBC_f	RBC Fall Time	0.7	-	2.4	ns

Table 2.1 Recommended Operating Conditions

Recommended Operating Conditions Notes

1. For Duty Cycle for TBC and RBC:
   a. at 26 MHz, a maximum 40/60 Duty Cycle is allowed
   b. at 53 MHz, a maximum 45/55 Duty Cycle is allowed
2. The Rise/Fall Time of SCLK must be kept under 3.5 ns. A Rise/Fall Time at 2.5 ns or lower is preferred to minimize insertion delays and phase error.

3. Revision 2 Discrepancies

Tachyon Revision 2 discrepancies are differences between the performance of Tachyon versus the information that is documented in the TACHYON User's Manual, First Edition.

3.1 Continuous Interrupts when Cable is Disconnected (B01)

If the cable is disconnected from the GLM or if the cable is defective, the host may receive continuous Frame Manager interrupts. The Frame Manager Status register will indicate a Loss of Signal and Elastic Store Error.

Suggested Resolution

Before attempting to go online, check the Frame Manager Status Register. If any of the conditions - Loss of Signal, Laser Fault or Out of Sync - exist, do not go online until the condition(s) has disappeared.

If Tachyon has just gone on-line, the following modification to the Tachyon C code will resolve any problems:

Chapter : Tachyon C Code
Section : SEST
Sub-section : `frame_manager_interrupt()`

Replace the code starting at

```
if (FM_STATUS_LF & status)
```

and ending just before

```
if (FM_STATUS_CE & status)
```

with the following:

```
if (FM_STATUS_LF & status)
   {
   /* We had a link failure.  Now check the upper bits of the status
    * to find out exactly what the problem is and report it back to
    * the user.  Then we can decide what to do.
    */
   if (FM_STATUS_LS & status || FM_STATUS_OS & status)
      {
      printf("LOSS OF SIGNAL, or Out of Sync.  ");
      printf("Check Link Module and its connections.\n");
      *ptr_fmcntl_control = FMCNTL_CMD_OFFLINE;
      sleep_until_fixed();
      }
   /* we have to go offline to prevent continuous interrupts from occurring.*/
   /* sleep_until_fixed() function sleeps until loss of signal, or
    * Out of Sync condition is corrected.  User must supply.
    * User can verify condition is fixed by polling the FM Status register
    * After it is fixed, we need to clear the link failure.
    * We still want to check for any additional faults.
    */
   if (status & FM_STATUS_LFLT)
      {
      /* Laser fault, so let's hit the ewrap bit on the GLM ser fault,
       * Note this happened while we were online supposedly,
       * so let's make sure we are OFFLINE.
```

PTI 172827

```
            * OFFLINE we go FIRST if we aren't there already.
            * If we're in a loop this MAY cause problems.
            * The other choice here is to reset the chip and apply EWRAP.
            */
           *ptr_fmcntl_control = FMCNTL_CMD_OFFLINE;
           /* Wait for us to go offline.  Sleep 1 second.
            */
           sleep(1);
           /* OK, turn on the Ewrap bit, wait one second for it to clear,
            * then turn it back on.  Control could be passed back to the
            * system here temporarily.
            */
           printf("LASER FAULT\n");
           *ptr_fmcntl_config = *ptr_fmcntl_config | FM_ENABLE_EXT_LB;
           sleep(1);
           *ptr_fmcntl_config = *ptr_fmcntl_config & (~FM_ENABLE_EXT_LB);
           link_is_up = FALSE;
           were_in_a_loop = FALSE;
           }
        if (FM_STATUS_TP & status)
           {
           /* A parity error.  We need to reinitialize the chip.
            * Proper thing to do is to assert the reset line on Tachyon
            * and rewrite the registers.  Instead, report the error and
            * just put the chip down.  Note, we probably need to free
            * headers and EDB data here as well.
            */
           free_all();
           printf("Parity Error\n");
           exit(PARITY_ERROR);
           }
        /* If we made it here then we are ready to try and bring the
         * link back up.  We need to clear the link failure to do this.
         */
        *ptr_fmcntl_control = FMCNTL_CMD_CLEAR_LF;
        if (!were_in_a_loop)
           {
           /* Reset the BB Credit, because we are connected to an N_PORT.
            */
           *ptr_fmcntl_config = *ptr_fmcntl_config | bb_credit;
           }
}
```

