Accredited Standards Committee
X3, Information Processing Systems

Doc: X3T10.1/96a172r2
Date: November 13, 1996
Project: X3T10.1 / 1147D & 1051D
Ref Doc.: SSA-TL2 rev 3
Reply to: Adge Hawes

To: X3T10.1 Membership
From: Adge Hawes

Subject: Simplification of Extended Distance Option

This proposal was discussed and agreed at the October X3T10.1 working group.

BACKGROUND

The "Extended Distance Option" (EDO) included in TL2 rev 3 requires devices that have the capability of extra transmitter and receiver buffering to report their buffer capability in the parameter called WINDOW SIZE. This parameter, which usually equates to one less than the number of buffers installed in the transmitter and receiver, is picked up by the master, and modified during configuration so that ports on either end of a link match their buffering; the result is that transmitters can send more than one outstanding frame to a receiver, and therefore greater distances can be achieved at maximum bandwidth.

During configuration, matching pairs of ports at each end of a link are enabled for the extended distance option and the appropriate WINDOW SIZE is programmed. Extended Distance Option ports implement counters instead of bits for the WAITING FOR ACK, WAITING FOR RR, ACK PENDING and RR PENDING flags. Ports programmed with a window size greater than one assume that they have received that number of RR's and begin transmitting frames; they may then have WINDOW SIZE ACKnowledgements outstanding. If an error occurs, an extra byte appended to the Link Status Byte provides the extra Receive Sequence Counter information necessary to identify the failed frame of all those outstanding; all outstanding frames are then retransmitted.

The drawbacks with this scheme are threefold:

1. The Master must adjudicate the permissable WINDOW SIZE,
2. Only one WINDOW SIZE is reported by a port, meaning that both the transmitter and receiver of that port must have the specified amount of buffering. If one has more buffering than the other, then the smaller value must be reported,
3. A port is not allowed to dynamically change the amount of buffering it has, i.e. a 'buffer pool' can not be used.

The October 1996 X3T10.1 working group meeting agreed to attempt to simplify this scheme to remove the confusion surrounding the WINDOW SIZE parameter and the initial values of the flow control flags, as well as address the drawbacks listed above.

PROPOSAL INTENTION

This proposal removes the WINDOW SIZE parameter altogether, and instead allows an Extended Distance Option port to automatically identify how much buffering the port at the other end of a link has, and by implication, how many frames it may send at a time without ACKnowledgement.

Ports that are capable of the extended distance option have the following features:

1. Both the transmitter and receiver may have any number of (up to 255) frame buffers available for service at any time,
2. They implement the WAITING FOR ACK and RR PENDING flags as up/down counters with enough bits to count the number of transmit and receive buffers respectively available, and implement the. WAITING FOR RR and ACK PENDING flags as 8 bit up/down counters, regardless of the number of buffers available,

3. They provide an eight-bit RECEIVE SEQUENCE COUNTER (RSC) for transmission in an extended Link Status Byte, rather than the usual two-bit RSC,
4. They implement a 100µS ACK timeout instead of 25µS, unless told to reduce back to 25µS by the RACK bit in Configure Port,
5. They report the capability by means of the EDO bit in the QUERY PORT REPLY SMS (and also report whether the feature is currently enabled), and they provide an EDO ENABLE bit in the CONFIGURE PORT SMS to enable the option.
6. They implement slightly different rules than normal regarding what constitutes an unexpected RR and an outstanding ACK.

The Master's duty is simply to enable the extended distance option on both ports at either end of a link when they both report EDO capability. No WINDOW SIZE is programmed. If the Master determines that a port is EDO capable while the port at the other end of the link is not, then it does not enable the EDO; instead it reduces the ACK timeout in the EDO-capable port back to 25µS, by setting the RACK bit in CONFIGURE PORT.

Protocol

Once enabled, an EDO receiving port sends out as many RR-pairs as it has buffers available. These could be sent out contiguously, but it is more robust if consecutive RR-pairs have at least one intervening character (data, ACK, FLAG, ABORT, NUL, SAT, SAT'), providing that all other protocol rules are followed. Therefore a sequence of three RR characters should be treated as the same error as an isolated RR. A non-EDO port would, of course, send out only a single RR-pair, waiting for the control byte of the frame to be received before transmitting another.

The transmitting port receives and counts these RRs. When it is ready, it may send as many frames as it has counted RR-pairs, or as many frames as it has transmit buffers for (whichever is the smaller), without waiting for the first to be ACKnowledged. It counts the frames it sends and the ACKs it receives; starts the ACK timeout counter whenever it transmits the trailing FLAG of a frame, and stops it when it receives the next ACK. As each ACK is received, the next buffer in sequence that is waiting for acknowledgement discards its contents and becomes available for use again.

The receiving port checks frames and ACKs them in the order it received them. With several frame acknowledgements outstanding, it would be possible for the ACK-pairs to be sent without intervening characters. For the same reasons as above, ACK pairs should not be sent continuously without at least one other character (data, ACK, FLAG, ABORT, NUL, SAT, SAT') between one ACK pair and the next.

If the receiving port detects an error then it does not ACKnowledge the frame, nor does it acknowledge any subsequent frames that may be en route. This will trigger an ACK timeout in the transmitting port. Ports then exchange the extended RECEIVE SEQUENCE COUNTER in the Link Status Byte during ERP and use that to determine which, if any, frames need to be retransmitted.

Lost ACKs

If an ACK is lost (i.e. transmitted but not received), a port may associate a subsequent ACK with the frame whose ACK is lost. The problem will not be discovered until a quiet period, when the ACK timeout is triggered by the failure to receive the last outstanding ACK. Because all ACK transmission stops when a frame is bad, reception of an ACK must mean that either:

1. the next frame awaiting acknowledgement has been acknowledged, or
2. a subsequent frame was acknowledged, meaning that the ACK for the next frame was transmitted but not received.

Either way, the next frame awaiting acknowledgement was received correctly and may be discarded. In any case, ACKs are unlikely to be lost; the most likely corruption is that they get distorted into some other character, possibly invalid, and cause an error by another means.

EDO disabled

A port that is capable of Extended Distance operation but in which the option is not enabled (e.g. just after reset) behaves just like a non-EDO port (and follows the normal rules for Waiting for ACK/RR and ACK/RR pending), with the following exceptions:

1. It must implement the 100µS ACK timeout, unless explicitly told to reduce it to 25µS by the RACK bit in Configure Port. This is because it could be connected via a >2.4Km cable, along which a 25µS timeout would trigger even in normal operation.
2. It must be prepared to receive up to 255 RR-pairs without generating an 'unexpected RR' error, and must count them as if it was enabled for EDO. This is because it could be connected to an EDO port that has already had its EDO enabled; the Master has not yet got round to enabling this port's EDO.

