IEEE 802.3 Interpretations Report

November 18th, 2004
San Antonio, TX

David Law
David_Law@3Com.com

IEEE-SA Standards Companion Text on interpretations

Interpretations are a unique form of commentary on the standard. They are not statements of what the standard should have done or meant to say. Interpretations cannot change the meaning of a standard as it currently stands. Even if the request points out an error in the standard, the interpretation cannot fix that error. The interpretation can suggest that this will be brought up for consideration in a revision or amendment (or, depending on the nature of the error, an errata sheet might be issued).

However, an interpretation has no authority to do any of this. It can only discuss, address, and clarify what the standard currently says. The challenge for the interpreters is to distinguish between their expertise on what 'should be,' their interests in what they 'would like the standard to be,' and what the standard says. Interpretations are often valuable, though, because the request will point out problems that might otherwise have gone unaddressed.

http://standards.ieee.org/guides/companion/part2.html#interpret

Interpretation Number: 1-11/04
Topic: 1Gb/s Pause resolution
Relevant Clause: 37
Classification:

Interpretation Request

We have found what we believe to be a small anomaly between "Table 37-4 Pause priority resolution" on page 85 of IEEE Std 802.3-2002 and "Table 37-2 Pause encoding" on page 81. I have scanned through the errata and the published interpretations, and I have found no references to this issue.

The problem is that the meaning of PAUSE = 1 and ASM_DIR = 1 in Table 37-2 appears to be somewhat ambiguous, and this leads to potentially conflicting interpretations that affect Table 37-4.

Table 37-2 says that "1 1" means "Both Symmetric PAUSE and Asymmetric PAUSE toward local device," but the text below the table clearly states that when ASM_DIR is set, the PAUSE bit merely indicates the direction in which PAUSE is supported on the link, and thus does not indicate support for symmetric PAUSE.

In other words, at least by the 37.2.1.4 text on page 81, "1 1" means "Asymmetric PAUSE toward local device."

With that interpretation, Table 37-4 becomes a problem. The problem is in the 6th and 9th entries in that table. These are reproduced in text form below.

Local Device		Link Partner			
PAUSE	ASM_DIR	PAUSE	ASM_DIR	Local Resolution	Link Partner Resolution
1	0	1	-	Enable PAUSE Transmit and Receive	Enable PAUSE Transmit and Receive
1	1	1	-	Enable PAUSE Transmit and Receive	Enable PAUSE Transmit and Receive

In the first case, this allows both "1 0 1 0" and "1 0 1 1" states to resolve with PAUSE in both directions. For "1 0 1 0," this makes sense, as both sides are advertising symmetric PAUSE capability. However, for "1 0 1 1," this would does not make sense. The local device is advertising symmetric PAUSE, but the link partner is saying that it supports only PAUSE receive, and not PAUSE transmit.

If the text below Table 37-2 were correct, I believe that "1 0 1 1" ought to resolve in a manner similar to "1 0 0 1" -- that is, if the link partner can support only asymmetric PAUSE, then none is possible.

The second case allows both "1 1 1 0" and "1 1 1 1." For the same reasons as above, I think that "1 1 1 0" would resolve as "Disable PAUSE Transmit and Receive." For "1 1 1 1," both sides are advertising the PAUSE receive capability, but neither is advertising PAUSE transmit. This would also be a "Disable PAUSE Transmit and Receive" result.

If, however, the text in 37.2.1.4 is misleading (and should instead indicate that all devices that support PAUSE receive *must* also support PAUSE transmit, and that the PAUSE bit does not just indicate direction but always signals the ability to transmit PAUSE frames), then both Table 37-2 and 37-4 are right.

My suspicion at this point is that the tables are right, the text is misleading, and that the ability to transmit PAUSE frames is the logical "OR" of the PAUSE and ASM_DIR bits -- when either is set, PAUSE transmit is available -- and that the ability to receive PAUSE frames is indicated by the PAUSE bit alone. Such an interpretation allows these two capabilities to function in a manner similar to other capability bits.

Interpretation Number 1-11/04 proposed response

Topic: 1Gb/s Pause resolution
Relevant Clause: 37

The requester is misinterpreting the text under Table 37-2. The text does not say that the Pause bit "merely“ indicates the direction in which PAUSE is supported.