PTI 172828

3.2 ERROR_L (B05)

The "Tachyon Response Under ERROR_L" section in "7.6.15 Error Handling" of the TACHYON User's Manual, First Edition, describes how Tachyon ideally responds under ERROR_L. The following text describes how Tachyon realistically responds under ERROR_L.

- If a host-mastered Write with bad data parity occcurs:
  - If ERROR_L is already asserted, Tachyon will receive the data and assert READY_L.
  - If ERROR_L has not been asserted, Tachyon will receive the data, assert READY_L, and assert ERROR_L on this transaction.
- If a host-mastered Read or Write with bad address parity occurs:
  - If ERROR_L is already asserted, Tachyon will send or receive the data and assert READY_L.
  - If ERROR_L has not been asserted, Tachyon will send or receive the data and assert ERROR_L, but will not assert READY_L.

Workaround: Instead of waiting for Tachyon to assert READY_L during ERROR_L assertion, use a hardware timer to generate READY_L back to the host. This allows the transaction to complete so that subsequent slave transactions can be performed.

3.3 TSI Protocol Violation when ERROR_L is Asserted (B10)

When Tachyon detects a parity error on any data word of a DMA mastered Read transaction, it asserts ERROR_L. Before the DMA mastered Read transaction completes, the TBR_L [1:0] signals may change from "01" to "10" or from "10" to "01". This situation violates TSI Protocol. At the end of this errored transaction, the violation is cleared (TBR_L [1:0] signals correctly change from either "01" or "10" to "11") and Tachyon enters a fatal error mode where no DMA activity occurs.

Workaround: Ignore TBR_L[1:0] when ERROR_L is asserted.

3.4 Forcing Frame Manager Offline if in Loop Fail State (B11)

If Frame Manager is configured for loop and is in the Loop Fail state (Refer to the Loop State Machine field in Frame Manager Status Register), writing the Offline command to the Frame Manager Control Register does not force Frame Manager offline.

Workaround:

Read the Frame Manager Status Register. If the Loop State Machine (lsfm) is in the Loop Fail state (0xd0), then perform the following steps to force Frame Manager offline:

1. Write the Host Control command to the Frame Manager Control Register to force Frame Manager into the Host Control state.
2. Write the Offline command to the Frame Manager Control Register to force Frame Manager offline.

3.5 Close Outbound Class 1 Connection Before Sending Class 2 Sequences (B13)

If an outbound Class 1 connection is open and Tachyon sends a Class 2 sequence with the X_ID Interlock (Lck) bit in the ODB set to one, Tachyon closes the outbound Class 1 connection if the first frame of the Class 2 sequence is BSY'd, RJT'd, or ACK_ABT'd.

Workaround: Close outbound Class 1 connections before sending Class 2 sequences or do not send Class 2 sequences with the Lck bit in the ODB set to one

PTI 172829

3.6 Tachyon Violates New Additions to the FC-AL Specification (B17a)

Tachyon violates FC-AL, Revision 4.5, Section 10.4.3, Step 2), sub-article d), p. 55, which states, "If anything else is received, it shall be discarded by the L_Port and the L_Port shall continue with steps (2)(a) to (2)(d)."

During Loop Initialization, if Tachyon receives an errored LISM frame from an L_Port, Tachyon recognizes it as an errored frame and discards it. However, if Tachyon receives another errored LISM frame from the same L_Port, it forwards it. This may cause the sending L_Port to receive its errored frame, assume that it is good, and continue to the next Loop Initialization phase (LIFA). At this time, Tachyon and the L_Port are "out of synchronization", since Tachyon is still waiting for a valid LISM frame. Loop Initialization times-out and retries.

