

High Speed Token Ring Alliance

October 30th - 31st, 1997
Oxford, United Kingdom

Minutes of Technical Meeting #2

Neil Jarvis

SilCom (A Microvitec Group Company), R'n'D, York, UK

Neil.Jarvis@mail.york.proteon.com
http://www.york.proteon.com/hstr/
http://www.york.proteon.com/802.5/

<i>Revision History</i>		
v1.0	01 Nov 97	First Release

Attendees

Karl Reinke,	Bay Networks	Dave Wilson,	Madge
Ken Wilson,	IBM	Andy Fierman,	Madge
Tam Ross,	IBM (Teleconference)	Neil Jarvis,	SilCom
Joe McDonald,	IBM (Teleconference)	John Messenger,	SilCom
Bob Love,	IBM	Ivar Jeppesen,	Olicom
Bob Hubbard,	Cisco	Jens Andreassen,	Olicom
Ivan Oakley,	Cisco (Part time)	Bo Thomsen,	Olicom
Mike Cohen,	Madge (Part time)	Mark French,	Treseq
Simon Harrison,	Madge	Mick Hanrahan,	Texas Instruments
Richard Knight,	Madge		

Agenda.

Thursday 30th October:

8.45	Coffee & Tea available in meeting room
9.00	Welcome & Introduction - Mike Cohen, Madge Networks
9.15 - 10.45	Technical Discussion on Lobe Test
10.45 - 11.00	Coffee Break
11.00 - 1.00	Technical Discussion on Lobe Test
1.00 - 2.00	Buffet Lunch
2.00 - 3.30	Technical Discussion on Phantom
3.30 - 3.45	Coffee Break
3.45 - 5.30	Technical Discussion on Phantom

Friday 31st October:

8.45	Coffee & Tea available in meeting room
9.00 - 10.45	Review of proposed clause 10 and layout options
10.45 - 11.00	Coffee Break
11.00 - 1.00	Review of Strawman
1.00 - 2.00	Buffet Lunch
2.00 - 4.00	Review of Strawman
4.00	Closing Remarks - Mike Cohen, Madge Networks.

DAY 1

Introduction.

Bob Love: Get it right, do it on time.

Mike Cohen: Welcome to Oxford.

Lobe Media Test.

Dave gave a recap of his proposal, as modified by Neil Jarvis and Ken Wilson.

The “Wrap Data” data stream received by the station is not guaranteed to be the same data it transmitted. Its length and the remainder of the frame shall be the same.

Q: Does 112 include FCS?

A: Yes

Q: Some classic Lobe Media Tests used DAT frames. Do we want to keep this?

A: No.

The standard shall mandate that the C-Port must repeat 1120 frames within 2.3s, and that the total length of the TEST MAC frame is 112 bytes (inclusive of FCS) (this prevent extended vectors & subvector). The standard shall show this lobe media test as an independent state table (Neil’s modifications).

C-Port support of Lobe Media Test at 100Mbit/s will not use phantom signal detection (if present) to create the repeat path.

[Resolved]

Phantom Drive and Wire Fault.

There are four types of Station:

1. Phantom signaling and wire fault detection
2. Phantom signaling and no wire fault detection
3. No phantom signaling and wire fault detection
4. No phantom signaling and no wire fault detection

There are four types of C-Port:

1. Phantom detection and wire fault load provision
2. Phantom detection and no wire fault load provision
3. No phantom detection and wire fault load provision
4. No phantom detection and no wire fault load provision

PD subvector and SPV(PD) and PPV(PD_MASK) need to reflect this. The following are clearer definitions of PD subvector values and PPV(PD_MASK).

Station PD Subvector Definition

Value	Definition
X'0001'	The Station supports the phantom signaling and wire fault detection method described in 7.2.1 in ISO/IEC 8802-5:1995.
X'0002'	The Station does not support the phantom signaling. The Station cannot support wire fault detection method.

C-Port PPV(PD_MASK) Definition

Value	Definition
X'0001'	The C-Port supports the phantom signaling detection and wire fault load provision.
X'0002'	The C-Port supports neither phantom signaling detection nor wire fault load provision.
X'0003'	The C-Port does not support phantom signaling detection, but supports wire fault load provision.

Lobe Media Test at 100Mbit/s will not use phantom signal detection (if present) to create the repeat path.

Q1: How does the C-Port detect a station closing?

Q2: How does the C-Port detect station power off?

Q3: How does the C-Port detect the cable being pulled out?

A1: Use phantom signal loss detection - protect with timer during beaconing.

A2: Use RMV_ALRT MAC frame.

A3: Use loss of link_status - protect with timer during beaconing.

