Accredited Standards Committee
X3, Information Processing Systems

Doc: X3T10.1/96a113r0
Date: February 1, 1996
Project: X3T10.1 / 1147
Ref Doc.: SSA-TL2
Reply to: John Scheible

To: X3T10.1 Membership
From: John Scheible

Subject: Simple speed negotiation proposal

BACKGROUND

I am concerned that the current discussions on speed negotiation proposals have several disadvantages.

1) They change the hardware. The latest discussions involve a 10 MHz carrier wave to synchronize the speed negotiation process (in case one port is still in POST or is doing link recovery). There is a push to get 40 MB/sec silicon out in the marketplace yesterday. If we change the hardware design, that invalidates that effort and sends a wrong message to the world (SSA is not stable).
2) They are overly complex. The speed negotiation process being discussed keeps getting additional "features" added and tries to do everything. If we adopt this now, we have to live with it forever (for real downward compatibility reasons) even if it is changed so much in the future that it is invalidated. The reality is that we have 20 MB/sec devices in the marketplace today, and are working hard to get 20/40 or 40 only devices out soon. If we create a complex algorithm for 20/40 to handle everything, and then change it to the next stage (20/40/60?, 20/40/80?, 20/40/80/100?, 100?, 20/40/45?, 20/80?, 20/100?, 40/80?) we need to carry needless baggage. The future is uncertain.
3) They are not coming to a consensus quickly enough. We need proposals that are in the proper format to review comprehensively. 20 MB/sec is in letter ballot and effectively frozen. 40 MB/sec needs hardware now so that we can properly evaluate it and get the spec complete. There is a strong push to get the 40 MB/sec spec in public review by the end of this year (the project proposals show 8 months difference between TL1/PH1 and TL2/PH2).

This proposal is in the format to include in the specification now and is the simple approach. If we decide later that a more complex process is needed, then we can add it at that time. If the more complex process fails to synchronize then the future devices would fall back on this much simpler process.

THEORY

A single speed device operates as it does today, whether it is 20 MB/sec only or 40 MB/sec only. It essentially sends FLAG characters indefinitely until synchronization is archived.

A dual speed device operates similarly except that it alternates between the two speeds, with one being twice as long as the other to make sure that nodes do not oscillate and never achieve synchronization. This occurs indefinitely until synchronization is achieved. If one port starts late, character synchronization will still occur at the highest speed. If some pathological case occurs where the link synchronizes at a lower speed, the last paragraph of the theory section describes the solution. If additional speeds are added, then this process can be easily expanded with a geometric progression to avoid the oscillation problem (no guarantees that strange complex speed devices could not be selected and fail).

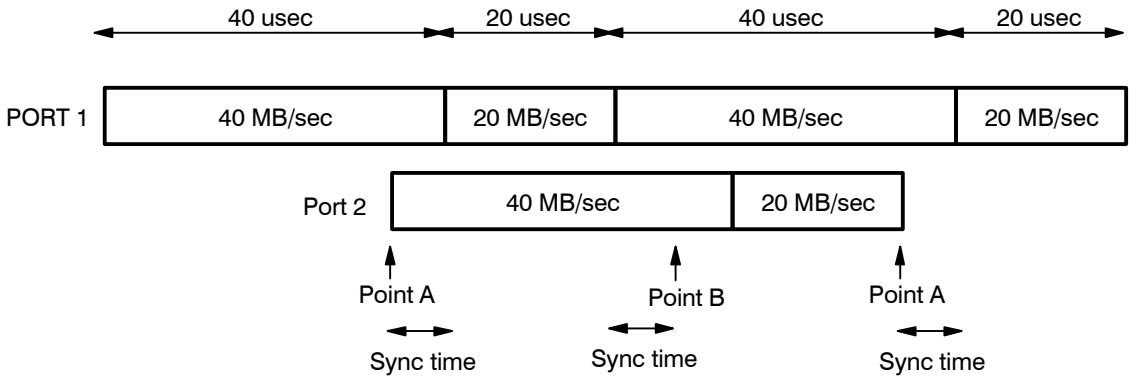

Figure 1 - Time diagram of two out of sync ports during speed negotiation

The 20 MB/sec case runs at the twice the maximum synchronization time of 200 characters (400 characters). The 40 MB/sec case runs twice as long (real-time) as the 20 MB/sec case (1600 characters) causing the condition shown in Figure 1. If Port 2 begins speed negotiation after port 1 but before point A, then synchronization occurs at 40 MB/sec. If Port 2 begins speed negotiation at point A, the synchronization occurs at 40 MB/sec at point B. If Port 2 begins speed negotiation after Point B but before point A, then synchronization occurs at 40 MB/sec.

If some pathological case (poor cables, long noise burst) that causes the link to synchronize at the slower speed (20 MB/sec) then the following OPTIONAL procedures could be used to speed negotiate again at 40 MB/sec. This process could also be used if a poor cable caused high error rates (alarm ~~alamrn~~ threshold exceeded), and the speed needed to be stepped down until the cable could be replaced.