3.7 Sending High Priority Frame May Block Outbound Channel (B22)

On a loop, if an active multiframe sequence is in progress and the host tries to send a high priority frame to a different AL_PA, then the high priority frame and the frames of the active multiframe sequence cannot be sent, the loop remains open, and the outbound channel hangs. B22 only applies to FC-AL.

Workaround:

To send high priority frames via the HPCQ while on loop, perform the following:

1. The host should send the high priority frames with outbound_high_priority completion messages and interrupts enabled.
2. After the outbound_high_priority completion message is received, the host must immediately issue a Close Loop Request (cl bit in the Frame Manager Control register should be set). This causes the Frame Manager to close the loop so that it can re-arbitrate to send the high priority frame.
3. Set ED_TOV so that the following operations can occur:
   a. The host responds to the outbound_high_priority completion message.
   b. The host issues the Close Loop Request.
   c. The Frame Manager re-arbitrates.
   d. The Frame Manager sends the high priority frame.
   e. The Frame Manager re-arbitrates.
   f. The Frame Manager sends the next frame of the MFS that was in progress.

PTI 172830

3.8 Assertion of READY_L May Cause a DMA WRITE1 to Address 0 (B24)

If the host interface asserts the READY_L signal during the data cycles of certain DMA Write operations, Tachyon may perform a DMA WRITE1 to address 0.

The DMA Write operations that are affected by this problem include the following.

1. Update of the OCQ Consumer Index (WRITE1)
2. Update of the HPCQ Consumer Index (WRITE1)
3. Update of a SEST entry (WRITE1, WRITE2, and WRITE4)

Workarounds:

1. Delay the assertion of READY_L until the final data cycle of all DMA Write operations (WRITE1, WRITE2, WRITE4, and WRITE8) completes.
2. Delay the assertion of READY_L until the final data cycle of WRITE1, WRITE2, and WRITE4 operations completes. Assert READY_L during any data cycle for WRITE8 operations. (Note: Typically, WRITE8 operations are used for most DMA Write operations, therefore, implementing this workaround causes only a small effect on performance.

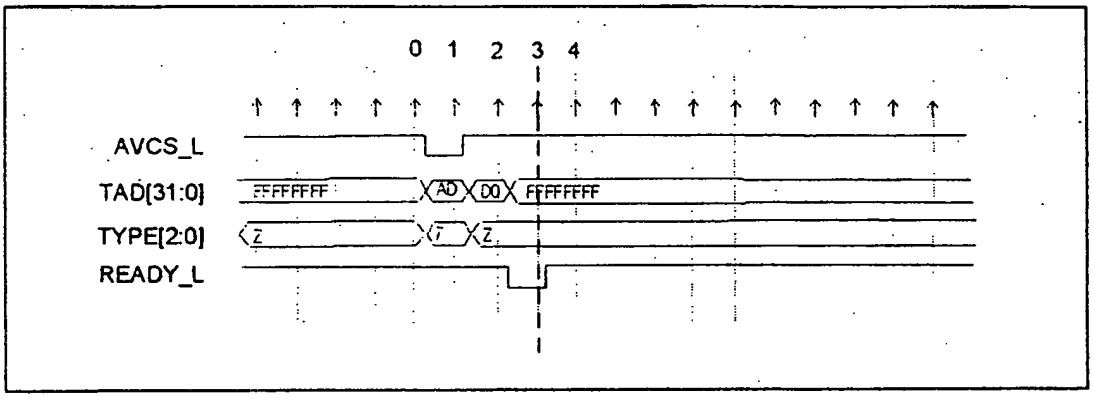

Figure 3.1 Assert READY_L After the Last Data Cycle of a WRITE1 Completes

In the Figure 3.1, the address cycle (TAD[31..0]) occurs on rising edge #1 of SCLK. The data cycle occurs on rising edge #2. If the host interface asserts READY_L so that Tachyon detects the signal on rising edge #3 (or later), then there is no problem.

