

Draft Minutes
Automation/Drive Interface (ADI) Working Group
Ad Hoc Meeting
T10/04-034r0
12-13 January 2004
9:00 AM – 7:00 PM (12 January)
9:00 AM – 11:00 AM (13 January)

1. Introductions: Group

Paul Suhler called the meeting to order at 9:05 AM MST on 12 January 2004. He thanked Intel for hosting the meeting. He passed along Eric Hall's regrets for his absence. A table of the attendees appears at the end of these minutes.

2. Approval of the agenda: 04-035r0 Paul Suhler

Paul Suhler discussed the order of the discussion items.

He asked how much time we should allot to the ADC letter ballot comment resolution, 03-385r0. Rod Wideman suggested working through the questions from Microsoft. We agreed to limit ADC letter ballot comment resolution to one hour the morning of the 12th.

Paul Suhler reviewed the order of the discussion items. At Rod Wideman's request, we moved Kevin Butt's ADC/ADT comments by e-mail up to the second position.

Kevin Butt made a motion for acceptance of the modified agenda. Paul Entzel seconded the motion. The group passed the motion; no one objected or abstained.

3. Approval of previous meeting minutes: Paul Suhler

3-4 November 2003 meeting 03-378r0

17 November 2003 conference call 04-002r0

15 December 2003 conference call 04-017r0

5 January 2004 conference call 04-021r0

Paul Suhler requested comments for the minutes of the 3-4 November 2003 meeting, the 17 November 2003 teleconference, the 15 December 2003 teleconference, and the 5 January 2004 teleconference: 03-378r0, 04-002r0, 04-017r0, and 04-021r0 respectively.

Rod Wideman made a motion for acceptance of the minutes as published. Paul Entzel seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

4. Review of action items: Michael Banther

- a. Susan Gray will generate a proposal for link services error recovery. *Carryover*
- b. Susan Gray will revise 03-355r5 per discussion item (a). *Closed*

Kevin Butt made a motion to accept for incorporation into ADT the two changes proposed above by Paul Entzel. Rod Wideman seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

Kevin Butt pointed out an inconsistency in the text describing the Buffer Offset field in the SCSI Transfer Ready IU: the text requires in-order sending of data but it also allows requests for retransmission. Paul Entzel stated that requests for retransmission will occur due to a request by the primary port transport protocol making such a request in a bridged situation. He suggested removing the text that allows using the Buffer Offset of both the SCSI Transfer Ready and SCSI Data IU's to achieve retransmission. Lee Jesionowski quizzed Paul on the consequences of this change. Lee asked if we remove out-of-order capability, what happens if a port receives a Transfer Ready making an out-of-order request. We went down a rat hole discussing the possibility of a port having multiple SCSI Transfer Ready IU's outstanding.

After everyone became sufficiently confused, we returned to Kevin Butt's original question. Lee Jesionowski suggested having a new NAK status code so that the receiving transport layer can refuse an out-of-order request. Michael Banther suggested having the upper-layer abort the associated task. However Paul Entzel pointed out that the transport layer has to sniff the SCSI Transfer Ready and SCSI Data IU's to operate the buffer correctly because no signal goes to the upper-layer for every SCSI Data IU. The signal only goes to the upper-layer when all of the data for the associated Transfer Ready arrives. Michael asked, if we use a NAK to report out-of-order, how does the transport report the transmission error to the upper layer? Paul Entzel, Susan Gray, and Rod Wideman discussed this question with Michael and concluded that the transport would signal a transmission error by either failing to send a Data-In Delivered or Data-Out Received notification to the upper layer or by including a status code in these notifications and sending them with a non-Good status. The upper layer can then do what it chooses (probably abort the SCSI task).

