Western Digital Corporation
5863 Rue Ferrari
San Jose, CA, 95138

To: **T10 SAS Protocol Working Group**
Contact: **Mark Evans**
Phone: 408.363.5257
Email: mark.evans@wdc.com
Date: **28 April 2008**

Subject: **SAS-2+, low power transceiver options, phy and link states**

1 Related documents

SAS-2r14 - Serial Attached SCSI - 2, revision 14
T10/08-015r3, SAS: Add low power transceiver options

2 Introduction

This proposal is a first pass at adding what is required in 08-015r3 for the phy and link layer state diagrams to include the low power transceiver options of partial and slumber modes for SAS.

3 Proposal

1) The following are the proposed changes to the SP state machine.

6.8.1 SP state machine overview

The SP state machine controls the phy reset sequence. This state machine consists of ~~three~~ four sets of states:

a) OOB sequence (OOB) states;
b) SAS speed negotiation (SAS) states~~.~~; and
c) SATA host emulation (SATA) states; and
d) SAS phy low power transceiver recovery (PM) states.

This state machine consists of the following states

…..

ae) SP29:SAS_Train (see 6.8.4.12);~~ and~~
af) SP30:SAS_TrainingDone (see 6.8.4.13)~~.~~;
ag) SP31:PM_Recovery_Start state (see 6.8.5.2);
ah) SP32:PM_ALIGN0 state (see 6.8.5.3).; and
ai) SP33:PM_ALIGN1 state (see 6.8.5.4).

…..

6.8.3.1 OOB sequence states overview

…..

[In figure 154 – SP (phy layer) state machine - OOB sequences states: add a transition arrow from SP0:OOB_COMINIT to SP31:PM_Recovery_Start.]

…..

6.8.3.2 SP0:OOB_COMINIT state

6.8.3.2.1 State description

This state is the initial state for this state machine.

Upon entry into this state, the phy shall:

a) set the COMWAKE_Received state machine variable to zero;
b) send a Stop DWS message to the SP_DWS state machine;
c) send a Phy Layer Not Ready confirmation to the link layer;
d) set the ATTACHED SATA DEVICE bit to zero in the SMP DISCOVER response (see 10.4.3.10);
e) if this state was entered due to power on, then set the ATTACHED SATA PORT SELECTOR bit to zero in the SMP DISCOVER response (see 10.4.3.10); and
f) if this state was not entered because of a Disable Phy request, then send a Transmit COMINIT message to the SP transmitter.

If this state was entered because of a Disable Phy request without a PhyPwrMgmt argument, then upon entry into this state, this state shall:

a) ignore COMINIT Detected messages until this state is re-entered due to a power on, hard reset, or Management Reset request; and
b) set the ResetStatus state machine variable to DISABLED.

If this state was entered because of a Disable Phy request with a PhyPwrMgmt argument, then upon entry into this state, this state shall:

a) not ignore COMINIT Detected messages; and
b) not change the ResetStatus state machine variable.

[Editor's note: the above changes allow this state to receive COMINITs initiating a link reset sequence that overrides the transceiver power management state.]

If this state was entered due to power on or hard reset, then upon entry into this state, this state shall set the ResetStatus state machine variable to UNKNOWN.

If this state was entered because of a Management Reset request, then upon entry into this state, this state shall:

a) if the ResetStatus state machine variable is not set to RESET_IN_PROGRESS, SPINUP_HOLD, G1, G2, or G3, then set the ResetStatus state machine variable to UNKNOWN; or
b) if the ResetStatus state machine variable is set to RESET_IN_PROGRESS, SPINUP_HOLD, G1, G2, or G3, then set the ResetStatus state machine variable to RESET_IN_PROGRESS.

[Editor's note: the above actions for a Management Reset request should allow this request to be used to initiate a link reset if the other actions for recovery from a low power transceiver mode fail.]

If this state was not entered due to a power on, hard reset, Disable Phy, or Management Reset request, then upon entry into this state, this state shall:

a) if the ResetStatus state machine variable is not set to PHY_RESET_PROBLEM, SPINUP_HOLD, or UNSUPPORTED_PHY_ATTACHED, then set the ResetStatus state machine variable to UNKNOWN; or
b) if the ResetStatus state machine variable is set to PHY_RESET_PROBLEM, SPINUP_HOLD, or UNSUPPORTED_PHY_ATTACHED, then not change the ResetStatus state machine variable.