In fact an EDO-capable port that receives more than one outstanding RR could naturally assume that it is talking to an EDO-enabled port; however, it must not perform as an EDO port itself until explicitly told to do so by the EDO ENABLE bit in Configure Port. This means that although this port must count the received RR's as for normal EDO operation, it must follow non-EDO frame transmission rules, i.e. only one outstanding ACK is allowed. Once its own EDO is enabled, it may send as many frames as it has counted RRs.

ACK timeout

A Master would normally leave ACK timeouts in EDO-enabled port pairs at 100µS, as the distance between such ports could exceed 2.4Km. However, a particular SSA implementation might have knowledge of the link distances involved by other means; if it was known that the EDO enabled links were less than 2.4Km then the Master could reduce each port's ACK timeout to 25µS by means of the RACK bit. It is therefore not an invalid combination to have EDO enabled and RACK set.

Rules for Extended Distance Option ports

The WAITING FOR ACK, WAITING FOR RR, ACK PENDING and RR PENDING flags are implemented for the Extended Distance Option as follows:

	FLAG	Initialized to:	Incremented when:	Decremented when:
Tx port	WAITING FOR RR	all 1's	Control field transmitted	RR-pair received
	WAITING FOR ACK	0	Trailing FLAG transmitted	ACK-pair received
Rx port	RR PENDING	number of receiver buffers available	Receiver buffer becomes available for use, e.g.: 1. Frame unloaded, or 2. New buffer brought in for service	1. RR-pair transmitted, or 2. Receiver buffer removed from service
	ACK PENDING	0	Trailing FLAG received	ACK-pair transmitted

A port may start sending a frame when either of the following conditions is satisfied:

a) the port is in the Ready state and its WAITING FOR RR FLAG is not all 1's (i.e. is less than 255),

b) the transmitted frame is a Control frame and the port is in the Enabled, Check or Ready states.

An RR pair shall be transmitted when the port is in the Ready state and the RR PENDING FLAG is greater than zero (although RR pairs shall not be transmitted contiguously - see above).

If a port's WAITING FOR ACK flag reaches the number of transmit buffers it had available when it finishes transmitting the trailing FLAG character of the last frame, then it shall send no more frames until either

a) the WAITING FOR ACK flag drops below that value (i.e. one or more ACKs is received), or

b) a new transmit buffer is brought into service and loaded.

Until another frame is sent the port shall transmit FLAGs. Alternatively, the port can hold off transmitting the trailing FLAG after the CRC of the last frame, sending NULs instead, and only transmitting the trailing FLAG

when either of the two events above occurs. The latter approach is similar to the non-EDO method of operation, however the former is recommended because:

a) the FLAG characters help maintain character synchronization at the receiver, and

b) the receiving port can get an earlier start on processing the last frame

When the ACK PENDING FLAG is greater than zero an ACK pair shall be transmitted as soon as possible.

Protocol Errors

The following changes in the protocol error apply to Extended Distance Option ports:

a) An RR character pair is only unexpected when WAITING FOR RR has reached zero and another RR occurs, i.e. more than 255 outstanding RRs are received. (Note that this applies even if EDO is disabled).

b) A third RR character in sequence shall be treated the same as an isolated RR character.

c) An ACK character pair is unexpected when the WAITING FOR ACK flag is zero.

d) A third ACK character in sequence shall be treated the same as an isolated ACK character.

Miscellany

It should be noted that a port that implements the EDO rules may enjoy performance benefits over longer distances even if it implements only two buffers. This is because the two buffers can send out both their RR-pairs earlier than a non-EDO port; the latter requires reception of the control field of a frame before it can send out the second RR pair. This ‘head start’ means that the second frame may be sent earlier. Also, the second frame may be completely transmitted (without holding off the trailing FLAG) while the first ACK is outstanding, allowing the receiver to process the incoming frame sooner. This is illustrated in Figure 1, which shows how the traffic flows on a long SSA link (where the propagation delay is of similar size to the frame time) with each type of protocol.