Paul Entzel proposed replacing the second sentence in the second paragraph of ADT-r09a sub-clause 7.1.5, 'This field can be used to recover from an error detected in transmission by allowing the receiver of the data to request re-transmission of the previous burst of data' with 'A port that receives a SCSI Transfer Ready IU requesting data out of order shall send a NAK IU with a status code of INVALID BUFFER OFFSET.' He proposed replacing the second sentence of the second paragraph of ADT-r09a sub-clause 7.1.6, 'Data shall not be sent out of order, however, a port may re-send a data burst should it detect an error in transmission' with 'Data shall not be sent out of order. A port that receives a SCSI Data IU containing data out of order shall send a NAK IU with a status code of INVALID BUFFER OFFSET.' Finally, he proposed adding the INVALID BUFFER OFFSET NAK frame STATUS CODE as value 4Ah in ADT-r09a table 16.

Kevin Butt made a motion for incorporation of the proposal described above into ADT. Paul Entzel seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

Kevin Butt described his concern about the exchange lifetime of Encapsulated SCSI and Fast Access protocols. Paul Entzel replied that we have to define the exchange lifetimes to end on the sending of the response in order to avoid the corner case found by Paul Suhler where a SCSI Transfer Ready IU crosses a SCSI Response IU. Kevin asked, what happens if a port sends a SCSI Response IU and either receives a NAK or doesn't receive any acknowledgement after the exchange has closed? Paul Entzel explained that all of

the auto-sensing transports have this problem. Michael Banther pointed out that the current definition of the initiator's exchange lifetime allows the initiator to open a second exchange with the same exchange ID even though the target has to still hold on to the SCSI Response IU for the old exchange which has, according to the standard, closed. Paul Entzel suggested that we change the definition of when the exchange closes in the initiator so that it happens after the initiator sends an acknowledgement. Rod Wideman explored the possibility of having the exchange close in the initiator after sending an ACK but to keep it open if the initiator sends a NAK. Paul Suhler explored reworking the definition of exchange lifetimes. He asked if we need text stating that the receiver of a SCSI Transfer Ready IU for an exchange when the receiver has already sent a SCSI response IU for that same exchange simply acknowledges the SCSI Transfer Ready IU and then discards it. Paul Entzel stated that SAM-2 already covers this ground.

Paul Entzel proposed changing item 1) in the first numbered list in ADT-r09a sub-clause 7.1.7 from 'The port receives a SCSI Response IU for that exchange' to 'The port receives a SCSI Response IU for that exchange and sends an ACK IU for it.' He also proposed changing item 1) in the second numbered list in ADT-r09a sub-clause of 7.1.7 from, 'the port transmits a SCSI Response IU for that exchange' to 'the port transmits a SCSI Response IU for that exchange and receives an ACK IU for it.'

Kevin Butt made a motion for incorporation of the proposal stated above into ADT. Michael Banther seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

Kevin Butt pointed out that corresponding problems exist in sub-clauses to 7.2.7 Fast Access Exchange Lifetime of ADT-r09a.

Kevin Butt proposed that the ADT editor will extend the exchange lifetimes for Fast Access exchanges to include transmission or reception of ACK IU's as appropriate.

Kevin Butt made a motion to incorporate the proposal stated above into ADT. Rod Wideman seconded the motion.

Susan Gray proposed a friendly amendment to change "ACK IU's" to "ACK IU's or NAK IU's with PR bit set to zero".

Paul Suhler made a motion to accept the friendly amendment to the previous motion. Kevin Butt seconded this motion.

Rod Wideman commented that the friendly amendment should apply to the Encapsulated SCSI protocol as well.

Subsequently Kevin Butt withdrew his original motion (and the friendly amendment dies with it).

Kevin Butt tried again proposing that the ADT editor extend the exchange lifetimes for Fast Access exchanges to include transmission or reception of ACK IU's and NAK IU's with the PR bit set to zero as appropriate and that the statements in Encapsulated SCSI defining the exchange lifetimes as ending with the transmission or reception of an ACK IU be modified to end the exchange also upon transmission or reception of a NAK IU with the PR bit set to zero.