If this state was not entered because of a Disable Phy request with a PhyPwrMgmt argument, the phy supports SATA port selectors, and this state receives a COMWAKE Detected message, then this state shall:

a) if the ResetStatus state machine variable is not set to PHY_RESET_PROBLEM, SPINUP_HOLD, or UNSUPPORTED_PHY_ATTACHED, then set the ResetStatus state machine variable to PORT_SELECTOR;
b) set the COMWAKE_Received state machine variable to one; and

c) if the ATTACHED SATA PORT SELECTOR bit is set to zero in the SMP DISCOVER response (see 10.4.3.10), then:
   A) set the ATTACHED SATA PORT SELECTOR bit to one in the SMP DISCOVER response; and
   B) send a SATA Port Selector Change confirmation to the link layer.

This state machine waits for receipt of a COMINIT Transmitted message and/or a COMINIT Detected message.

…..

6.8.3.2.4 Transition SP0:OOB_COMINIT to SP31:PM_Recovery_Start

This transition shall occur if this state was entered because of a Disable Phy request with a PhyPwrMgmt argument, and this state receives a COMWAKE Detected message.

…..

6.8.4.1 SAS speed negotiation states overview

…..

[In figure 155 – SP (phy layer) state machine - SAS speed negotiation states: add a transition arrow to SP15:SAS_PHY_Ready from SP33:SP_ALIGN1.]

…..

6.8.5 SAS phy low power transceiver recovery states

[Editor's note: this is a new clause with a new figure.]

6.8.5.1 SAS phy low power transceiver recovery states overview

Figure a shows the SAS phy low power transceiver recovery states. These states process the actions that return a phy from a low power state to participating in an operational logical link. These states are indicated by state names with a prefix of PM.