Note:

B24 is similiar, but not related, to B08 (Slave Must Not Assert READY_L Prior to Last Data Word). The differences are described below:

B08: The workaround for B08 was to delay the assertion or READY_L until the last data cycle (at the earliest) for all DMA Write operations.

B24: The workaround for B24 requires one extra cycle of delay, but only for WRITE1, WRITE2, and WRITE4 operations.

PTI 172831

3.9 Certain SFSs Using ACK_0 Cause Tachyon to Hang (B25)

Tachyon will hang if the EE_Credit field in the ODB is set to a value other than one when sending the following single frame sequences.

1. An outbound SFS with the SOFc1 and the ACK_0 bits (in the ODB) set to one.
2. An outbound SFS with the Sequence Interlock and the ACK_0 bits (in the ODB) set to one.

Workaround: Set the EE_Credit field in the ODB to one when sending the single frame sequences described above.

3.10 Tachyon May P_BSY an ABTS Frame for a Class 2 SFS (B26)

When Tachyon is a target and has a Class 1 connection open, and the initiator node sends a Class 2 SFS, Tachyon receives the frame correctly and responds with an ACK. If this ACK gets lost and the initiator sends an ABTS to Tachyon with a SOFn2, Tachyon receives the ABTS as part of an MFS and P_BSYs the ABTS. If the initiator keeps resending the ABTS in response to the P_BSY, this situation continues until:

1. The Class 1 connection completes and Tachyon responds to the ABTS appropriately, or
2. The initiator node exceeds its limit of retries

Initiator Workaround: Send the ABTS with an SOFi2.

Target Workaround:

Set the Disable AUTO P_BSY bit in the Tachyon Configuration register to one.

If Tachyon receives an ABTS with an SOFn2, it sends an inbound_busied_frame completion message to the host. The inbound_busied_frame completion message does not generate an interrupt, therefore, the host does not become aware of receiving the ABTS until it polls the IMQ Producer Index or receives another interrupt. When the host is aware of the ABTS, it manages the proper Reply Sequence, BA_RJT or BA_ACC, for any detected MFS collision that would normally receive a Reply Sequence of P_BSY.

PTI 172832

Appendix A. Revision 1 Discrepancies

The following Tachyon Revision 1 Discrepancies have been fixed and verified in Tachyon Revision 2. This section will be deleted in the next TACHYON Errata Revision.

A.1 Header Checksum Assists May Cause Error (B02)

Do not enable Header Checksum Assists at this time. More information about implementing Header Checksum Assists will be available in future versions of the 3.3V Production TACHYON Errata.

Status: B2 has been fixed in Tachyon Revision 2. No known workarounds are available for Tachyon Revision 1.

Revision 2 Verification: The B02 fix has been verified in simulation and in hardware.

A.2 Class 1 Connections (B03)

- To open a Class 1 connection, send Class 1 sequences using two Outbound Descriptor Blocks (ODBs). The first ODB should contain only the SOFc1 header (no data payload is included). The first ODB opens the connection and keeps the connection open (End Connection bit must be cleared to zero). The second ODB should contain the data payload with a SOFi1 header.
- Do not use Class 1 with outbound SCSI assists.

Status: B03 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B03 fix has been verified in simulation.

A.3 Class 1 Point-to-Point Standard/Profile Interpretation of SOFc1 (B04)

B04 is a documentation issue. Refer to "Class 1 Point-to-Point Standard/Profile Interpretation" of SOFc1 (A09).

A.4 Do Not Set Both Lck and F bits in ODB (B06)

When the Sequence Interlock (Lck) bit and the Frame Boundary (F) bit in the ODB are set and Tachyon receives a BSY for the frame, Tachyon may enter an unknown state.

Workaround: Do not set both Lck and F bits in the ODB.

Status: B06 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B06 fix has been verified in simulation.

A.5 SCSI Write Operations of 1 MByte or Greater (B07)

When an Initiator Tachyon receives an FCP_XFER_RDY frame during a SCSI H/W Assisted exchange, bits [31..20] of the BURST_LEN field are ignored. Thus, the Tachyon SCSI H/W Assistance does not work properly for SCSI Write operations of 1MByte or greater.