1) The Master can notice the link running at less that optimum (two 40/20 ports running at 20) by using the QUERY PORT SMS during the Configuration process. It can issue a CONFIGURE PORT SMS to the nearest port and instruct it to speed negotiate as a 40 MB/sec only devices (SPEED NEGOTIATE field of 0002h). If speed negotiation fails (poor cable), then it can re-instruct the nearest port to speed negotiate as a 40/20 device (SPEED NEGOTIATE field of 0003h) and synchronize at 20 (or 40) MB/sec.
2) A “smart” node which is 40/20 capable but has negotiated at 20, could generate an Asynchronous Alert with a TBD ALERT CODE value to instruct the Master to perform the steps shown in 1) above.
3) A “really smart” node which is 40/20 capable but has negotiated at 20, could determine if the remote port was 40 capable by issuing a QUERY PORT SMS. If the remote port was 40 capable, then the local port could restart the speed negotiation process. If a certain number of retries failed, then the speed could remain at 20 MB/sec. A 20 only, 40 only, or 40/20 device running at 40 would not need to do anything.

PROPOSAL

The proposal changes SSA-TL2 as follows.

1) Modify 9.3 Beginning Communications” as shown on the next page (with revision marks).
2) Add new sections 9.4, 9.4.1, 9.4.2 on speed negotiation as shown on the next page.
3) Modify 11.2.3 (CONFIGURE PORT SMS) as shown on the following pages with revision marks to allow for a Speed Re-negotiation function. Total Reset affects both ports and is more disruptive that this new function.

The proposal (as of yet) does not add anything from the Theory section.

Sincerely,

John Scheible
Voice: (512) 823-8208

FAX: (512) 823-0758
Email: Scheible@vnet.ibm.com

9.3 Beginning Communication process

A port is in the Disabled state and Privileged mode after the following conditions:

a) POST;
b) a Total Reset Control frame is processed;
c) an Absolute Reset Control frame is processed;
d) ~~or~~ the Link ERP process exits with an ALERT CODE value;
e) or a CONFIGURE PORT SMS is received with a non-zero NEGOTIATE SPEED field.

The port shall then use the following procedure to begin communication.

1) If the line fault detector indicates a line fault, then wait until there is no line fault. A line fault indicates that the remote node is not powered-on or that the cable is disconnected.
2) Invoke the Speed Negotiation process
3) S~~Wait for the port to enter the Ready state, then s~~et the port's OPERATIONAL FLAG.
4) Wait for at least 10 FLAG characters to be sent
5) Ensure that buffer space is available to receive at least 1 frame of any type. Then send an RR character pair and clear the RR PENDING FLAG (see 8.4.2.)
6) Wait to receive an RR character pair from the remote port. Then clear the WAITING FOR RR FLAG.

Privileged frames are now able to be received and transmitted.

9.4 Speed Negotiation process

The Speed Negotiation process is exited whenever synchronization is achieved (when the port enters the Ready state). The speed Negotiation process can be broken down into two versions, single speed and dual speed.

9.4.1 Single speed

The single speed Speed Negotiation process sends at least 200 DIS characters to allow the remote port to synchronize prior to potentially forwarding a Control frame and then sends FLAG characters indefinitely until synchronization is achieved.

1) Transmit at least 200 DIS characters.
2) Put the port into the Enabled state.
3) Wait until character synchronization is achieved.

9.4.1 Dual speed

The dual speed Speed Negotiation process alternates between the two speeds. In each speed, the process sends 200 DIS characters to allow the remote port to synchronize prior to possibly forwarding a Control frame and then sends FLAG characters until either synchronization is achieved or 800 FLAG characters have been sent. There is a 10% tolerance to the numbers to allow firmware implementation rather than hardware

The following process is repeated indefinitely until character synchronization is achieved.

1) Set the transmitter at the faster speed (i.e., 400 MHz).
2) Transmit 180-220 DIS characters.
3) Put the port into the Enabled state.
4) Transmit FLAG characters until synchronization is established or 910-1090 FLAG characters have been sent.
5) Set the transmitter at the slower speed (i.e, 200 MHz).
6) Transmit 180-220 DIS characters.
7) Put the port into the Enabled state.

8) Transmit FLAG characters until synchronization is established or 80-220 FLAG characters have been sent.

11.2.3 CONFIGURE PORT SMS

The CONFIGURE PORT SMS defined in Table 28 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 Priviledged mode. In both cases the destination node returns a RESPONSE SMS.

Table 28 - 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			
12	reserved							
13	ALARM THRESHOLD							
14	ALARM THRESHOLD							
15	NEGOTIATE SPEED							
16	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 identifies the port that is being configured.

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 Master Asynchronous Alert table (see 9.2.2). 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 is allowed to 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 is allowed to 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 originate, nor forward any user defined characters from any other port of the node. If the EUDC bit is set, the port is allowed to originate user defined characters, and to forward user defined characters from other ports of the node that also have their respective EUDC bit set.

The REFLECT bit is one of the controls affecting the REFLECTION FLAG node function (see 9.2.6). 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 29.

Table 29 - 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 9.1.4.3). 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 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 11.2.8).

The NEGOTIATE SPEED field defined in Table 30 indicates what speeds the port should use when speed negotiating during the Beginning Communication process (see??.?). The changes to speed negotiation caused by the state of the NEG40 and NEG20 flags is defined in Table 31.

Table 30 - NEGOTIATE SPEED field

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

Table 31 - NEGOTIATE SPEED field effects on speed negotiation

NEG20	NEG40	Speeds to be negotiated
Cleared	Cleared	No change to speed negotiation
Cleared	Set	Speed negotiate only at 40 MB/sec
Set	Cleared	Speed negotiate only at 20 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 11.2.14).