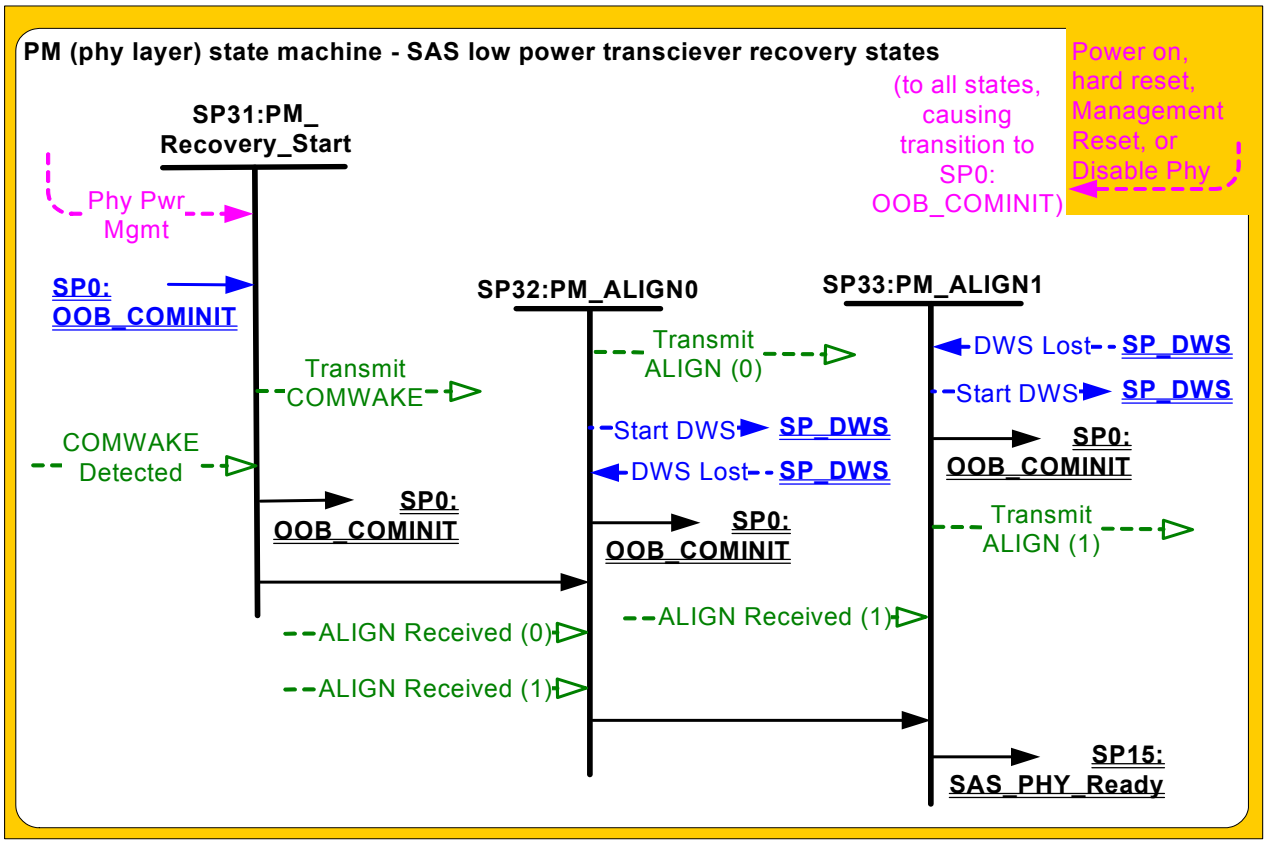

Figure a — PM (phy layer) state machine - SAS low power transceiver recovery states

6.8.5.2 SP31:PM_Recovery_Start state

6.8.4.2.1 State description

This is the initial state for the SAS low power transceiver negotiation sequence.

If this state was entered as the result of this state receiving a Phy Pwr Mgmt request, then this state shall transmit a COMWAKE.

6.8.4.2.2 Transition SP31:PM_Recovery_Start to SP0:OOB_COMINIT

This transition shall occur if this state transmits a COMWAKE and does not receive a COMWAKE within a COMWAKE timeout.

6.8.4.2.2 Transition SP31:PM_Recovery_Start to SP32:PM_ALIGN0

This transition shall occur after:

a) this state sends a Transmit COMWAKE message to the SP transmitter; or
b) this state receives a COMWAKE detected message from the SP receiver.

6.8.5.3 SP32:PM_ALIGN0 state

6.8.5.3.1 State description

Upon entry into this state, the phy shall:

a) initialize and start the SNLT timer;
b) send a Start DWS message to the SP_DWS state machine; and
c) repeatedly send Transmit ALIGN (0) messages to the SP transmitter.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

6.8.5.3.2 Transition SP32:PM_ALIGN0 state to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message, if this state does not send a Start DWS message.

Before the transition, this state shall set the ResetStatus state machine variable to UNKNOWN.

6.8.5.3.3 Transition SP32:PM_ALIGN0 to SP33:PM_ALIGN1

This transition shall occur if this state receives an ALIGN Received (0) message or an ALIGN Received (1) message before the SNLT timer expires.

6.8.5.4 SP33:PM_ALIGN1 state

6.8.5.4.1 State description

Upon entry into this state, the phy shall:

a) initialize and start the SNLT timer;
b) send a Start DWS message to the SP_DWS state machine; and
c) repeatedly send Transmit ALIGN (1) messages to the SP transmitter.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

6.8.5.4.2 Transition SP33:PM_ALIGN1 state to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message.

Before the transition, this state shall set the ResetStatus state machine variable to UNKNOWN.

6.8.5.4.3 Transition SP33:PM_ALIGN1 state to SP15:SAS_PHY_Ready

This transition shall occur if this state receives an ALIGN Received (1) message before the SNLT timer expires.

This indicates that the other phy has been able to achieve dword synchronization with the previously negotiated settings.

[Editor's note: after the transition, the SP15:SAS_PHY_Ready state sends a Phy Layer Ready (SAS) confirmation after establishing multiplexing, if multiplexing is enabled.]

.....

2) The following are the proposed changes to the SL state machine.