Initiator Workarounds:

Two possible work-arounds exist for Tachyon Initiators:

1. Break up large SCSI commands (1MByte or greater in length) into smaller SCSI commands [(1MB-4bytes) or less in length]. Tachyon will handle these smaller SCSI commands properly.
2. Do not use the SCSI H/W Assistance for SCSI Write operations of 1MByte or greater. If no valid SEST Entry exists for the received FCP_XFER_RDY frame, Tachyon will pass it to the host as a Bad_SCSI frame. The host can then build an OCQ entry which causes Tachyon to send the requested data to the Target. Note that it is possible to send data via the OCQ with a Tot_Len value of up to 32 bits.

Target Workaround:

If it is known that the Initiator is a Tachyon, then the Target can avoid this problem by breaking the transfer up into two data phases:

1. When the Target receives a SCSI Write command of 1MByte or greater, the Target should return an FCP_XFER_RDY frame with BURST_LEN set to less than 1MByte. The SCSI H/W Assistance in the Initiator Tachyon will send this SCSI Write data automatically.
2. When the first data phase is complete, the Target returns a second FCP_XFER_RDY frame requesting the remaining SCSI Write data. Since the Initiator Tachyon will pass to its host any FCP_XFER_RDY frame in which the DATA_RO field is not equal to zero, SCSI H/W Assistance will be bypassed in this case. The value of BURST_LEN in the second FCP_XFER_RDY frame can be any 32 bit value. The Initiator software will build an OCQ entry to send the remaining SCSI Write data.

Status: B07 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B07 fix has been verified in simulation.

PTI 172834

A.6 Slave Must Not Assert READY_L Prior to Last Data Word (B08)

If Tachyon is mastering a Write transaction and the slave asserts READY_L prior to the last word of the Write transaction, then Tachyon may continuously write to address zero.

Workaround: The slave must assert READY_L on the last data word of the Write transaction or later. To maximize TSI performance, the slave should assert READY_L on the last data word of the transaction. The slave may need to monitor the transaction TYPE [ ] signals to assert READY_L on the last data word. The slave can assert READY_L after the last data word, but TSI performance will be compromised.

Example of a slave asserting READY_L on the last data word of a Tachyon-mastered WRITE8 transaction:

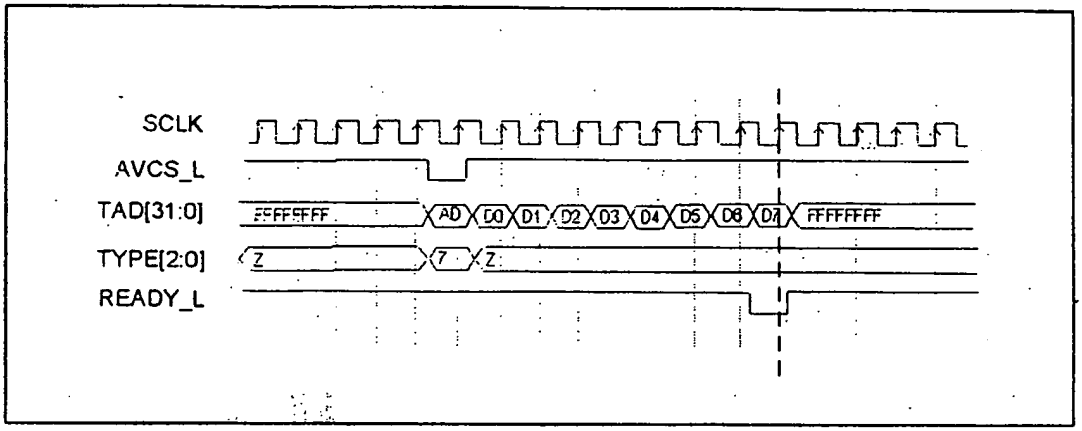

Figure A.1 Slave Asserts READY_L on Last Data Word of a WRITE8 Transaction

Status: B08 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B08 fix has been verified in simulation.