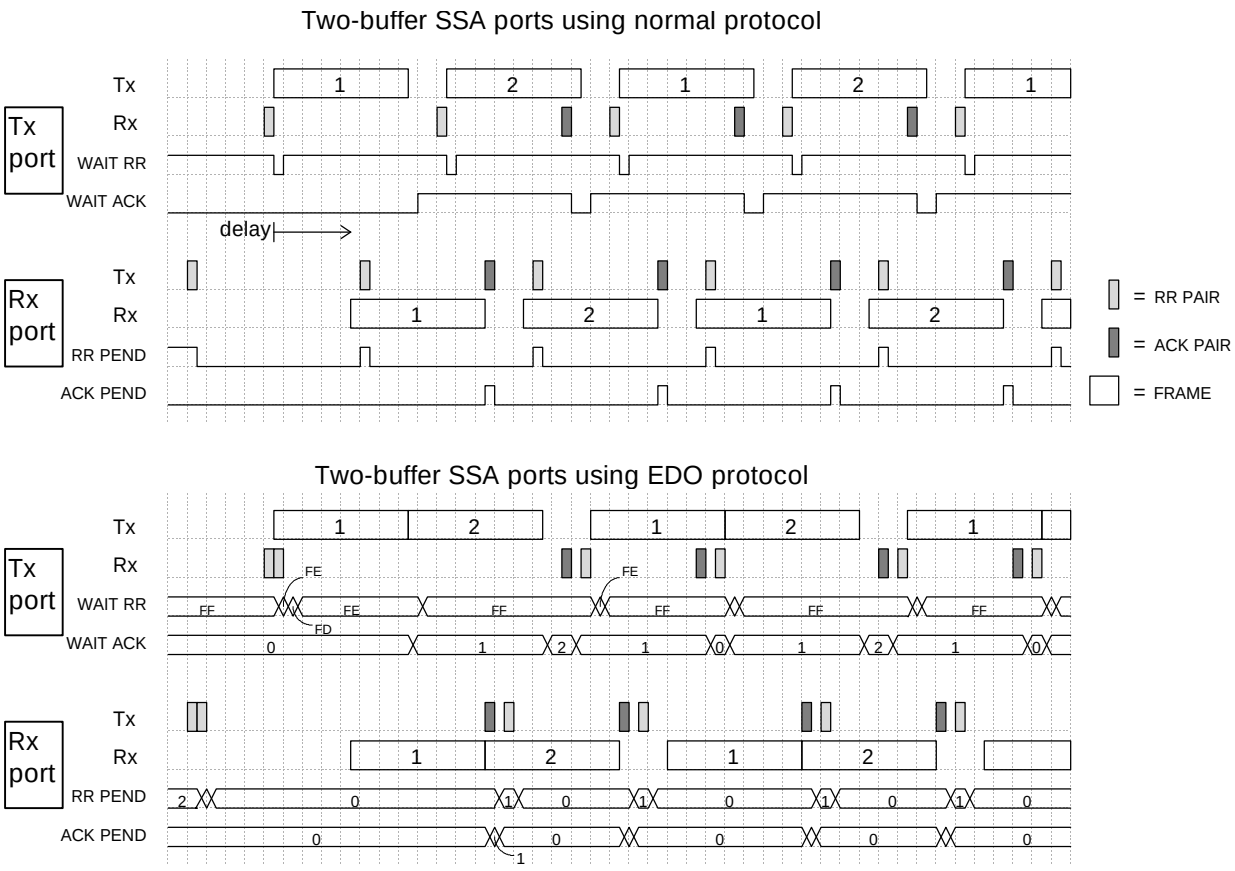

Figure 1. Comparison of normal and EDO protocols

Therefore, if the cost of the counters for the flags is acceptable, EDO capability may be the preferred design method for ports with even modest buffering.

The extra buffering capability provided by the extended distance option may be useful in circumstances that do not involve particularly long distances. The improvement is gained in cases where the propagation delay is similar to the frame time, so EDO operation could be useful when short frames are being transmitted over relatively short distances. Also, extra buffering can be employed to prevent blockage during peak traffic flow, as well as reducing the likelihood of deadlock.

Note: the exact changes to the TL2-r03 spec shall follow in revision 1 of this proposal.

PROPOSAL

Make the following changes to the clauses as indicated:

7.5 Permitted Special character sequences

... shall be at least one data character between an ABORT - FLAG pair and the last preceding FLAG character (e.g., FLAG, ABORT, FLAG is not allowed).

There shall be at least one other character between an RR pair and the last preceding RR pair, and between an ACK pair and the last preceding ACK pair (e.g., RR,RR,RR,RR is not allowed).

Violation of these rules while in the Ready state shall cause the receiver to invoke the Link ERP process and exit with a meaning code of PROTOCOL ERROR (see 11.1.1.6).

8.1.3 fsn field

... Since the TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER are independent on each link, the router in a dual port or switch node shall regenerate the FSN field as it forwards a frame.

Ports that implement the Extended Distance Option may have to implement additional bits for the TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER, to account for the higher number of unacknowledged frames permitted. A maximum of eight bits per counter are allowed. Only the two least significant bits are reported in the FSN field. The full value of the RECEIVE SEQUENCE COUNTER is exchanged in the EXT RSC field of the Extended Link Reset frame when the Extended Distance Option is enabled.

To facilitate future extensions, ...

8.3 ext rsc field

The EXT RSC field is a one byte field that immediately follows the STATUS field in an Extended Link Reset frame, and contains the extended RECEIVE SEQUENCE COUNTER. The Extended Link Reset frame is only used for link recovery when the Extended Distance Option is active for a link.

9.4.1 Acknowledgments

The link protocol requires a port to acknowledge every valid received frame, except for Total Reset, Reserved or Absolute Reset Control frames. A frame is valid if it does not contain any of the errors listed in 11.1.1. The destination port transmits an ACK character pair when it receives a valid frame. When the source port receives the ACK character pair, it shall discard the transmitted frame from its frame buffers.

This protocol guarantees that the transmitter always associates each ACK character pair unambiguously with the corresponding frame independently of propagation delays, the transmission speed and the frame length. When the Extended Distance Option is active, a transmitter cannot unambiguously associate an ACK character pair with the corresponding frame, and therefore must retransmit all unacknowledged frames during error recovery.

Each port has two associated flags, WAITING FOR ACK FLAG and ACK PENDING FLAG, that control acknowledgments. Ports that are not capable of the Extended Distance Option shall implement these flags as single bits. Ports that are capable of enabling the Extended Distance Option shall implement the WAITING FOR ACK FLAG and ACK PENDING FLAGs as counters, with enough bits (up to 8) to count the number of transmit and receive frame buffers, respectively, the port provides. The flags operate as follows:

I. When a port enters the Disabled state it sets the WAITING FOR ACK FLAG and ACK PENDING FLAG to zero.

II. A port shall increment the WAITING FOR ACK flag after it finishes transmitting the trailing FLAG character of any frame except Total Reset, Reserved or Absolute Reset Control frames. A port decrements the WAITING FOR ACK FLAG when it receives an ACK character pair. The local port shall discard the corresponding transmit data. The error conditions shall be handled as follows:

   A. After incrementing the WAITING FOR ACK FLAG a port shall wait to receive the ACK character pair for a minimum of 25 µs and a maximum of 50 µs before setting the ACK bit in the Link Status Byte and invoking the Link ERP process (this range allows a simple binary counter to be used for the timer) for the default SSA link behavior.

      Ports capable of running the Extended Distance Option shall use a 100 µs to 200 µs timer, and reset this timer each time an ACK character pair is received while the WAITING FOR ACK flag is

greater than zero. Extended distance nodes come out of reset with their ACK timeouts set between 100 µs and 200 µs. If an extended distance capable port is connected to a port not capable of extended distance, then the Master shall issue a CONFIGURE PORT SMS with the RACK bit set causing the extended distance capable node to reduce its ACK timeout to a minimum of 25 µs and a maximum of 50 µs. In order to guarantee that a port does not indicate an ACK TIME-OUT error for a frame, the delay from the port incrementing the WAITING FOR ACK FLAG to receiving the second ACK character shall be less than 25 µs for the normal operating mode, and 100 µs for the Extended Distance Option. This limits the maximum operating distance of the link (see 5)

B. If a non-Extended Distance enabled port's WAITING FOR ACK FLAG is equal to 1 when it finishes transmitting the CRC field of the next frame, then it shall not transmit the trailing FLAG character. Instead it sends NUL characters until either the ACK character pair is received or an ACK time-out occurs. If an ACK time-out occurs in this state, then the port shall abort the frame (see 7.2.6), set the ACK bit in the Link Status Byte and invoke the Link ERP process.