7.14.1 SL state machines overview

……

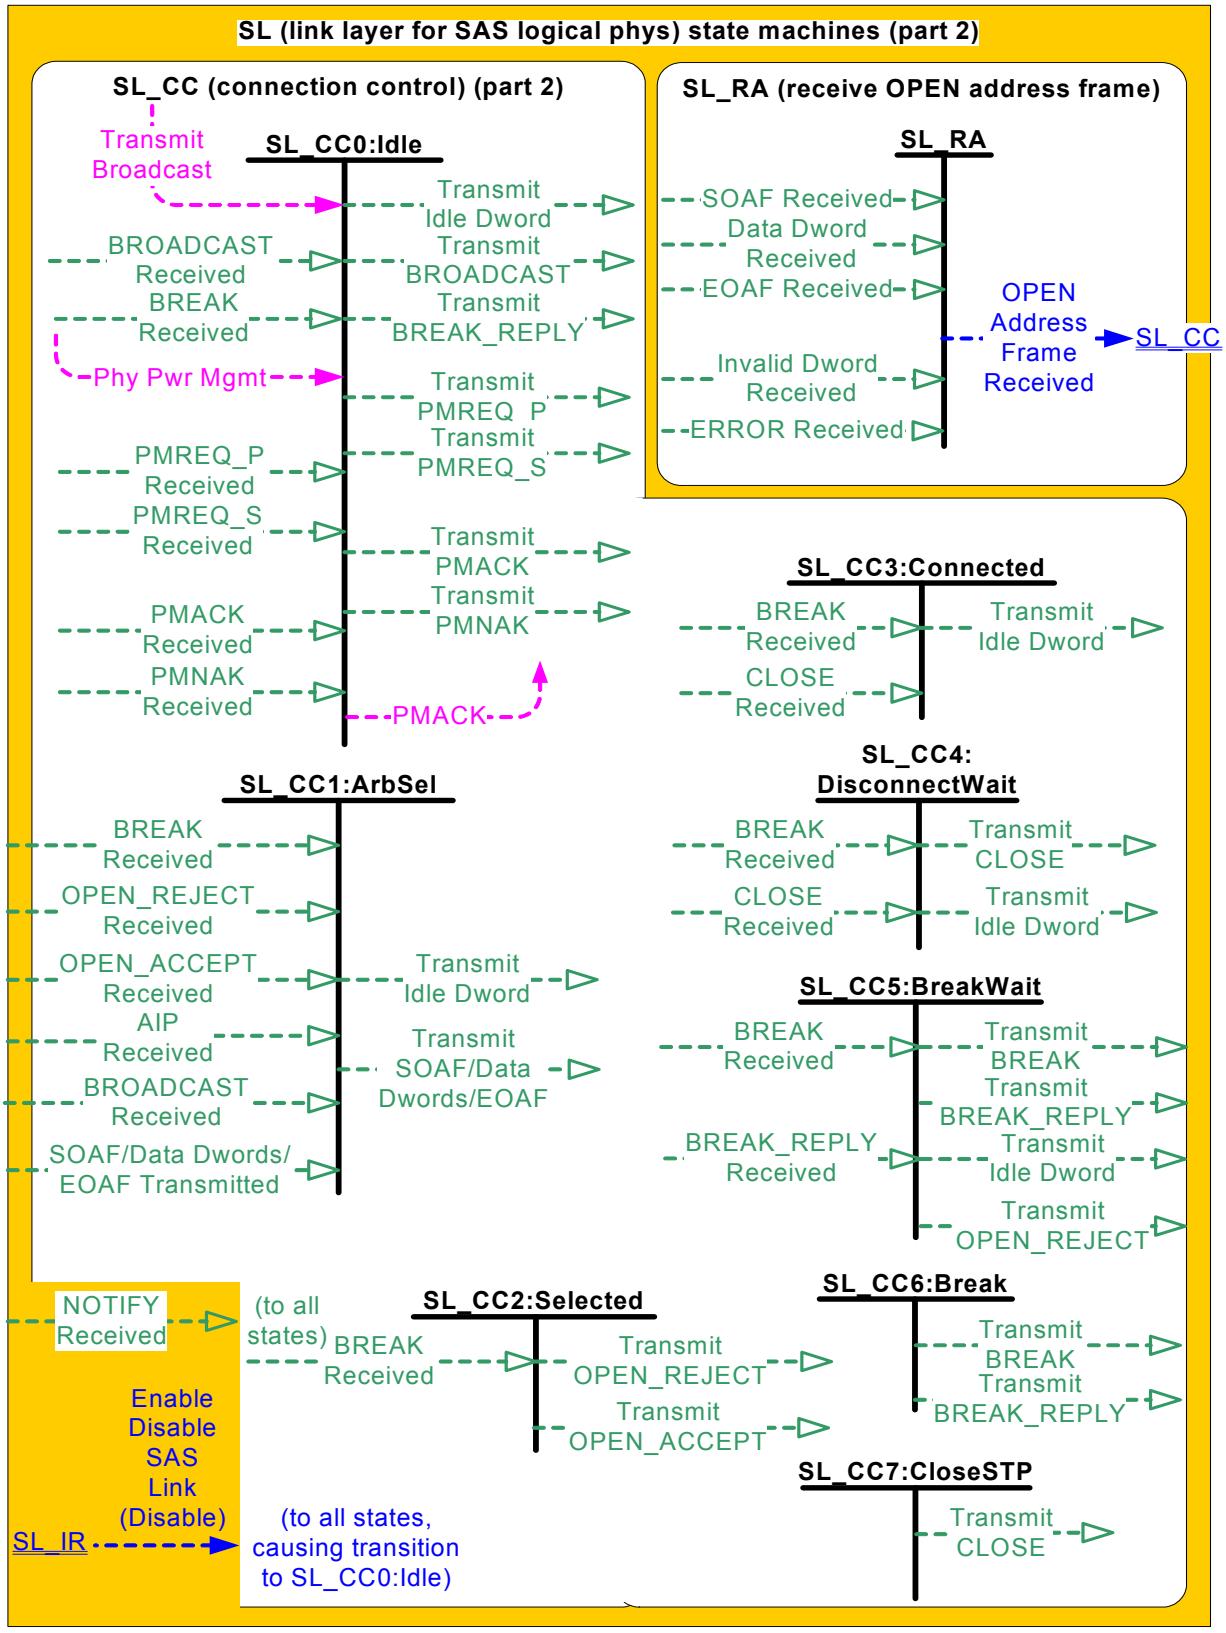

Figure 184 — SL (link layer for SAS logical phys) state machines (part 2) [replacement]

7.14.2 SL transmitter and receiver

The SL transmitter receives the following messages from the SL state machines specifying primitive sequences, frames, and dwords to transmit:

a) Transmit Idle Dword;
b) Transmit SOAF/Data Dwords/EOAF;
c) Transmit OPEN_ACCEPT;
d) Transmit OPEN_REJECT with an argument indicating the specific type (e.g., Transmit OPEN_REJECT (Retry));
e) Transmit BREAK;
f) Transmit BREAK_REPLY;
g) Transmit BROADCAST; and
h) Transmit CLOSE with an argument indicating the specific type (e.g., Transmit CLOSE (Normal)).

The SL transmitter receives Phy Pwr Mgmt Req requests with one of the following arguments:

a) Transmit PMREQ_P;
b) Transmit PMREQ_S;
c) Transmit PMACK;
d) Transmit PMNAK; or
e) Phy Power Recovery complete.

When the SL transmitter is requested to transmit a dword from any state within any of the SL state machines, it shall transmit that dword. If there are multiple requests to transmit, the following priority should be followed when selecting the dword to transmit:

1) BREAK_REPLY;
2) BREAK;
3) CLOSE;
4) OPEN_ACCEPT or OPEN_REJECT;
5) SOAF or data dword or EOAF; ~~then~~
6) PMREQ_P;
7) PMREQ_S;
8) PMACK;
9) PMNAK; then
10) idle dword.