Kevin Butt made a motion to incorporate the proposal above into ADT. Rod Wideman seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

- c. ADT SCSI Application Diagrams 04-014r1 Susan Gray

Susan Gray recapped the changes in this revision.

Rod Wideman objected to the text in 8.2.12 being too long. Paul Entzel agreed that he will break the sentence into an unordered list as an editorial change.

Kevin Butt asked about the lack of all ACK IU's in the second and third diagrams. Several people explained our decision at the last teleconference to support Susan's choice to limit the number of ACK IU's shown. Regarding the diagram for Data Out, Rod Wideman pointed out that the different ordering of the SCSI Transfer Ready and SCSI Data IU's for the two SCSI data movements will only raise questions unless we add some explanatory text. After some haranguing for asking for a vague change, Rod requested adding a note to the text that states that the figure shows either ordering.

Lee Jesionowski asked if we need a similar change to the Send Data-In. We debated and concluded that no change is necessary.

Susan Gray made a motion for incorporation of 04-014r1 as revised into ADT. Kevin Butt seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

- d. ADT State Machine 03-369r1 Susan Gray

Susan Gray reviewed our progress on this proposal from the last teleconference. She also reviewed the Link Negotiation state diagram.

Susan noted that the transition from N4: Complete to N0: Idle needs the text 'with accept = 1' added.

Some concern arose regarding the transition from N2: Negotiating to N4: Complete. Paul Entzel pointed out that the sending of the Port Login IU with Accept equal to one is actually an action. It should be moved to the text. Susan agreed.

Rod Wideman raised questions about the state the machine is in upon reception of an ACK IU for a Port Login IU. Is it in N2: Negotiating or one of N3: Accept or N4: Complete? After some discussion, Rod concluded that the actual event for the transition from N3: Accept to N0: Idle is the reception of a Port Login IU with Accept equal one and unchanged parameters. Paul Entzel expanded this idea stating that we need to split N3: Accept into two states. One transition occurs upon reception of the Port Login IU with unchanged parameters and Accept equal to one. The second transition occurs upon completion of sending the ACK to the Port Login IU just received.

Someone asked what notation we should use to avoid letter ballot comments. Kevin Butt responded with a pointer to SAM-2. The example state diagram there and in SAM-3 includes the actions taken upon entry into a state immediately underneath the name of the state. We debated the wisdom of following SAM-2 and found two examples of other standards that do not follow it: SAS and SSC-2. Paul Entzel suggested noting in the Conventions that our state diagrams do not show actions upon entry underneath each state's name. Ralph Weber wandered in and made a point that we can use whatever state diagram notation we wish but we'd better include a section on it in the Conventions.

Discussion continued on what conventions we will follow. Susan Gray and Paul Entzel debated whether the sending of an IU should be counted as an event. We agreed to table this discussion until the morning of 13 January. At that time we will seek agreement on the conventions used in our state diagrams.

Upon resumption of this discussion, we debated the state diagram convention that we will use. Prompted by George Penokie, we examined the SAS state diagram conventions. He noted that the templates are available in Visio format.

Susan Gray described the subset of the SAS conventions she expects to use: state designator and state name, no labels on state-to-state transitions, use of messages to/from the state machine (with labels) to other state machines or entities, text for each state describing:

- The purpose of the state,
- The transitions from that state, and
- The messages to/from that state.

Having settled the notation question, we discussed the alterations to the Port state machine diagram based on these conventions. And then we moved on to considering the Transmission Error state machine and the Receiver Error state machine diagrams.

Michael Banther brought up a concern with the transition from R1 to R0. He would like to see the transition fire for any frame received that matches the expected frame number. The group debated the wisdom of this idea. Michael's concern centered on the ability of the transmitter to retry the sending of a frame an indefinite number of times and the mechanism that the receiver uses to detect when the transmitter stops retrying. After much discussion, we concluded that the existing text for transmitter retries specifies something that cannot actually happen.