PTI 172835

A.7 TBR_L Signals Not De-asserted (B09)

If the arbiter does not keep the bus grant signal (TBG_L) asserted during the entire bus tenancy, Tachyon may not de-asserted both bus request signals (TBR_L[1:0]) for at least one clock cycle when switching from a bus tenancy under one bus request signal to a new bus tenancy under the other bus request signal. The TBR_L[1:0] signals can change from "01" to "10" without transitioning to "11". This situation violates TSI Protocol. Refer to "Workaround" below.

If the TBR_L signals are combined logically to form a single bus request signal, then de-assertion of bus request may not be seen. The arbiter does not know to assert the bus grant signal for the second TSI tenancy. Tachyon hangs.

If the arbiter keeps the bus grant signal (TBG_L) asserted during the entire bus tenancy, this TSI violation does not occur. Tachyon de-asserts one bus request signal, waits for the bus grant signal to be deasserted, then asserts the other bus request signal. This guarantees that the arbiter knows that Tachyon de-asserted the first bus request signal.

Workaround: Add external circuitry to correct the situation. Refer to the following verilog RTL code.

```
reg  [1:0]  TSI_TBR_L;
reg  [1:0]  delay_int_tbr_l;
wire int_tbr_l_1;  /*  connect to Tachyon's p_tbr_l [1]  */
wire int_tbr_l_0;  /*  connect to Tachyon's p_tbr_l [0]  */

always  @  (int_tbr_l_1 or int_tbr_l_0 or TSI_TBR_L)
begin
   case  (TSI_TBR_L)
   2'b11:    if (~int_tbr_l_0)
         delay_int_tbr_l = 2'b10
           else
         if (~int_tbr_l_1)
            delay_int_tbr_l = 2'b01
             else
            delay_int_tbr_l = 2'b11
   2'b10:    if (int_tbr_l_0)
         delay_int_tbr_l  = 2'b11
           else
         delay_int_tbr_l = 2'b10
   2'b01:   if (int_tbr_l_1)
         delay_int_tbr_l  =  2'b11
           else
         delay_int_tbr_l  =  2'b01
   endcase
end

always @ (posedge TSI_SCLK or negedge TSI_RESET_L)
if   (~TSI_RESET_L)
   TSI_TBR_L = 2'b11;
else
   TSI_TBR_L = delay_int_tbr_l;
```

Status: B09 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B09 fix has been verified in simulation.

PTI 172836

A.8 NOP Causes Inbound SEST Entry Invalidation (B12)

When Tachyon receives an NOP on an Inbound SEST Entry, Tachyon invalidates the Inbound SEST Entry, and sends all subsequent frames to the host as bad SCSI frames.

Status: B12 has been fixed in Tachyon Revision 2. No known workarounds are available for Tachyon Revision 1.

Revision 2 Verification: The B12 fix has been verified in simulation.

A.9 Shift Word Problem for Incoming Frame during a Link Reset (B14)

If a Link Down condition occurs while Tachyon is writing the first word (SOF) of the FC Header of an inbound frame into the Inbound Data FIFO, the SOF is not purged. This causes the next frame received (after the link comes back up) to have an extra SOF word, which shifts all remaining frame header/payload data by one word. This results in the frame being passed to the host as an unknown frame. Frames received after this frame will be received correctly.

Workaround: After the link comes back up, the host should discard the first frame received if it is an unknown frame.

Status: B14 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B14 fix has been verified in simulation and in hardware.

A.10 Fabric Login Required bit is Set in Internal Loopback Mode (B15)

B15 is a documentation issue. Refer to 2.8 Fabric Login Required bit is Set in Internal Loopback Mode (A08).

A.11 Internal Loopback Test May Fail due to RBC Violating GLM Specifications (B16)

If the Receive Byte Clock (RBC) violates GLM RBC specifications, an internal loopback test failure may occur. Use of a loopback hood guarantees that the RBC meets GLM specifications.