What happens when the C-Port starts beaconing, but neither link_status nor phantom go away? A timer (mandatory at 100Mbit/s, optional at 4/16) will start on entry to beaconing, and on expiry cause the C-Port to return to bypass. (How long?). This timer should be the exception, not the norm. With correct cable termination, link_status staying asserted should not happen (but is still possible).

Q: Do we need a PHY section entry to define what a station should do on power off? (Termination etc.)

A: Yes. Define receive return loss termination in the Station at power-off.

<Andy to check if the FDDI spec has a typo in its definition of this receive return loss>

[Unresolved]

Q: Do we need a clear statement of clock tolerances for HSTR?

A: Yes.

<Andy will write some words>

[Unresolved]

Q: If an entity has phantom circuitry, it may violate voltage isolation requirements that are standardised for 100BASETX. Do we need a statement to allow this violation?

A: Yes

<Andy will write some words>

[Unresolved]

Q: How many RMV_ALRT MAC frames do you send?

A: n7, using a pacing timer.

Both the Station and C-Port will send this frame in response to Disconnect.xMAC

All three answers (A1 to A3) are to be added to the next draft of the standard.

[Resolved]

Reduced MII (RMII) Presentation (Mark French, Treseq and Bo Thomsen, Olicom).

RMII is aimed at multi-port PHY devices (8+ ports). In the future, some PHY vendors may move more and more towards selling RMII vs. MII devices.

From the presentation, issues specific to RMII:

- Elastic buffer designed to work with 1500 byte frames. Worst case 18k frames cause elastic buffer over/under-flows. There are 10 bits in the ethernet EB. At least 40 bits are required for token ring.
 - Increase EB size. This must be switchable, since the large buffer would introduce delays for ethernet. **POSSIBLE**
 - Specify higher tolerance clocks. 8ppm for 18k, 25ppm for 4.5k. All entities must use these better clocks. **BAD**
 - Use loop timing. If the receive clock goes out of range, they attempt to resynchronise to transmit. This does not work on some PHYs. Cannot be implemented? **BAD**
- No abort sequence support.
 - TX_ER not present.
 - Could use invalid FCS and E bit.
 - Requires E bit in ET.
 - Make this the only abort sequence for HSTR, otherwise RMII is dead.
 - Make PHY vendors change their PHYs to allow token ring to indicate an abort across the RMII to produce the /H/H/T/R/ abort sequence currently defined for HSTR.
- Start of frame detection is difficult.
 - False carrier and frames could look the same?
 - Possible separate CRS and RX_DV, by masking out CRS from CRS_DV.
 - Could insert Ethernet SFD before AC to help identify start of frame.
 - Must be done now, otherwise RMII is dead.

Summary of the impact of supporting RMII to the PHY and HSTR standard

Issues	PHY Changes	HSTR Changes
Elastic buffer	✓	✗
Abort to wire without TX_ER	✗	Invalid FCS & 'E' bit definition
CRS_DV/False carrier	✓	Define SFD
Frame destruction after violation (Destroys ET)	✓	✗

List of changes proposed for RMII:

- 1) Token Ring mode selection.
- 2) Large elastic buffer size.
- 3) Remove CRS signalling on CRS_DV signal.
- 4) Make frame sequence compatible with MII frame sequence. Don't trash frames after code violations.
- 5) Signal abort to generate /H/H/T/R/ abort sequence (out-of-band RMII signal).

or

- Add back the TX_ER pin.
- 6) Programmable auto negotiation selector field.

Q: Will PHY vendors be prepared to do this?

[Unresolved. We need responses before the plenary meeting.]

DAY 2

RMII Vendor response (Karl Reinke, Bay Networks)

Karl has had some response from an RMII vendor (National Semiconductors) on the list of RMII issues. They are positive on being able to accommodate the token ring mode changes.

RMII Token Ring Requirements	Vendor Response
Programmable Auto-negotiation selector field	Bits [10:0] of REG 07 (ANNPTR)Auto-Negotiation Next Page Transmit Register are available to program as you wish per the Standard
Token Ring Mode enable bit	Without knowing the availability of unused register bits, no commitment can be made. However, if there are unused bits, and notice is given with sufficient time prior to tape-out, yes; it can be done... <<This could be the selector field>>
Larger Elastic Buffer (40 bits?)	An additional depth to 32-64 would not be an issue.
CRS_DV is DV only	Yes
Don't trash rest of frame after code violation	I believe that the 100baseX PCS does not "'trash' the rest of the frame". An assertion of RX_ER signals a code violation and the CRC at the end of the packet invalidates the frame contents at the MAC level. If bits are either dropped or picked up such that another valid data code is produced, the responsibility again falls upon the CRC at the MAC layer to resolve the error. << Clarification required >>
Add TX_ER. If not, another Mechanism to do TX_ER (out-of-band RMII signal?)	Yes to both. TX_ER could conceivably be restored by combining with TX_EN and used as a method of encoding as the fifth bit via the bypass 4B/5B feature. Again, dependent upon a 'TR' enable bit.