Paul Entzel proposed modifying the second sentence of the first paragraph of 4.7.2.2 in ADT-r09a to remove the phrase, 'at least once and no more than four times.'

Paul Entzel made a motion to incorporate the proposal stated above into ADT. Michael Banther seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

We continued massaging the Receiver Error state machine diagrams. Upon completing that we moved on to consider the Negotiating state machine diagram, stopping when we ran out of meeting time.

Susan Gray will revise 03-369r1 based on comments received.

- e. ADT: Example Containing Figure 3 in ADT w/o the Figure

04-009r0

Kevin Butt

Kevin Butt introduced the proposal.

Michael Banther objected to the use of 'will' throughout. David Hawks pointed out that misplacement of 'both' in the sentence starting, 'in this case both the ADT port' David also pointed out a change from 'chose' to 'choose'. Paul Entzel raised a more fundamental concern with the wording of the first sentence. He and Michael suggested replacement text. As a group we worked through subsequent text 'improving' it.

Paul Entzel questioned the purpose of this text. Lee Jesionowski suggested limiting the text severely to state only that both ADT ports act as SCSI target/initiator ports. Seriously shortened text resulted.

Michael Banther made a motion for incorporation of 04-009r0 as revised into ADT. Paul Entzel seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

- f. ADT Informative Annex, Port Login Example 04-033r0 Paul Entzel

Paul Entzel introduced the proposal and explained his reasoning for the example proposed. He walked us through the text with various members throwing editorial suggestions his way. Kevin Butt suggested adding reference to the port and negotiation states. Paul Entzel agreed to do so after the state machines stabilize.

Kevin Butt made a motion for incorporation of 04-033r0 as revised into ADT. Lee Jesionowski seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

- g. ADT questions e-mail 6 Jan 04 Michael Banther

Michael Banther asked how a port should respond if it receives a SCSI Data IU when it expects a SCSI Transfer Ready IU and vice versa. Paul Entzel suggested that the receiver of an unexpected SCSI Data IU responds with a NAK with status code 43h, Invalid or Illegal IU received. The sender of the errant frame (the receiver of the NAK IU) reports the bad status to the upper layer which then probably aborts the associated task.

Rod Wideman raised concerns about using this same strategy for SCSI Transfer Ready IU's. Paul Entzel suggested that the receiver of a SCSI Transfer Ready should ACK it and performs no further action. The receiver of the SCSI Transfer Ready IU has received permission to send data when it has no data to send. However the receipt of the SCSI Transfer Ready IU has no other detrimental effect. The group agreed with these suggestions. Michael Banther will generate a proposal incorporating this change.

Michael raised the question about the correct behavior of a port that has received multiple non-acknowledged IU's and has NAK'ed the first one. Paul Entzel gave the example of out-of-resources. Paul argued that we need a way to re-synchronize the pipeline of data flowing across the interface. For every other case, the exchange is probably on its last legs. Paul Entzel will investigate further.

Michael Banther raised the mismatched length in the SCSI Data IU's Length field versus the Payload Size field. We made this choice to match other transport protocols and keep the fields on a four byte boundary. No change to the text will occur.

Michael Banther pointed out the circular reference in the handling of the Expected Frame Number Counter. After a bit of discussion, the group reached a consensus that the text has a problem.

Susan Gray and Paul Entzel proposed adding an Initiate Recovery IU to item 2) in the numbered list in 4.6.3 of ADT-r09. They also proposed modifying item 3) of this same list from, 'it shall be set using the rules in clause 4.7 when an Initiate Recovery IU is received or when ...' to 'it shall not be adjusted when'

Kevin Butt made a motion to incorporate the proposal stated above into ADT. Paul Entzel seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

6. **Unscheduled business:**a. **Buffer Offset discussion**

Lee Jesionowski

Lee Jesionowski pointed out the existence of the Incorrect Relative Offset Value in SCSI Data IU in the Response Code Values table (table 24 in ADT-r09). He stated that it's a holdover from earlier discussions that has been superseded. Lee proposed removing this code and renumbering the list in table 24.