Ports that are enabled for the Extended Distance Option may transmit the trailing FLAG character of the next frame even when there are ACKs outstanding. The WAITING FOR ACK FLAG shall keep count of the frames that remain unacknowledged.

C. If a port receives an ACK character pair when the WAITING FOR ACK flag is zero, or it receives only a single ACK character, then it shall set the PROTOCOL ERROR value in the RECEIVER ERRORS field of the Link Status Byte and invoke the Link ERP process.

III. A port increments the ACK PENDING FLAG when it receives the trailing FLAG character of a valid frame in either of the following cases:

A. the port is in the Ready state and the frame is not a Total Reset, Reserved or Absolute Reset frame;

B. the port is in the Check state and the frame is a Link Reset frame.

IV. When the ACK PENDING FLAG is non-zero the port shall transmit an ACK character pair as soon as possible, but not immediately following the previous ACK character pair. The ACK PENDING flag is decremented when the ACK character pair has been transmitted.

9.4.2 Acknowledgment example

The timing requirements for acknowledgments and the use of NUL characters for non-Extended Distance operation is illustrated in Figure 11, where port 1 and port 2 are on opposite ends of a link (i.e., on different nodes). Port 1 finishes sending a first frame and then it immediately sends a second short frame. Due to the delay before port 2 acknowledges the first frame port 1 has to insert a NUL character in the second frame.

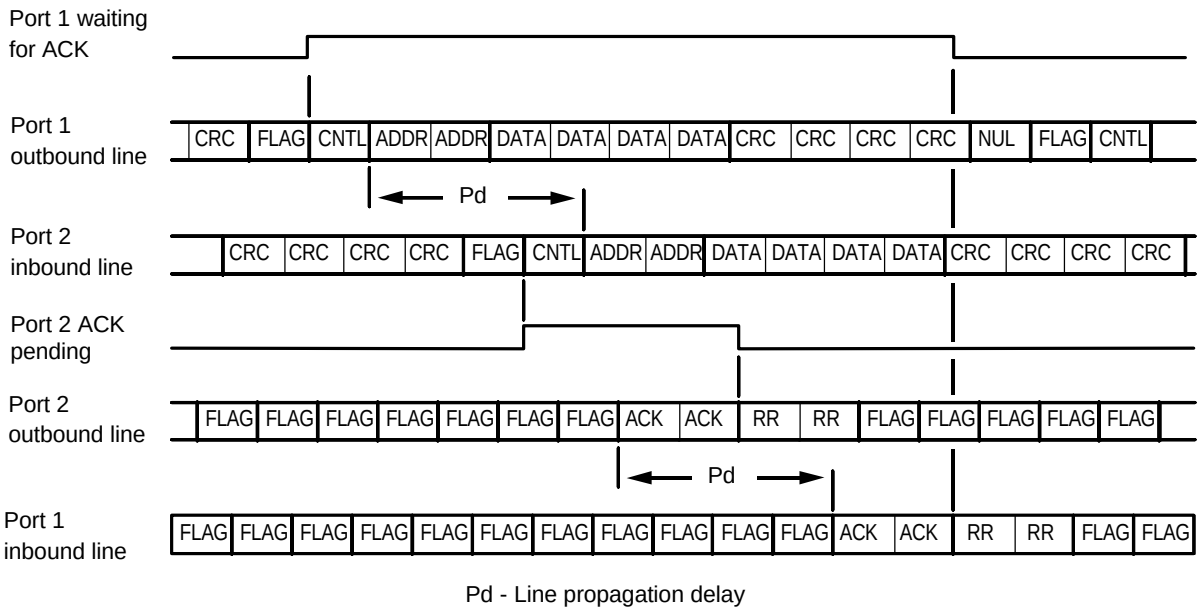

Figure 11 - Timing of acknowledgments (non extended distance operation)

9.4.3 Receiver Ready Pacing

Pacing ensures that a transmitting port does not overrun the available buffer space in the receiving port. The unit of pacing is a frame. The receiver shall have buffer space for at least 1 frame, although additional buffer space is required to sustain the full data rate of the link.

Each port has two associated flags, the WAITING FOR RR FLAG and the RR PENDING FLAG, that control pacing. Ports that are not capable of the Extended Distance Option shall implement these flags as single bits. Ports that are capable of enabling the Extended Distance Option shall implement the WAITING FOR RR FLAG as an 8-bit counter, and the RR PENDING FLAG as a counter with enough bits (up to 8) to count the number of receive frame buffers the port provides, and additionally a sign bit to allow the counter to go negative. The flags operate as follows:

I. When a port enters the Enabled state it sets the WAITING FOR RR FLAG and RR PENDING FLAG according to the following table, depending on whether the port is capable of Extended Distance operation, and whether or not the Extended Distance operation is enabled:

	For non Extended Distance capable port	**For Extended Distance capable port (Extendeded Distance operation not enabled)**	**For Extended Distance capable port (Extendeded Distance operation enabled)**
WAITING FOR RR FLAG	1b	FFh	FFh
RR PENDING FLAG	1b	1	Number of receiver frame buffers available

I. A port may start sending a frame when either of the following conditions is satisfied:

A. the port is in the Ready state and it's WAITING FOR RR FLAG is as shown in the following table:

	For non Extended Distance capable port	**For Extended Distance capable port**
WAITING FOR RR FLAG	0b	less than FFh;

A. the transmitted frame is a Control frame and the port is in the Enabled, Check or Ready states.

II. The WAITING FOR RR FLAG is incremented when a port transmits the CONTROL field of a Privileged, Reserved or Application frame and is decremented when the port receives an RR character pair.

III. When all of the following conditions are satisfied, a port shall transmit an RR character pair after the current character:

A. the port is in the Ready state and the RR PENDING FLAG is greater than zero;

B. buffer space is available to receive at least one frame on any active Channel, or shall be available before an incoming frame is received;

C. the immediately preceding character it transmitted was not an RR.

IV. The RR PENDING FLAG is controlled as follows, depending on whether the port is operating with Extended Distance Option enabled:

RR PENDING FLAG	**For non Extended Distance operation**	**For Extended Distance operation**
Incremented when	Port receives the CONTROL field of a Privileged, Reserved or Application frame	Receive buffer becomes available
Decremented when	Port transmits an RR character pair	Port transmits an RR character pair, or receive buffer is removed from service

When a port capable of Extended Distance operation has the option enabled, it shall increment the RR PENDING FLAG by the number of extra receive buffers it has, i.e. it shall consider those buffers to have become available for use.