It is not necessary to have a GLM present to perform an internal loopback test. If RBC is being violated, remove the GLM before performing the test.

Status: B16 is not a Tachyon bug. Nothing will be changed in Tachyon Revision 2.

A.12 Tachyon Violates New Additions to the FC-AL Specification (B17b)

Tachyon violates FC-AL, Revision 4.5, Section 10.4.1, Figure 4, p. 52, which depicts Word 3 of the Frame_Header being cleared to 0x00000000. Tachyon Revision 1 sends all Loop Initialization frames with SEQ_ID (bits 31..24 of Word 3) = 0x01.

Status: B17b has been fixed in Tachyon Revision 2. Tachyon Revision 2 will send Loop Initialization frames with a SEQ_ID = 0x00.

Revision 2 Verification: The B17b fix has been verified in simulation and in hardware.

A.13 Shift Word Problem for Incoming Frames during a Link Reset (B18)

B18 and B14 describe the same bug. Refer to A.9 Shift Word Problem for Incoming Frames during a Link Reset (B14).

PTI 172837

A.14 Link Down/Link Up Transistion May Cause Inbound Channel to Become Blocked (B19)

When the link goes down and then comes back up, the Inbound Data FIFO may not reset due to a synchronization problem.

If the host has not received inbound completion messages for awhile (provided there are SFS and MFS buffers available to Tachyon), but the Receive FIFO Empty bit (mt) in the Tachyon Status register is clear, then the inbound channel is blocked. When this occurs, the host must perform a soft reset.

Status: B19 has been fixed in Tachyon Revision 2. No known workarounds are available for Tachyon Revision 1.

Revision 2 Verification: The B19 fix has been verified in simulation and in hardware.

A.15 Tachyon Cannot Clear the Non-Participating bit (B20)

Once the Non-participating bit (np) is set to one in the Frame Manager Status register, Tachyon cannot clear this bit to zero. The host must perform a hard or soft reset for Tachyon to exit non-participating mode.

Status: B20 has been fixed in Tachyon Revision 2. Tachyon Revision 2 will clear the np bit to zero when Tachyon begins to participate on the loop.

Revision 2 Verification: The B20 fix has been verified in simulation and in hardware.

A.16 Outbound FC Header Corruption May Occur at 266 MBaud (B21)

With the following conditions...

1. Operating at 266 Mbaud data rate mode (PAR_ID [1..0] = 01)
2. Using Class 1 or Class 2
3. Having an outbound FC frame size greater than 2012 bytes

...a single word of the outbound FC header may be corrupted if Tachyon receives an inbound ACK occurring approximately at or prior to the start of the transmission of the outbound FC header.

Workaround: Limit the maximum outbound FC frame size to 2012 bytes or less.

Status: B21 has been fixed in Tachyon Revision 2.

Revision 2 Verification: The B21 fix has been verified in simulation.

A.17 Data Corruption in Pre-Production Tachyon Revision 1 Chips (B23)

Pre-Production 3.3V Tachyon Revision 1 chips contain an internal noise susceptibility which occasionally results in corruption of inbound data.

This information has been previously communicated with all Tachyon customers. The "data corruption white paper" that was sent to all customers outlines the system test configuration which most readily causes the data corruption to occur. It also includes information about the relative susceptibility of the various lot numbers for Revision 1 silicon. This information is helpful for those customers who want to screen Revision 1 chips as an interim measure prior to receiving Revision 2 chips. To receive a copy of this "data corruption white paper", e-mail a request to tachyon@hp.com or call 1-800-TACHYON (1-800-822-4966) and ask for the "data corruption white paper".

Status: B23 has been isolated, understood, and resolved. B23 has been fixed in Tachyon Revision 2. No known workarounds are available for Tachyon Revision 1.

Revision 2 Verification: The B23 fix has been verified in hardware.

TACHYON
Errata

HEWLETT® PACKARD

HP Errata Part No. 5965-1433E
Printed in USA May 13, 1996

Revision 5.0

PTI 172839

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.
As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.