When there is no outstanding message specifying a dword to transmit, and a PMACK has neither been received nor transmitted, the SL transmitter shall transmit idle dwords.

The SL transmitter sends Phy Pwr Mgmt Conf confirmations with one of the following arguments based on primitive sequences received from the SP_DWS receiver (see 6.9.2):

a) PMREQ_P received;
b) PMREQ_S received;
c) PMACK received;
d) PMNAK received; or
e) PMACK timeout.

If a PMACK has either been received or transmitted, then the SL transmitter shall transmit no dwords (i.e., the transmitter shall be at idle) (see the state description for SL_CC0:Idle).

The SL transmitter sends the following messages to the SL state machines based on dwords that have been transmitted:

a) SOAF/Data Dwords/EOAF Transmitted.

The SL receiver sends the following messages to the SL state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2):

a) SOAF Received;
b) Data Dword Received;
c) EOAF Received;

d) BROADCAST Received with an argument indicating the specific type (e.g., BROADCAST received (Change));
e) BREAK Received;
f) BREAK_REPLY Received;
g) OPEN_ACCEPT Received;
h) OPEN_REJECT Received with an argument indicating the specific type (e.g., OPEN_REJECT Received (No Destination));
i) AIP Received;
j) CLOSE Received with an argument indicating the specific type (e.g., CLOSE Received Normal));
k) NOTIFY Received (Power Loss Expected);
l) ERROR Received; and
m) Invalid Dword Received.