The text states:

"The PAUSE bit indicates that the device is capable of providing the symmetric PAUSE functions as defined in Annex 31B. The ASM_DIR bit indicates that asymmetric PAUSE operation is supported. The value of the PAUSE bit when the ASM_DIR bit is set indicates the direction PAUSE frames are supported for flow across the link."

This last sentence is an indication of the "direction" of the asymmetric PAUSE "should" the devices resolve to "only" using asymmetric PAUSE. This does not restrict the advertisement. This does not say that both capabilities can't be advertised at the same time. This does not say that advertising asymmetric PAUSE negates the advertisement of symmetric PAUSE. This becomes even more clear when looking at Table 37-4.

Interpretation Number 1-11/04 proposed response (cont’)

Entry 4 in this table is a good example. The local device is advertising that it only supports asymmetric PAUSE toward the link partner - "01". Thus, when the link partner advertises both symmetric PAUSE and asymmetric PAUSE toward itself - "11", since the local device only supports asymmetric PAUSE then the link partner uses the asymmetric PAUSE portion of its advertisement and resolves to asymmetric PAUSE toward the link partner.

Entry 6 in this table is another example. The local device is advertising that it supports both symmetric PAUSE and asymmetric PAUSE toward itself - "11". Thus, when the link partner advertises either that it supports only symmetric PAUSE - "10" or both symmetric PAUSE and asymmetric PAUSE toward itself - "11", they can't resolve to asymmetric PAUSE since the direction they are both advertising is opposite to each other so they must resolve to symmetric PAUSE.

Other entries are resolved in a similar fashion.

Interpretation Number: 02-11/04
Topic: Transmitter electrical specifications
Relevant Clause: 39.3.1
Classification:

Interpretation Request

Can you please help us on the folowing 802.3 interpretation requests?

(1) Section 39.3.1 of the IEEE Std 802.3, 2002 edition.
In table 39-2, Transmitter output type is is (P)ECL. However, output driver is ac-coupled. Is it required by standard that transmitter DC common-mode voltage needs to meet (P)ECL?

(2) Section 39.3.1 of the IEEE Std 802.3, 2002 edition.
In Figure 39-2, is there any requirement on transmitter impedance or Transmitter return loss @ TP2?

39.3.1 Transmitter electrical specifications

The output driver is assumed to have output levels approximating those of Emitter Coupled Logic (ECL), as measured at TP1. The transmitter shall meet the specifications in Table 39–2.

Table 39–2—Transmitter characteristics at TP2

Description	Value	Unit
Type	(P)ECL	
Data rate	1000	Mb/s
Clock tolerance	±100	ppm
Nominal signalling speed	1250	MBd
Differential amplitude (p-p)		
Max (worst case p-p)	2000	mV
Min (opening)	1100	mV
Max (OFF)[a]	170	mV
Rise/Fall time (20-80%)		
maximum	327	ps
minimum	85	ps
Differential skew (max)	25	ps

[a]Examples of an OFF transmitter are no power supplied to the PMD and PMA transmit output being driven to a static state during loopback.

For all links, the output driver shall be ac-coupled to the jumper cable assembly through a transmission network, and have output levels, measured at the input to the jumper cable assembly (TP2), meeting the eye diagram requirements of Figure 39–3 and Figure 39–4, when terminated as shown in Figure 39–2. The symbols X1 and X2 in Figure 39–3 and Figure 39–4 are defined in Table 39–3.

Interpretation Number 2-11/04 proposed response

[1] The standard is unclear on this issue, and no distinction can be made between alternative implementations. This is being referred to the sponsor for possible action at the next revision.

[2] The standard does not speak to these issues, and as such no distinction can be made between alternative implementations. This is being referred to the sponsor for possible action at the next revision.

Interpretation Number: 03-11/04
Topic: Frame Reception
Relevant Clause: 4.2.9
Classification:

Interpretation Request

We have a question with regard to IEEE 802.3 Part 3: CSMD/CD Access Method and Physical Layer Specification (dated 2002). The question concerns the Frame Reception pseudo code provided in 4.2.9. The function RemovePad is copied below.

