



Maxtor Corporation
500 McCarthy Boulevard
Milpitas, CA 95035 USA

To: Serial Attached SCSI Protocol Working Group
From: Mark Evans
Phone: 408-894-5310
Email: mark_evans@maxtor.com
Date: 08 July 2002

Subject: SAS transport layer state machines for SSP

Introduction

In this revision, only the SCSI initiator device elements are included. This proposal has been revised based on discussion at the SAS working group meetings, June 24 through 26, 2002. In addition, the “application client” state machine has been renamed “application client”, and the “application client” state machine has been renamed “application client thread” to conform to a proposed letter ballot resolution for SAM-2. The SCSI target device elements are under construction and will be included in a subsequent revision of this proposal.

9.2.7 SSP transport layers

9.2.7.1 Overview

The SSP transport layer contains state machines that process service requests from the application layer and return service responses to the application layer. These state machines reside in the SCSI initiator device and the SCSI target device. The SSP transport state machines are as follows:

- a) Application client (AC state machine);
- b) Application client thread (ACT state machine);
- c) Task router (TR state machine);
- d) Task manager (TM state machine); and
- e) Device server (DS state machine).

9.2.7.2 Initiator device state machines

9.2.7.2.1 Initiator device state machine overview

The initiator device state machine processes service requests from the application layer. These service requests may be tasks (i.e., commands or linked commands) or task management functions. The AC (application client) state machine receives service requests from the application layer and creates an ACT (application client thread) state machine for each request. The AC state machine may also receive service requests from the application layer after an AC state machine has been created (e.g., an ABORT TASK task management function) and uses these to further manage the ACT state machine. The AC state machine also communicates with the port layer to specify when resources are available so that connections may be accepted or rejected.

The AC state machine contains the following state:

- a) ST_AC:Application_Client state.

Once created, the ACT state machine:

- a) generates a SSP frame containing a device service request or device task management request and sends a Transmit Frame(Interlocked) request to the port layer;
- b) creates any SSP data frames associated with the request and sends a Transmit Frame(Non-interlocked) request to the port layer for each data frame;
- c) receives and processes any SSP Receive Frame confirmations from the port layer;
- d) receives and processes any transmission information received from the port layer (e.g., Transmission Status confirmation); and
- e) sends a response to the application layer and the application client when the task or function is complete.

The ACT state machine contains the following states:

- a) ST_ACT1:Send state;
- b) ST_ACT2:Prepare_Send_Data state;
- c) ST_ACT3:Receive state;
- d) ST_ACT4:Process_Receive_Data state; and
- e) ST_ACT5:Process_Response state.

[Editor's note: I think that the following list of requests and confirmations (and all other inter-layer requests and confirmations) should be in one location and referenced in the corresponding state diagrams to avoid duplication.]

The AC state machine receives the following requests from the application layer:

- a) Task service request; and
- b) Task management service request.

The AC state machine sends the following requests to the port layer:

- a) Accept_Reject OPENS(Accept); and
- b) Accept_Reject OPENS(Reject).

The AC state machine receives the following confirmation from the port layer:

- a) Unknown Tag.

The ACT state machine receives the following request from the application layer:

- a) Data-out Available.

The ACT state machine sends the following confirmations to the application layer:

- a) Service Response; and
- b) Data-in Data Available.

The ACT state machine sends the following requests to the port layer:

- a) Transmit Frame(Interlocked);
- b) Transmit Frame(Non-interlocked); and
- c) Cancel.

The ACT state machine receives the following confirmations from the port layer:

- a) Transmission Status(Frame Transmitted)
- b) Transmission Status(Open Failed-No Destination)
- c) Transmission Status(Open Failed-Open Timeout)
- d) Transmission Status(Connection Lost)
- e) Transmission Status(Disconnect in Process)
- f) Transmission Status(Open Failed-Retry)
- g) Transmission Status(Open Failed-Wrong Destination)
- h) Transmission Status(Open Failed-Link Rate Not Supported)
- i) Transmission Status(Open Failed-Protocol Not Supported)
- j) Transmission Status(Open Failed-Pathway Busy)
- k) Transmission Status(Open Failed-Open Timeout Occurred)
- l) Transmission Status(Open Failed-Port Layer Request)
- m) Transmission Status(Open Failed-Break Received)
- n) Transmission Status(Open Failed-No Destination)
- o) Transmission Status(Open Failed-Bad Destination)
- p) Transmission Status(Open Failed-STP Resources Busy)
- q) Transmission Status(Open Failed-I_T Nexus Loss Timeout)
- r) Transmission Status(Open Failed-Physical Not Ready)
- s) Transmission Status(Credit Timeout)
- t) Connection Closed(Break Received)
- u) Connection Closed(Close Timeout)
- v) Connection Closed(Link Broken)
- w) ACK received;
- x) NAK received;
- y) ACK/NAK timeout;
- z) Cancel Acknowledge;
- aa) Frame Received(ACK/NAK Balanced); and
- bb) Frame Received(ACK/NAK Not Balanced).

Figure 1 describes the SSP initiator device state machine.