Lee Jesionowski made a motion for incorporation of the proposal above into ADT. Kevin Butt seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously.

b. **Ordering of Initiate Recovery IU and outstanding NAK IU's**

Kevin Butt

Kevin Butt found a problem in sub-clause A.3 of ADT-09a. Figure A.2 shows the sending of an Initiate Recovery IU before receiving the outstanding NAK's from the other port due to frames sent after the frame received in error. In fact, in order to ensure that the sender of the Initiate Recovery IU doesn't mistake a NAK for a previous IU as being for the Initiate Recovery IU, the port must wait for all outstanding acknowledgements or an acknowledgement timeout before sending the Initiate Recovery IU.

Kevin Butt will generate a proposal to address this problem.

7. **Next meeting requirements:**

Paul Suhler

The group will hold teleconferences tentatively on 26 January and 9 and 23 February 2004, from 8:00 to 10:00 PST. We are soliciting hosts for those teleconferences.

The group will hold a meeting 8-9 March 2004 during T10 plenary week in Dana Point, CA. Subject to approval by T10, the meeting will begin on the 8th at 9:00 AM and conclude at 7:00 PM PST. The meeting will reconvene on the 9th at 9:00 AM and adjourn at 11:00 AM PST.

8. **Review new action items:**

Michael Banther

a. Rod Wideman will post 03-385r1.

b. Rod Wideman will send an e-mail to the T10 reflector that tags possibly controversial technical comments in 03-385r1.

c. Paul Entzel will incorporate the proposals contained in discussion item (b) into ADT.

d. Paul Entzel will incorporate the proposal contained in unscheduled business item (a) into ADT.

e. Susan Gray will revise 04-014r1 per discussion item (c).

f. Paul Entzel will incorporate 04-014r1 as revised into ADT.

g. Ralph Weber will send the SAM-3 state diagram conventions figure to Paul Entzel.

h. Kevin Butt will generate a proposal to address the problem described in unscheduled business item (b).

- i. Kevin Butt will post the revised 04-009r0 to the T10 web site.
 - j. Paul Entzel will incorporate 04-009r0 as revised into ADT.
 - k. Paul Entzel will revise 04-033r0 per discussion item (f).
 - l. Paul Entzel will incorporate 04-033r0 as revised into ADT.
 - m. Michael Banther will generate a proposal to deal with how a port responds if it receives a SCSI Data IU when it expects a SCSI Transfer Ready IU and vice versa (see discussion item [g]).
 - n. Paul Entzel will generate a proposal to deal with recoverable transport layer errors with PR equal zero, e.g., out of resources (see discussion item [g]).
 - o. Paul Entzel will incorporate the proposal described in discussion item (g) into ADT.
 - p. Paul Entzel will incorporate the proposal described in discussion item (d) into ADT.
 - q. Susan Gray will revise 03-369r1 per discussion item (d).
9. Adjournment: Group

Kevin Butt made a motion for adjournment. Susan Gray seconded the motion. In the absence of objections or abstentions, the group passed the motion unanimously. Paul Suhler adjourned the group at 10:53 AM MST on 13 January 2004.

Attendees:

Name	Organization	E-mail
Rod Wideman	ADIC	rod.wideman@adic.com
Paul Suhler	Certance	paul.a.suhler@certance.com
Ralph Weber	ENDL	roweber@ieee.org
Michael Banther	HP	michael.banther@hp.com
George Penokie	IBM	gop@us.ibm.com
Kevin Butt	IBM	kdbutt@us.ibm.com
Lee Jesionowski	IBM	ljesion@us.ibm.com
David Hawks	Iomega	hawks@iomega.com
Paul Entzel	Quantum	paul.entzel@quantum.com
Susan Gray	Quantum	susan.gray@quantum.com