The basic question is whether a received frame with the following attributes should be discarded:

- lengthOrTypeParam is < maxValidFrame.
- Received frame is > minFrameSize.
- Received frame is <= maxValidFrame.
- Received frame contains padding.
- CRC is valid.

The "twist" in the above frame is that the transmitter has sent a frame with padding even though the clientDataSize is >

minFrameSize. In the function RemovePad(), validLength is set based on "Length/Type field matches the received clientDataSize". So it seems like such a frame is discarded. Is this correct?

```
function RemovePad(
        var lengthOrTypeParam:LengthOrTypeValue; dataParam:DataValue):DataValue;
begin
        if lengthOrTypeParam ≥ minTypeValue then
        begin
            validLength:= true; {Don't perform length checking for Type field interpretations}
            RemovePad := dataParam
        end
        else
        begin
            if lengthOrTypeParam ≤ maxValidFrame then
            begin
                validLength := {For length interpretations of the Length/Type field, check to determine if
                                value represented by Length/Type field matches the received
                                clientDataSize};
                if validLength then
                    RemovePad:={truncate the dataParam (when present) to value represented by
                                lengthOrTypeParam (in octets) and return the result}
                else
                    RemovePad:=dataParam
            end
```

Interpretation Number 3-11/04 proposed response

The standard is unclear on this issue, and no distinction can be made between alternative implementations based on this.

- Note however that the example frame in the request should not exist. This is due to the requirement in subclause 3.2.7 that states that ‘The length of PAD field required for MAC client data that is n octets long is max [0, minFrameSize – (8 × n + 2 × addressSize + 48)] bits.’ A frame in excess of minFrameSize should always have a PAD field of size 0.

Interpretation Number 4-11/04, Q1

Question 1: Minimum Fragment Size:

Can the last fragment of a Ethernet frame be smaller than 64 Bytes?

If not, what is the padding mechanism and the mechanism to strip off the extra padding at the receiver?

If it can, then the 64B/65B Transmit state machine (p. 364 Fig 61.18) is not equipped to handle a case where a fragment is shorter than 64 Bytes.

There is nothing specific in any subclause, the only reference how to handle the last fragment are with respect to multiples-of-four: 61.2.2.6 PME Aggregation transmit function restrictions (page 354):

> "c) The fragment size, not including PAF header, shall be a multiple of 4 octets except for the last fragment of a data frame."

61.2.2.3 PME Aggregation Transmit function

The PME Aggregation transmit functions uses the following algorithm:

a) Select an active PME (i.e., one with TC_link_state asserted, see 61.3.1) for the next transmission.

b) Select the number of octets to transmit on that PME (shall not be less than minFragmentSize nor greater than maxFragmentSize, see 61.2.2.6).

c) Increment by one (modulo 2^{14}) and set fragment sequence number in the Fragmentation Header. There is a single sequence number stream for each aggregation, not one per PME. It is this sequence number stream that the receiver uses for fragment reassembly.

d) Set the start-of-packet and end-of-packet bits in the Fragmentation Header as appropriate.

e) Transmit fragment to the TC sublayer.

Interpretation Number 4-11/04, Q1 proposed response

Unambiguous

The standard states in subclause 61.2.2.3 that the PME Aggregation transmit functions uses an algorithm that selects the number of octets to transmit on that PME for every fragment such that this number shall not be less than minFragmentSize.

Interpretation Number 4-11/04 Q2

Question 2: PCS TC_link_state handling:

With respect to: 61.2.1.1 MAC-PHY Rate Matching functions

"If the PAF is disabled or not present, transmit frames shall not be forwarded to the TC sublayer unless TC_link_state is true for the whole frame. If the PAF is enabled, transmit fragments shall not be forwarded from the PAF to a TC sublayer unless the TC_link_state value of that TC sublayer instance is true for the whole fragment.

NOTE?This implies that in the absence of an active PAF, frames being transmitted over the MII when TC_link_state becomes true are never forwarded to the TC sublayer. A frame being transmitted over the MII when TC_link_state becomes false is aborted."

When the PCS receives a TC_link_state = false from the TPS-TC receiver layer, it is supposed to drop the frame towards the TX direction (i.e., towards the TPS-TC sublayer). Why is that? And, what is sent out towards the MAC I/F?