If an Extended Distance operating port subsequently has the option disabled by a CONFIGURE PORT SMS, then it shall consider the extra receiver buffers to be removed from service, and decrement the RR PENDING FLAG by that amount. If this causes the RR PENDING FLAG to become negative, then the port shall not transmit any more RR character pairs until the RR PENDING FLAG is again greater than zero.

10.1.1 Frame buffers

Since the unit of pacing is a frame each port shall provide buffering to receive at least 1 frame. The buffer shall accommodate the maximum-length CONTROL, ADDRESS and DATA fields (i.e., 1+6+128 = 135 bytes total).

In certain occurrences of a transmission error the Link ERP process retransmits the last frames sent. In nodes implementing the Extended Distance Option, this can be up to 256 frames, limited by the 8 bit Extended FSN passed during Link ERP. Therefore a port shall also retain each frame it transmits until it receives the corresponding ACK character pair.

To achieve continuous communication at the full bandwidth of the link it is necessary for each port to have additional buffering. If a port has a pair of transmit buffers available and a pair of receive buffers to provide A/B buffering, one buffer of each pair is emptied/filled by the link while the other is filled/emptied by the application. If the Extended Distance Option is implemented, then the additional buffers permit the sending of additional frames while waiting for earlier transmitted frames to be acknowledged. The number of buffers (B) required to maintain full speed data transmission depends on the data rate(R_{DATA}), the round trip propagation delay of the transmission medium(T_{RT}), and the size of the frame(S_{FRAME}). The relationship is as follows:

$$B \geq (T_{RT} * R_{DATA} / S_{FRAME}) + 1$$

NOTE 1 - To maintain full speed (40 MB/s) at full size data frames (128 bytes of data) at the maximum distance (10 KM) require 30 buffers. Additional dedicated buffers are optional for SMSs to avoid impacting data transfers due to SMS processing delays.

The number and type of buffers in a port are vendor specific. However the following requirements shall be met to ensure performance and compatibility with other nodes in a Web:

a) All ports shall be able to sustain data transfers with 128 byte DATA fields at the full bandwidth of the link;

b) When there is no contention, the Router in dual port or switch nodes shall be able to forward the maximum length frames continuously;

c) All ports shall have an additional buffer to hold a Control frame when all other buffers are full.

10.1.2 transmit pointer and retry pointer

The transmit buffering shall be managed to allow correct recovery after a link error. In the case of recovered errors, the Link ERP process retransmits or discards the last one or two frames that were transmitted before an error. Each port shall maintain sufficient status to identify the buffers containing these frames and the order that they were transmitted.

For example, if there are *N* transmit buffers and the transmitter always accesses them in a cyclic sequence then the following two pointers provide sufficient information:

a) The TRANSMIT POINTER points to the buffer that is to be transmitted next. It is incremented by 1 modulo *N* each time the transmitter sends a trailing FLAG character without a preceding ABORT character;

b) The RETRY POINTER points to the next buffer to be acknowledged. It is incremented by 1 modulo *N* each time an ACK character pair is received while the WAITING FOR ACK FLAG is non-zero. Normally the RETRY POINTER follows the TRANSMIT POINTER closely but when an error occurs it is possible to lag by 1, or by the number of frame buffers available (Extended Distance operation).

The above scheme is an example used later in the description of the Link ERP process. However, sometimes it is impossible to guarantee cyclic use (e.g., if some buffers are reserved for SMSs). In this case, an alternate solution is to keep a log of the buffers that have been transmitted but not acknowledged, and the appropriate actions to be taken when TRANSMIT POINTER and RETRY POINTER are referenced.

10.1.7 waiting for rr flag and rr pending flag

The port shall maintain two flags, WAITING FOR RR FLAG and RR PENDING FLAG, to control the pacing of frames (see 9.4.3). If the Extended Distance Option is not supported, then these flags can be single bits. If the port implements the Extended Distance Option, the WAITING FOR RR FLAG shall be 8 bits, and the RR PENDING FLAG shall contain enough bits to count the number of receive buffers implemented, along with a sign bit.

10.1.8 waiting for ack flag and ack pending flag

The port maintains two flags, WAITING FOR ACK FLAG and ACK PENDING FLAG, to control the acknowledgment of frames (see 9.4.1). If the Extended Distance Option is not supported, then these flags can be single bits. If the port implements the Extended Distance Option, the WAITING FOR ACK FLAG shall contain enough bits to count the number of transmit buffers implemented. The ACK PENDING FLAG shall contain enough bits to count the number of receive buffers implemented.

11.1.1.6 Protocol Error

The PROTOCOL ERROR indicates the receiver has neither detected a LOSS OF SYNCHRONIZATION ERROR nor has detected a CODE VIOLATION ERROR, and a port receives an incorrect sequence of valid characters as follows:

a) a frame with less than 6 data characters between 2 FLAG characters (e.g., noise corrupting or manufacturing a FLAG character);

NOTE 2 - It is recommended that a short, and therefore invalid, Privileged, Reserved or Application frame with less than 7 data characters between 2 FLAG characters also be identified as a PROTOCOL ERROR.

b) a Privileged, Reserved or Application frame and no buffer is available, (i.e., when the RR PENDING FLAG is non-zero);

c) an unexpected RR character pair (i.e., when the WAITING FOR RR FLAG is zero);

d) an isolated RR character, or 3 RR characters. One half of a link hangs if an RR character pair is lost without any errors being detected, (e.g., if both RR characters are changed to FLAG characters while

the link is idle). This is extremely unlikely but the Master Alive process shall detect and correct this case;

e) an unexpected ACK character pair (i.e., when the WAITING FOR ACK FLAG is zero);

f) an isolated ACK character, or 3 ACK characters. If an ACK character pair is corrupted then the transmitter also detects an ACK Time-out error;

g) a NUL character with no intervening data character since the last FLAG character:

h) an ABORT character with no intervening data character since the last FLAG character:

i) an ABORT character that is not immediately followed by a FLAG character;

j) an isolated DIS character.

11.1.3 Link ERP process definition

To facilitate cross-referencing, the Link ERP exits are identified by an ALERT CODE value that corresponds to the ALERT CODE field used by the Asynchronous Alert process.