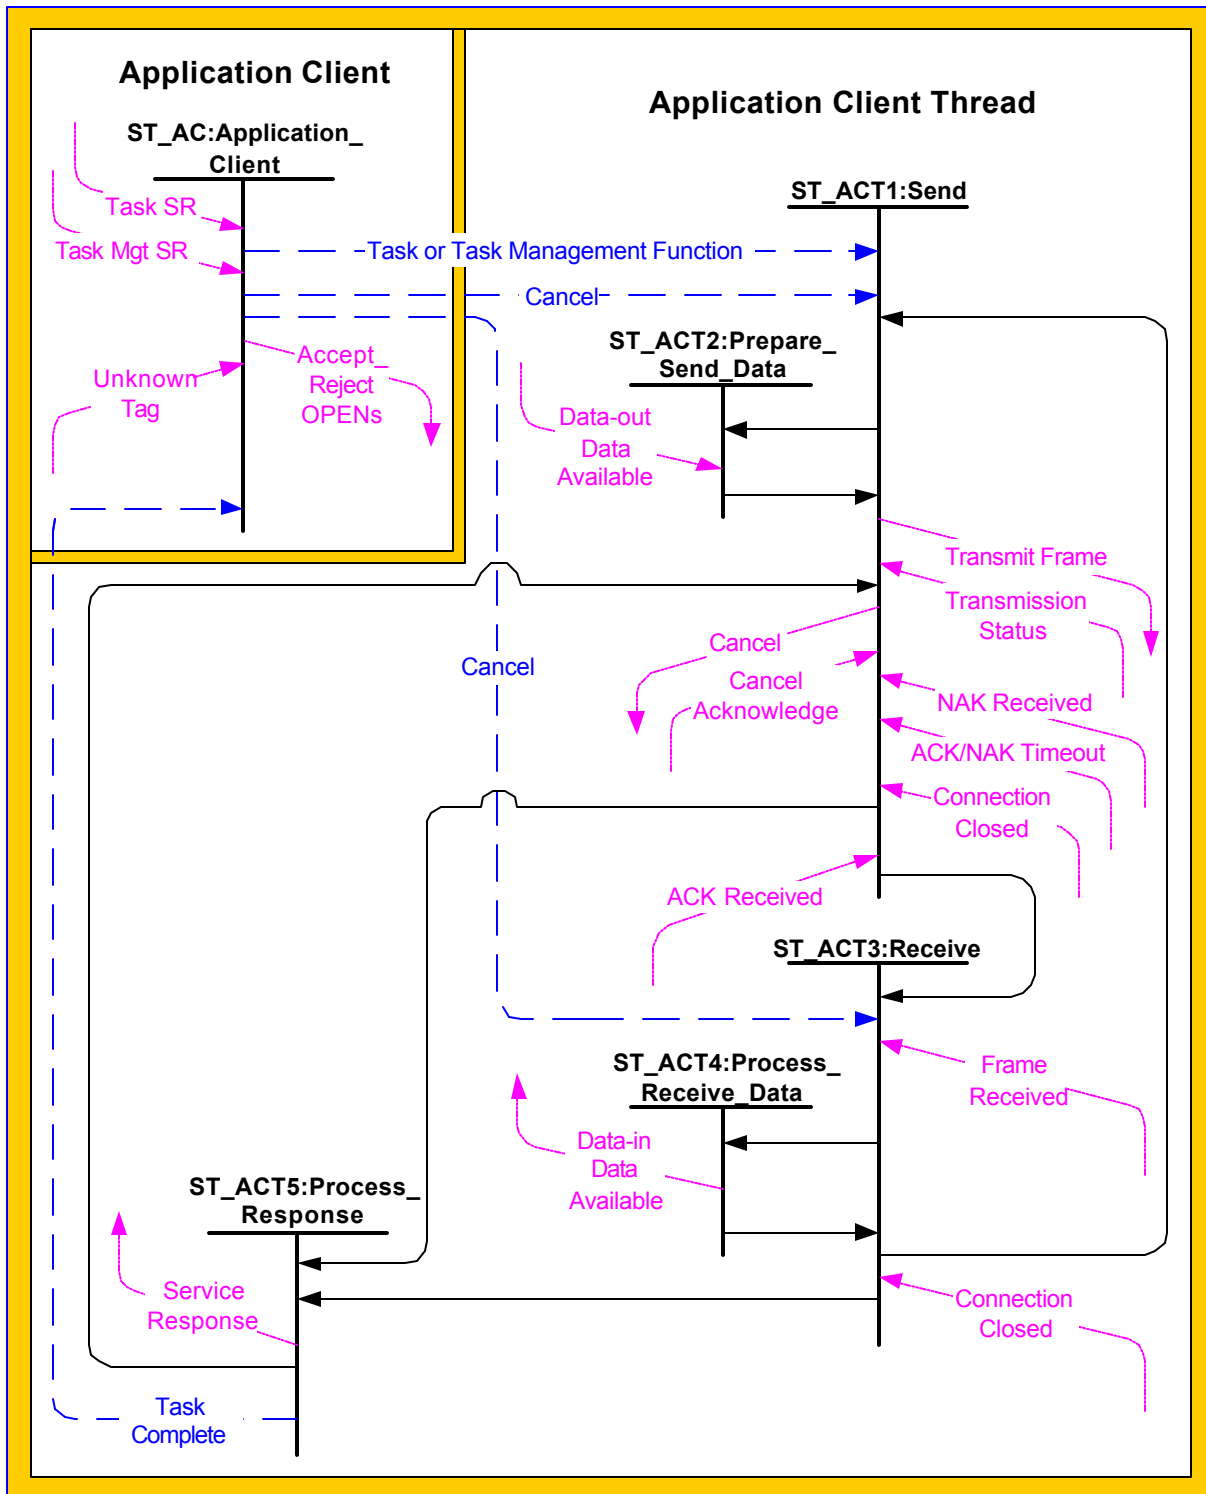


Figure 1. SSP initiator device state machine

9.2.7.2.2 ST_AC:Application_Client state

9.2.7.2.2.1 ST_AC:Application_Client state description

This state machine receives service requests for tasks or task management functions from the application layer. A request for a task or task management function from the application layer shall consist of values to be used in the LINK RATE, INITIATOR CONNECTION TAG, DESTINATION DEVICE NAME, and SOURCE DEVICE NAME fields in any OPEN address frame, and the LUN, CDB, AND TASK ATTRIBUTE fields in the SSP frames for the task or task management function. An ACT state machine shall be created to process the task or task management function.

When creating an ACT state machine the AC state machine shall send one of the following sets of parameters to the ST_ACT1:Send state:

- a) A request for a task to be processed (either a non-linked or linked command); or
- b) A request for a task management function to be processed.

For either type of request, the sets of parameters passed to the ST_ACT0:Send state shall include the values to be used in the LINK RATE, INITIATOR CONNECTION TAG, DESTINATION DEVICE NAME, and SOURCE DEVICE NAME fields in any OPEN address frame, and values for the COMMAND ID, TAG, LUN, CDB, AND TASK ATTRIBUTE fields in the SSP frames for the task or task management function. The AC state machine shall manage all COMMAND ID and TAG values for all task or task management requests for the SCSI initiator device.