.

Interpretation Number 4-11/04, Q2 proposed response

Request for consulting.

Interpretation Number 4-11/04 Q3

Question 3: Y-Byte issue:

In Fig 61-18, the states SYNC_IDLE, UPDATE_K, IDLE: it is implied by transmitZ(k, loop) that the Y-BYTE is only sent when the local 64B/65B receiver is not synchronized as indicated by TC_synchronized = FALSE. If TC_link_state is FALSE, but TC_synchronized is TRUE (implying that the remote receiver is out-of-sync) then only idle codewords are sent.

This is in direct conflict with the following paragraph, which makes sending of the Y-BYTE only dependent on TC_link_state: Which one is correct?

> 61.3.3.1 TC encapsulation and coding (the paragraph before table 61.11)
> "No new fragment shall be transmitted when TC_link_state = FALSE (TC_link_state is defined in 61.3.3.7).
> If a fragment is being transmitted when TC_link_state becomes false, the End of Frame codeword completing the fragment shall not contain an S symbol after the end of the fragment. If an Idle codeword is being transmitted when TC_link_state becomes false, it shall be completed with Z symbols only. After the completed End of Frame or Idle codeword, only Out-of-Sync Idle codewords shall be transmitted until TC_link_state becomes true again."

Interpretation Number 4-11/04, Q3 proposed response

Unambiguous

The standard states in paragraph 7 of subclause 1.2.1 ‘The state diagrams contain the authoritative statement of the functions they depict; when apparent conflicts between descriptive text and state diagrams arise, the state diagrams are to take precedence.’

The paragraph referenced is being referred to the sponsor for possible action at the next revision.

Interpretation Number 4-11/04 Q4

Question 4: Sync word TC_coding Error handling:

With respect to:

> 61.3.3.1 TC encapsulation and coding:
> "When any of the following events occur, signal TC_coding_error shall be asserted:
> a) An incorrect octet is received when a Sync Octet is expected."

According to the 64B/65B receiver synchronization FSM, there can be multiple wrong synchwords before the receiver is deemed out-of-sync. What's the rationale for asserting TC_coding_error (and therefore RX_ER) on the gamma I/F, which will lead to the whole frame being dropped?

Interpretation Number 4-11/04 Q4 proposed response

This request is being returned to you because the questions asked do not constitute a request for interpretation but instead a request for consultation. Generally, an interpretation request is submitted when the wording of a specific clause or portion of a standard is ambiguous or incomplete. The request should state the two or more possible interpretations or the lack of completeness of the text. While you referred to subclause 61.3.3.1, you have not indicated any problem with the text.

Interpretation Number: 5-11/04
Topic: Gate Processing ONU Activation State Diagram
Relevant Clause: 64
Classification:

Interpretation Request

Classification of a grant as a hidden grant or a back-to-back grant in state diagram shown in "Figure 64–29—Gate Processing ONU Activation State Diagram" is not very clear.

The transition from "check next grant" state to the "Back to back grant" state is shown when (apart from other conditions) *"(nextStopTime > stopTime)*".

NextStoptime is assigned as :-

nextStopTime = nextGrant.start + nextGrant.length - laserOffTime

And **stopTime** is assigned as :-

stopTime = currentGrant.start + currentGrant.length - laserOnTime - LaserOffTime - syncTime

In this case if :-

1) *NextStoptime* is < *currentGrant.start + currentGrant.length – LaserOffTime* **AND**
2) *nextStopTime > stopTime*

Then the grant will be treated as a back-to-back grant. But this grant will end before the current grant ends (*currentGrant.start + currentGrant.length > nextGrant.start + nextGrant.length)* Hence is a hidden grant.

Interpretation Number 5-11/04 proposed response

Concerns have been raised about this issue that are being referred to the sponsor for possible action at the next revision.

IEEE 802.3 Motion

IEEE 802.3 approves the proposed Interpretation responses to Interpretation requests 1-11/04 through 5-11/04 as presented without the need for a 30 day letter ballot.

M: W. Diab S: S. Carlson Tech 75%/~~Proc 50%~~

PASSED/FAILED

Y: 29 N: 0 A: 0