The Link ERP process shall not be invoked from within itself (i.e., it is non-recursive). The implementation shall protect against the Link ERP process looping if there is a permanent error.

NOTE 10 - The following is an example of one method that is used to protect against the Link ERP process from looping. Each invocation of the Link ERP process increments a RETRY COUNTER that is reset to zero periodically by a timer. If the number of retries in one period of the timer exceeds a maximum value then the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, it places the port in the Disabled state, and it exits with an ALERT CODE value of RETRY LIMIT EXCEEDED. This example also protects against excessive use of the ERP in the event of severe external noise.

The port that detects the error enters the Check state and invokes its Link ERP process. Any Link ERP exit shall terminate the Link ERP process and invoke the Asynchronous Alert process as defined in 11.3. The Link ERP process is as follows in order:

a) The Link Status Byte is built. If the port detects a line fault that persists for more than 1 ms the port enters the Disabled state, and Link ERP process exits with an ALERT CODE value of PERMANENT LINE FAULT.

b) If the port is not receiving characters, the port enters the Disabled state, and the Link ERP process exits with an ALERT CODE value of NO CHARACTERS RECEIVED.

c) If the port is detecting DIS characters the port enters the Disabled state and the Link ERP process exits with an ALERT CODE value of REMOTE PORT DISABLED.

d) Either a Link Reset frame (normal mode) or an Extended Link Reset frame (Extended Distance Option) is transmitted. If no acknowledgment is received the Link Reset or Extended Link Reset frame is retransmitted once. If no acknowledgment is received to the retransmitted Link Reset or Extended Link Reset frame, then the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, sets the port to the Disabled state, and exits with an ALERT CODE value of LINK RESET FAILED.

e) If the LINK RESET RECEIVED FLAG is cleared and a Link Reset frame has not been received within 5 ms, the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, places the port in the Disabled state, and exits with an ALERT CODE value of LINK RESET FAILED.

f) If a hardware error is detected, the Link ERP process sets the port to the Disabled state and exits with an ALERT CODE value of HARDWARE ERROR.

g) If FRAME REJECT ERROR is indicated by the port's Link Status Byte, the Link ERP process sets the port to the Disabled state, and exits with an ALERT CODE value of FRAME REJECT.

h) The number of outbound frames with an acknowledgment outstanding is calculated. When using strict cyclic buffers this is obtained as follows (see 10.1.2):

$$Q = (\text{TRANSMIT POINTER} - \text{RETRY POINTER}) \bmod N$$

where N is the number of buffers in use by the transmitter. Q is integral number of frames up to the number of buffers.

The number of outbound frames that the local port is expecting an acknowledgment for but have not been received by the remote port is also calculated, where local refers to the ports TRANSMIT SEQUENCE COUNTER and remote refers to the RSC field received in the Link Status Byte.

P = (local TRANSMIT SEQUENCE COUNTER - remote RSC) mod 4 (or mod 256 for Extended Distance operation)

where *P* is less than or equal to *Q*.

If either *Q* > number of buffers or *P* > Q then the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, sets the port to the Disabled state and it exits with an ALERT CODE value of INVALID RETRY STATUS.

i) The Link ERP process arranges to resend the number of frames indicated by subtracting *P* from its TRANSMIT POINTER, modulo N. Those outbound buffers that do not need to be retransmitted shall be discarded using the following algorithm.

```
Do while RETRY POINTER ≠ TRANSMIT POINTER;
      Discard buffer at RETRY POINTER;
      Increment RETRY POINTER modulo N;
End;
```

The frame buffers and TRANSMIT POINTER are set to retransmit the proper number of frames after the Link ERP re-establishes communication following the port being set to the Disabled state.

j) If the port has received a frame containing any of the errors listed in Table 18, then the appropriate inbound buffer contents shall be discarded. The ERP places the port in the Disabled state. The ERP waits until the receiver detects DIS characters. This is required to synchronize the local and remote ports Link ERP process's and prevent the transmitter sending RR while the remote port is still in the Check state.

If the receiver does not detect DIS characters within 5 ms after the local port is disabled then the ERP exits with an ALERT CODE value of TIME-OUT WAITING FOR DISABLED STATE.

k) The port is set to the Enabled state.

If the receiver detects a FLAG character within 5 ms of placing the port in the Enabled state, the port is set to the Ready state, otherwise the Link ERP process exits with an meaning code of TIME-OUT WAITING FOR READY STATE.

l) The ERP terminates successfully.

11.1.3.2. ERP exits

If the Link ERP process exits unsuccessfully then each node shall take the following actions in order:

a) Clear the port's OPERATIONAL FLAG.

b) Disable the Extended Distance Option and set the ACK timeout to between 100 and 200 μsec.

c) Set the port to Privileged mode. The transmitter discards Application frames. This avoids frames backing up and blocking other traffic. It also prevents any further data from being transferred through the affected port.

d) Invoke the Asynchronous Alert process (see 11.3) with the ALERT CODE specified in the Link ERP exit.

e) Restart the affected port, as described in 10.3.

12.2.3 CONFIGURE PORT SMS

The CONFIGURE PORT SMS defined in Table 29 is sent by the Master to each port on every other node in the Web during the configuration process. The Master also issues a CONFIGURE PORT SMS to return a port to Normal mode after it has processed an Asynchronous Alert that placed the port in the Privileged mode. In both cases the destination node returns a RESPONSE SMS.

Table 29 - CONFIGURE PORT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (02h)							
1	PORT							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	reserved							
9	A QUOTA							
10	B QUOTA							
11	EUDC	REFLECT	MODE		reserved		EDO	RACK
12	reserved							
13	ALARM THRESHOLD							
14	ALARM THRESHOLD							
14	NEGOTIATE SPEED							
15	NEGOTIATE SPEED							

A CONFIGURE PORT SMS specifies the return path and the tag to be used when a node sends the Master an ASYNC ALERT SMS. It also establishes the operating parameters of the port.

The PORT field contains an unsigned integer to identify the affected port.

The TAG field is returned in the RESPONSE SMS. The same tag is also used in an ASYNC ALERT SMS if the node subsequently reports an Asynchronous Alert.

The TAG value is assigned by the Master and shall be unique among the TAG values used by other ports in the Web as indicated by the Asynchronous Alert table (see Table 13). The receipt of the resulting RESPONSE SMS shall release the TAG value for use by SMSs other than the CONFIGURE PORT SMS, and the tag value shall no longer be considered active.