The AC state shall send an Accept_Reject OPENS(Accept) request to the port layer to inform that machine that it may accept an attempt to open a connection by a target port. The AC state shall send an Accept_Reject OPENS(Reject) request to the port layer to inform that machine that it shall reject any attempt to open a connection by a target port.

The AC state may receive a task management function request from the application layer to cancel an application client thread in progress. The AC state shall then send a Cancel parameter to the ST_ACT1:Send and ST_ACT3:Receive states.

The AC state shall receive a Task Complete parameter from the ST_ACT5:Process_Response state when the ACT state machine has completed a task or task management function.

9.2.7.2.3 ST_ACT1:Send state

9.2.7.2.3.1 ST_ACT1:Send state description

This state is initiated when a task or task management function is received from the application client.

This state shall construct the frame containing a COMMAND IU or the TASK IU (including calculating the values for the HASHED DESTINATION DEVICE NAME and the HASHED SOURCE DEVICE NAME, and setting the value of the TARGET PORT TRANSFER TAG field to the value from a received XFER_READY information unit) and shall send a Transmit Frame(Interlocked) request to the port layer. Along with the request, this state shall send to the port layer the values to be used in the LINK RATE, INITIATOR CONNECTION TAG, DESTINATION DEVICE NAME, and SOURCE DEVICE NAME fields in any OPEN address frames. This state shall receive confirmation (i.e., ACK Received) from the port layer before sending another Transmit Frame request.

If the ACK Received confirmation is received for a task or task management frame is and additional frames are required to be transmitted for the task or task management function, this state shall request data from the ST_ACT2:Prepare_Send_Data state and construct the DATA frames.

If the negotiated value in FIRST BURST SIZE field in the Disconnect-Reconnect mode page for the SCSI target device or the unused portion of the value from the WRITE DATA LENGTH field in an XFER_RDY received for the command is smaller than the size of the next frame to be transmitted, this state shall transition to the ST_ACT3:Receive state and wait for receipt of an XFER_RDY frame.

If the negotiated value in `FIRST BURST SIZE` field in the Disconnect-Reconnect mode page for the SCSI target device or the unused portion of the value in the `WRITE DATA LENGTH` field if an `XFER_RDY` frame has been received is larger than the size of the next frame to be transmitted, this state shall send a `Transmit Frame(Non-interlocked)` request to the port layer for each `DATA` frame. This state machine shall receive a confirmation of `Transmission Status(Frame Transmitted)` from the port layer for each frame before sending another `Transmit Frame(Non-interlocked)` request to the port layer.

If the confirmation for any frame is `Transmission Status` with any argument other than (`Frame Transmitted`), or `Close Connection` with any argument, `NAK received`, or `ACK/NAK timeout`, this state shall transition to the `ST_ACT5:Prepare_Response` state.

If confirmation from the port layer for all frames to be transmitted for the task or task management function is `Transmission Status(Frame Transmitted)` and an `ACK` confirmation for each transmitted frame has been received from the port layer, this state shall transition to the `ST_ACT5:Prepare_Response` state.

If a `Cancel` parameter is received from the `ST_AC` state machine this state shall send a `Cancel` parameter to the port layer and transition to the `ST_ACT5:Process_Response` state.

9.2.7.2.3.2 Transition `ST_ACT1:Send` to `ST_ACT2:Prepare_Send_Data`

This transition shall occur whenever the `ST_ACT1` state requires data to include in a frame.

9.2.7.2.3.3 Transition `ST_ACT1:Send` to `ST_ACT3:Receive`

This transition shall occur whenever the `ST_ACT1` state has sent a `Transfer Frame` request to the port layer and one of the following is true:

- a) A `Transmission Status(Frame Transmitted)` and an `ACK` confirmation have been received from the port layer for a frame containing a `COMMAND` information unit requesting data from the SCSI target device;
- b) A `Transmission Status(Frame Transmitted)` and an `ACK` confirmation have been received from the port layer for a frame containing a `COMMAND` information unit requesting to send data to the SCSI target device and the value in `FIRST BURST SIZE` field or the unused portion of the value from the `WRITE DATA LENGTH` field is smaller than the size of the next frame to be transmitted; or
- c) `Transmission Status(Frame Transmitted)` and `ACK` confirmations have been received from the port layer for all frames for a task or task management request and a `RESPONSE` frame is required from the SCSI target device.