Q: How is the /H/ symbol indicated across the RMII interface?

[Unresolved]

RMII Issues and Considerations (Bob Love, IBM).

- Abort delimiter is still /H/H/T/R/
- Support for E bit in the standard

Scrambler reset timer is needed for both the MII and the RMII.

[Wait until Tam tele-conference]

The abort sequence is /H/H/T/R/.

On receipt of this sequence, the frame is considered aborted.

In addition, the reception of an invalid frame (code violation, invalid FCS) with the E bit set, the frame is ignored, and not counted as a line error.

Aborting a frame using an invalid FCS and the E bit set *will not be defined in the standard today*. However, if in the future, if the RMII cannot be modified to support the transmission of /H/H/T/R/ sequence, then the invalid FCS with the E bit set will be added to the standard.

[Resolved]

Publishing 100 Mbit/s Frames and Facilities clause.

Strong opinion has been expressed that the 100 Mbit/s modification for Frames and Facilities should be published as a new clause, rather than modifications to the existing clause 10. This is a good position for existing implementers, but leaves new implementers having to read clause 3, 10 and 12 to understand how to build a 100Mbit/s entity.

Decision will be made before the end of the meeting about which format the modifications should be published in.

Merge into Clause 10	3
Create new Clause 12	9
Don't care	3

Editor will create a clause 12.

[Resolved]

This is now completely different from how clause 9 will be published. It will remain a merged clause showing 4/16 and 100 Mbit/s transitions. Go figure...

FR, FR_WITH_ERR and FR_MAC definitions.

Deferred until the IEEE 802.5 Plenary meeting, to allow Neil to write a presentation.

< Note to Mick: Add agenda item for this. >

[Resolved]

Code violations outside the frame (CxIDE errors).

CxIDE needs 2 non adjacent bits in 10, in the IDLE stream to be corrupted into 0s before the error condition occurs.

Do we need this new counter? We can monitor the link condition by looking at the other error conditions.

< Remove CxIDE >

[Resolved]

HSTR Speed Trade-up procedure (Ivar Jeppesen, Olicom).

In a DTR link, where both entities are capable of 4/16/100 Mbit/s, it would be nice if there were a method of determining that a link running at a lower speed, could "trade-up" to a higher media rate.

There are two parts to this problem:

- 1) If a 4/16/100 Station attempts to join at 100 into a 4/16 C-Port, then there is currently a 5s timer, whose expiry indicates a failure to join. This will delay the entity joining correctly at 16.
< Ken will change this timer, and correct the use of the FxANO option flag. >
- 2) If a link comes up at a lower speed than it is capable of, a new protocol could trade-up. Ivar will present a paper at the IEEE 802.5 Plenary meeting.
< Agenda item required. >

[Unresolved]

HSTR Fibre.

We must ensure that technical decisions made in relation to 100BASETX do not adversely affect the definition and publication of the HSTR over fibre document. For example, phantom vs. fibre keying.

< *Committee to investigate these issues* >

[Unresolved]

HSTR Gigabit.

We need to review the differences between MII and GMII.

< *Richard Knight has been volunteered to champion this investigation.* >

[Unresolved]

Answers from MII vendors to meeting #1 questions (Tam).

PHY vendors questions from the minutes to meeting #1:

1) Length of frame (do long frames upset synchronisation?), MLT-3: more baseline wander with longer frames.

1ms synchronisation timer length in existing PHYs. We need 1.5ms for 18k frames. Nat Semi has an extended timer (2ms) which would support 18k frames.

2) Inter-frame gap of 12 bytes

This is enough, even if the synch timer is increased.

3) Preamble requirements

No.

4) Baseline wander

No issue identified yet, but don't close the issue.

5) Phantom signalling into ethernet or token ring connection

Devices should not be harmed, but they are not happy either.

Not answered: Phantom into a token ring C-Port which does or does not support phantom.

[Unresolved]

6) 18 byte frames?

No trouble.

[Mostly Resolved]

ISO/IEC 8802-5:1997.

The proposal is that the IEEE will publish the following documents

ISO/IEC 8802-5:1998 consisting of

ISO/IEC 8802-5:1995

802.5s

Supplement to ISO/IEC 8802-5:1998 consisting of

802.5r

802.5j

The overwhelming opinion in the committee is that this is not how the standards should be published.

However, if this method of publication does go ahead, then Bob must ensure that 802.5r and 805.5j are **not** merged. They must remain as separate documents within one document.

[Unresolved]

Tutorial at next Plenary.

Scott Valcourt will present the Tuesday tutorial at the Montreal Plenary meeting about the HSTR PARs.

[Resolved]

Meeting adjourned.