The RETURN PATH field specifies the path component that shall be placed in the address field of the resulting RESPONSE SMS and any subsequent Asynchronous Alert.

The A QUOTA field specifies the frame quota that a node may originate before it is satisfied. The value of the A QUOTA field shall be greater than zero.

The B QUOTA field specifies the maximum frame quota that a node may originate for each rotation of the SAT token. The value of the B QUOTA field shall be greater than or equal to the value of the A QUOTA field. If the B QUOTA value is less than the A QUOTA value, then the CONFIGURE PORT SMS shall be ignored and a RESPONSE SMS generated to the Configutor node issuing the CONFIGURE PORT SMS with the appropriate TAG value and a RETURN CODE of INVALID FIELD.

The EUDC bit (Enable User Defined Characters) specifies how the port handles User Defined characters. If the EUDC bit is cleared, the port shall not transmit any User Defined characters and shall discard received User Defined characters. If the EUDC bit is set, the port may transmit User Defined characters, and shall forward User Defined characters.

The REFLECT bit is one of the controls affecting the REFLECTION FLAG node function (see 10.2.5). If the node is a dual port node and the REFLECT bit is set, the REFLECTION FLAG is set for the node. If the node is a dual port node with the OPERATIONAL FLAG set for both ports and the REFLECT bit is cleared, the REFLECTION FLAG is cleared for the node. The REFLECT bit shall be ignored for single port nodes or switch nodes.

The MODE field selects the port mode as defined in Table 30.

Table 30 - MODE field values

Value	Description
00b	NO CHANGE TO THE CURRENT MODE.
01b	SET WRAP MODE.
10b	SET NORMAL MODE.
11b	SET PRIVILEGED MODE.

If a node receives a CONFIGURE PORT SMS with MODE field value of NO CHANGE TO THE CURRENT MODE and the specified port has an Asynchronous Alert pending then it shall send all pending ASYNC ALERT to the Master before sending the RESPONSE SMS for the CONFIGURE PORT. A MODE field value of SET WRAP MODE sets the Wrap mode and is only valid if the node received the CONFIGURE PORT SMS via a different port than that being configured. Otherwise the node returns a RESPONSE SMS with RETURN CODE VALUE of INVALID FIELD and the port mode is unchanged. When a port enters the Wrap mode, the port also becomes not Operational (see 10.1.4.3). Once a port has been placed in Wrap mode, the Master causes the node to exit the Wrap mode on the port by sending a CONFIGURE PORT SMS via a different port with a mode field value of SET PRIVILEGED MODE. The port shall exit the Wrap mode and invoke the Beginning Communication process (see 9.4). If a port currently configured in Wrap mode receives a CONFIGURE PORT SMS via a different port with a mode field value of SET NORMAL MODE, the node returns with a RESPONSE SMS with RETURN CODE VALUE of INVALID FIELD and the port remains in the Wrap mode.

the EDO (Extended Distance Option) bit causes the port to enable Extended Distance operation (if supported).

The RACK bit (Reduce ACK timeout) causes a port to set the ACK timeout to between 25 µs and 50 µs.

The ALARM THRESHOLD field specifies the number of times the Link ERP process is invoked before the Asynchronous Alert process is invoked with an ALERT CODE value of ALARM THRESHOLD EXCEEDED. An ALARM THRESHOLD field set to zero disables the alarm threshold function. When the LINK ERP COUNTER = ALARM THRESHOLD value the node shall issue an Asynchronous Alert. Subsequent increments to the LINK ERP COUNTER above the ALARM THRESHOLD value shall not invoke the Asynchronous Alert process. If a valid QUERY PORT SMS is received with the CLE flag set, the specified port shall set the LINK ERP COUNTER to zero and thereby restart the alarm threshold function (see 12.2.9).

The NEGOTIATE SPEED field defined in Table 30 indicates what speeds the port should use when speed negotiating during the Beginning Communication process (see 9.3). The changes to speed negotiation caused by the state of the NEG40 and NEG20 flags is defined in Table 31. If a port is instructed to operate at only 20 MB/sec or to operate at only 40 MB/sec, then the node behaves as a single speed device (see 9.4.1).

Table 30 - NEGOTIATE SPEED field

Byte	Bit 7	6	5	4	3	2	1	Bit 0
14	reserved							
15	reserved						NEG40	NEG20

Table 31 - NEGOTIATE SPEED field effects on speed negotiation

NEG40	NEG20	Speeds to be negotiated
Cleared	Cleared	No change to speed negotiation
Cleared	Set	Speed negotiate only at 20 MB/sec
Set	Cleared	Speed negotiate only at 40 MB/sec
Set	Set	Speed negotiate at both 40 MB/sec and at 20 MB/sec

A port only saves the parameters of the most-recent CONFIGURE PORT that it has received.

The acknowledgment for the CONFIGURE PORT SMS is the receipt of a RESPONSE SMS with the same TAG value as was used for the CONFIGURE PORT SMS. The RETURN CODE field of the RESPONSE SMS shall indicate whether or not the CONFIGURE PORT was processed successfully (see 12.2.16).

Following the generation of a RESPONSE SMS for the CONFIGURE PORT SMS, a port that is running at 20 M/sec but is capable of running at both 40 MB/sec and 20 MB/sec shall invoke the Asynchronous alert process with an ALERT CODE value of PORT OPERATING AT SLOWER THAN OPTIMAL SPEED.

12.2.7 QUERY NODE REPLY SMS

The QUERY NODE REPLY SMS as defined in Table 35 is returned when a QUERY NODE SMS is received. The QUERY NODE REPLY SMS is returned on the same port that received the corresponding QUERY NODE SMS.

Table 35 - QUERY NODE REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (01h)							
1	PORT							
2	TAG							
3	TAG							
4	UPPER LEVEL PROTOCOL							
5	ITF	MASTER PRIORITY			reserved			
6	TOTAL OTHER PORTS							
7	SSA-TL VERSION (10h)							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							
16	RETURN PATH ID							
17	RETURN PATH ID							
18	RETURN PATH ID							
19	RETURN PATH ID							
20	P1O	P2O	LONG	CM	reserved			

The PORT field indicates the number of the port currently being used.

The TAG field is copied from the QUERY NODE SMS. It identifies the QUERY NODE SMS that this reply is being generated for.

The UPPER LEVEL PROTOCOL field identifies the upper-level protocol that the node shall respond to. The UPPER LEVEL PROTOCOL field shall contain a value from Table 36.