9.2.7.2.3.4 Transition `ST_ACT1:Send` to `ST_ACT5:Process_Response`

This transition shall occur whenever an error has occurred during transmission of a frame or a `Cancel` parameter has been received from the `ST_AC` state machine.

9.2.7.2.4 `ST_ACT2:Prepare_Send_Data` state

9.2.7.2.4.1 `ST_ACT2:Prepare_Send_Data` state description

In this state the data required for inclusion in an SAS frame to be transmitted to the SCSI target port is prepared.

9.2.7.2.4.2 Transition `ST_ACT2: Prepare_Send_Data` to `ST_ACT1:Send`

This transition shall occur when data is available to be included in a frame for transmission.

9.2.7.2.5 ST_ACT3:Receive state

9.2.7.2.5.1 ST_ACT3:Receive state description

This state is entered when there are frames to be received from the SCSI target device. This state receives confirmation from the port layer that a frame has been received. The possible confirmations are:

- a) Frame Received(ACK/NAK Balanced);
- b) Frame Received(ACK/NAK Not Balanced); and
- c) Frame Received(Frame Failed).

If a Frame Received(ACK/NAK Balanced) confirmation is received, this state shall process the frame.

If a Frame Received(ACK/NAK Not Balanced) confirmation is received and the frame contains a DATA information unit, this state shall process the frame.

If a Frame Received(ACK/NAK Not Balanced) confirmation is received and the frame contains XFER_RDY or RESPONSE information unit, this state shall transition to the ST_ACT5:Process_Response state.

When processing the frame this state may verify that the INFORMATION UNIT TYPE, HASHED DESTINATION DEVICE NAME, HASHED SOURCE DEVICE NAME, COMMAND ID, and TAG fields contain the values associated with that task or task management function. If the state checks the values and they do not match, this state shall transition to the ST_ACT5:Process_Response state.

If the frame was received correctly, all of the checked values are correct, and the frame contains an XFER_RDY information unit, this state shall transition to the ST_ACT1:Send state.

If the frame was received correctly, all of the checked values are correct, and the frame contains a DATA information unit, this state shall transition to the ST_ACT4:Process_Receive_Data state.

If the frame was received correctly, all of the checked values are correct, and the frame contains a RESPONSE information unit, this state shall transition to the ST_ACT5:Process_Response state.

If a Cancel parameter is received from the ST_AC state machine this state transition to the ST_ACT5:Process_Response state.

9.2.7.2.5.2 Transition ST_ACT3:Receive to ST_ACT4:Process_Receive_Data

This transition shall occur whenever the a frame containing a DATA information unit has been received without error.

9.2.7.2.5.3 Transition ST_ACT3:Receive to ST_ACT1:Send

This transition shall occur whenever the a frame containing an XFER_RDY information unit has been received without error.

9.2.7.2.5.4 Transition ST_ACT3:Receive to ST_ACT5:Process_Response

This transition shall occur whenever there is service response data to send to the application layer.

9.2.7.2.6 ST_ACT4:Process_Receive_Data state

9.2.7.2.6.1 ST_ACT4:Process_Receive_Data state description

In this state data from a received SAS frame is processed.

9.2.7.2.6.2 Transition ST_ACT4:Process_Receive_Data to ST_ACT3:Receive

This transition shall occur after the data from a received SAS frame has been processed.

9.2.7.2.7 ST_ACT5:Process_Response_Data state**9.2.7.2.7.1 ST_ACT5:Process_Response_Data state description**

In this state a service response for command completion shall be sent to the application layer and the AC state machine. In this state a service response with intermediate response data for a linked command shall be sent to the application layer. After sending intermediate response data this state shall transition to the ST_ACT1:Send state.

9.2.7.2.7.2 Transition ST_ACT5:Process_Receive_Data to ST_ACT1:Send

This transition shall occur after intermediate response data for a linked command has been sent to the application layer.