The SL receiver shall ignore all other dwords.

The SL transmitter relationship to other transmitters is defined in 4.3.2. The SL receiver relationship to other receivers is defined in 4.3.3.

…..

7.14.4.1 SL_CC0 state machine overview

…..

The SL_CC state machine shall maintain the timers listed in table 150.

Table 150 — SL_CC timers

Timer	Initial value
Open Timeout timer	1 ms
Close Timeout timer	1 ms
Break Timeout timer	1 ms
PMACK Received Timeout timer	1 ms

…..

7.14.4.2 SL_CC0:Idle state

7.14.4.2.1 State description

This state is the initial state and is the state that is used when there is no connection pending or established.

Upon entry into this state, this state shall send:

a) an Enable Disable SSP (Disable) message to the SSP link layer state machines;
b) an Enable Disable SMP (Disable) message to the SMP link layer state machines;
c) an Enable Disable STP (Disable) message to the STP link layer state machines; and
d) a Connection Closed (Transition to Idle) confirmation to the port layer.

If a PMACK has neither been received nor transmitted, then ~~T~~this state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL transmitter (see 7.4).

If a BROADCAST Received (Change) message, BROADCAST Received (Reserved Change 0) message, or BROADCAST Received (Reserved Change 1) message is received, this state shall send a Change Received confirmation to the management layer.

If a Transmit Broadcast request is received with any argument, this state shall send a Transmit BROADCAST message with the same argument to the SL transmitter.

If a BREAK Received message is received and the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), then this state shall send a Transmit BREAK_REPLY message to the SL transmitter.

After this state receives an Enable Disable SAS Link (Enable) confirmation, this state shall:

a) set the Reject SSP Opens state machine variable to a vendor-specific default value (i.e., YES or NO);
b) set the Reject SMP Opens state machine variable to a vendor-specific default value (i.e., YES or NO); and
c) set the Reject STP Opens state machine variable to a vendor-specific default value (i.e., YES or NO).

If this state receives a NOTIFY Received (Power Loss Expected) message and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this state shall send a Notify Received (Power Loss Expected) confirmation to the port layer.

If this state receives a Phy Pwr Mgmt Req request with a Transmit PMREQ_P argument, then this state shall:

a) transmit a PMREQ_P; and
b) initialize and start the PMACK Received timer.

If this state receives a Phy Pwr Mgmt Req request with a Transmit PMREQ_S argument, then this state shall:

a) transmit a PMREQ_P; and
b) initialize and start the PMACK Received timer.

If this state receives a Phy Pwr Mgmt Req request with a Transmit PMACK argument, then this state shall:

a) transmit a PMACK; and
b) stop sending Transmit Idle Dword messages to the SL transmitter (see 7.4).

If this state receives a Phy Pwr Mgmt Req request with a Transmit PMNAK argument, then this state shall transmit a PMNAK.

If this state receives a PMACK after transmitting a PMREQ_P or PMREQ_S, then this state shall:

a) stop the PMACK Received timer;
b) send a Phy Pwr Mgmt Conf confirmation with a PMACK Received argument; and
c) stop sending Transmit Idle Dword messages to the SL transmitter (see 7.4).

If this state receives a PMNAK after transmitting a PMREQ_P or PMREQ_S, then this state shall:

a) stop the PMACK Received timer; and
b) send a Phy Pwr Mgmt Conf confirmation with a PMNAK Received argument.

If the PMACK Received timer expires, then this state shall:

a) stop the PMACK Received timer; and
b) send a Phy Pwr Mgmt Conf confirmation with a PMACK Timeout argument.

If this state is sending Transmit Idle Dword messages to the SL transmitter (see 7.4) and receives a PMREQ_P, then this state shall send a Phy Pwr Mgmt Conf confirmation with a PMREQ_P Received argument.

If this state is sending Transmit Idle Dword messages to the SL transmitter (see 7.4) and receives a PMREQ_S, then this state shall send a Phy Pwr Mgmt Conf confirmation with a PMREQ_S Received argument.

If this state receives a Phy Pwr Mgmt Req request with a Phy Power Recovery Complete argument, then this state shall start sending Transmit Idle Dword messages to the SL transmitter (see 7.4).