Table 36 - Upper-level protocol code values

ULP code	Protocol	Notes
00h	USE THE QUERY PROTOCOL SMS TO LIST UPPER LEVEL PROTOCOLS SUPPORTED	1
01h	SHALL RESPOND TO NO UPPER-LEVEL PROTOCOL	1
02h	Vendor specific	1,2
03h-7Fh	reserved	
80h	SSA-IA / 95SP	1,2
81h	reserved	
82h	SSA-S2P	1,2
83h	SSA-S3P	1,2
84h-FAh	reserved	
FBh-FFh	Vendor specific	2

NOTES
1) Valid in QUERY NODE REPLY.
2) Valid in QUERY PROTOCOL REPLY

The Configutor table full (ITF) bit is set when there is no space left in the Configutor table to make an entry for the Configutor node that sent QUERY NODE SMS.

The MASTER PRIORITY field defines the priority of the node for becoming the Web Master. A value of 000b indicates that the node is a Responder node and is not capable of functioning as a Master. Any value greater than 000b indicates the node's priority for becoming the Master. A value of 001b is the lowest priority and 111b the highest

NOTE 11 - A Configutor node optionally fixes its MASTER PRIORITY at 4 (the default priority) or provides some mechanism outside of SSA to change its MASTER PRIORITY dynamically.

The TOTAL OTHER PORTS field contains a value that is one less than the number of ports implemented. If this value exceeds two, the QUERY SWITCH SMS is used to retrieve a port mask for the switch.

NOTE 12 - The TOTAL OTHER PORTS field value shall be 0 (single port), 1 (dual port) or an odd number greater than 1 (switch) as per clause 9.2.

The SSA-TL field defined in Table 34 identifies the version of SSA-TL being used by the sender. If multiple levels of SSA-TL are supported, then highest value shall be reported that is equal to or less than the SSA-TL field in the associated QUERY NODE SMS

The UNIQUE ID field contains the node's Unique ID.

The RETURN PATH ID contains a value created by the node and returned to the Configutor node. This field shall be used by the Configutor node in any future application SMSs that are utilizing the same path as used by the QUERY NODE SMS that caused this QUERY NODE REPLY SMS to be generated.

The P1O bit is only valid for single port and dual port nodes, Switch nodes shall clear the P1O bit. For single port or dual port nodes, the P1O bit is set if port 1 is operational, and is cleared if port 1 is not operational.

The P2O bit is only valid for single port and dual port nodes, Switch nodes shall clear the P2O bit. For single port or dual port nodes, the P2O bit is set if port 2 is operational, and is cleared if port 2 is not operational.

The LONG bit is set if any port in the node is equipped with the Extended Distance Option, and cleared otherwise. If the long bit is set, a QUERY PORT SMS may be used to determinewhich ports support this option.

If the CM bit is cleared, the node that issued the QUERY NODE REPLY SMS is not currently the Master. If the CM bit is set, the node that issued the QUERY NODE REPLY SMS is currently the Master. If the current Master receives a QUERY NODE SMS with a MASTER PRIORITY field value greater than its own, then it responds with a QUERY NODE REPLY SMS with the CM bit set and its master priority in the MASTER PRIORITY field, and then resigns mastership.

The Path component of the ADDRESS field in a QUERY NODE REPLY SMS frame is copied from the RETURN PATH field in the corresponding QUERY NODE SMS. All padding bytes shall be discarded.

12.2.9 QUERY PORT REPLY SMS

The QUERY PORT REPLY SMS defined in Table 38 is sent from a node to a Configutor node in response to a QUERY PORT SMS.

Table 38 - QUERY PORT REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Bh)							
1	reserved							
2	TAG							
3	TAG							
4	LINK ERP ERROR COUNT							
5	LINK ERP ERROR COUNT							
6	A QUOTA							
7	B QUOTA							
8	EUDC	REFLECT	MODE		reserved	EDO	EDOE	
9	reserved							
10	ALARM THRESHOLD							
11	ALARM THRESHOLD							
12	SUPPORTED SPEED							
13	SUPPORTED SPEED							
14	CURRENT SPEED							
15	CURRENT SPEED							

The TAG field is copied from the QUERY PORT SMS. It identifies the QUERY PORT SMS that this response is being generated for.

The LINK ERP ERROR COUNT field is the count of the number of times the Link ERP has been invoked for the selected port since the last power on, Total Reset Control frame, Absolute Reset Control frame or the last QUERY PORT MS with the CLE bit is set. A value of FFFFh is returned if the LINK ERP COUNTER is greater than or equal to 65535 (the count stops incrementing at 65535).

The A QUOTA field specifies the current frame A quota for the port set by the CONFIGURE PORT SMS for the selected port.

The B QUOTA field specifies the current frame B quota for the port set by the CONFIGURE PORT SMS for the selected port.

The EUDC bit specifies the current setting of the EUDC bit for the port set by the CONFIGURE PORT SMS for the selected port.

The REFLECT bit specifies the current setting of the REFLECT FLAG for the node.

The MODE field specifies the current port mode for the port set by the CONFIGURE PORT SMS for the selected port.

the EDO field specifies that the port is capable of supporting the Extended Distance Option

The EDOE field specifies that the Extended Distance Option is currently enabled for this port.

The ALARM THRESHOLD field is the current setting of the ALARM THRESHOLD set by the CONFIGURE PORT SMS for the selected port. See 11.2.3.

The SUPPORTED SPEED field is defined in Table 39. The SUP40 bit shall be set if the port is capable of operating at 40 MB/s and cleared otherwise. The SUP20 bit shall be set if the port is capable of operating at 20 MB/s and cleared otherwise.

Table 39 - SUPPORTED SPEED field

Byte	Bit 7	6	5	4	3	2	1	Bit 0
12	reserved							
13	reserved						SUP40	SUP20

The CURRENT SPEED field is defined in Table 40. If the port is currently operating at 40 MB/s, then the CUR40 bit shall be set and the CUR20 bit shall be cleared. If the port is currently operating at 20 MB/s, then the CUR40 bit shall be cleared and the CUR20 bit shall be set. If the port's OPERATIONAL FLAG is cleared, then both the CUR40 and the CUR20 bits shall be cleared.

Table 40 - CURRENT SPEED field

Byte	Bit 7	6	5	4	3	2	1	Bit 0
14	reserved							
15	reserved						CUR40	CUR20

Adge Hawes
Voice: +44 (0)1705 486363
FAX: +44 (0) 1705 499278
Email: adge@uk.ibm.com