100 Mbit/s 802.5 PHY Layer – Changes to Standard

T. J. Ross
IBM, NHD, D63/656
RTP, NC

Phone: (919) 543-2633
FAX: (919) 253-9508

tamross@vnet.ibm.com

September 29, 1997

Editorial Comments

(Items in Itlaics are comments and explanations. They are NOT intended to be included in the standard)

(Bob, Ken: the 802.3 world sometimes publishes a "BLUE BOOK" which contains the explanations and history of how and why items in the standard were chosen. It is meant to help implementers and future revisors to understand the reasoning behind the standard. Usually the standard itself is too terse to get a feel for the these issues. Do we wish to do anything of this sort?)

MAC/PHY Interface

Overview

Scope

This clause describes the interface between the MAC and PHY layers of the 100Mbps dedicated token ring. The PHY interface is precisely defined as the Media Indedendent Interface (MII) . {802.3 clause 22 “Reconciliation Sublayer} both defines the MII and maps the signals into primitives used by the 802.3 MAC. This clause maps the MII signals into MAC primitives used by the 802.5 MAC.

Objectives

- Define a simple mapping of MII signals into MAC primitives
- Maintain, as closely as possible, the 4/16 802.5 signalling interface.

9.a.1 100 Mbps Protocol Support (Common)

x.1 Overview

An explicit embodiment of the MII is not required. If a physical embodiment of the MII is not present, then it is strongly recommeded that the implementation provide control and status mechanisms equivalent to those described in {802.3-22}. If an exposed MII does exist, it shall meet all requirements of {802.3-22.x}.

x.2 MAC Interface Service Specification

The following service primitives specify the required information that is passed between the MAC and the PHY. This service specification replaces the service specification defined in 5.1.2, 5.1.3, and 5.1.4. (or should this be in x.1 about all of chapter 5).

The service primitives are grouped into two classes: physical signalling (PS) and physical station management (PSMT). All PHY activities related to the control, transmission, and reception of data are included in the PS primitives. Activities related to the configuration and negotiation of the PHYs and related status are included in the PSMT primitives. Implementation of the set of PSMT functions is optional, however if implemented the functions shall be compliant with the description given.

(Question: how parallel do we wish to be with the 802.3? All of these primitives could be references to those defined in 802.3 22.2.1 “Mapping of MII signals to PLS service primitves and station management”. But some concepts might need to be changed, e.g. how we handle inter-frame bits. A level of abstraction will provided us with the ability to change definitions, but might also become very confusing, especially since the names will be so similar.)

I suppose, in a perfect world, I would need to completely re-write the reconciliation layer specification. That would conform nicely to the OSI model layers. Instead, I’m really writing a layer the “reconciles” the 802.3 reconciliation layer to the 802.5 MAC. (Ken – how picky is the standards group about well-structured documents versus ones that just work?)

In the following, a nibble is a 4-bit interface, exchanged with the MAC at a 40 ns time interval.

x.2.1 PS_UNITDATA.indication

This primitive defines the transfer of data from the PSC to the MAC.

PS_UNITDATA.indication[Rcv_symbol (ref)]

The Rcv_symbol specified is one of the following:

- Start_delimiter
- Data_nibble
- End_delimiter
- Idle

Start_Delimiter is a 2-nibble field which indicates the reception of the beginning of a valid frame. The Start_Delimiter is equivalent to the {802.3-24.2.2.1.4 Start-of-stream Delimiter (SSD)}.

The analogous 802.3 primitive is PLS_DATA.indicate(INPUT_UNIT) which has values 0 or 1. No value is defined for SDEL or EDEL. I assume (always dangerous) that they are not needed because data is only indicated between these boundaries, i.e. there is no "indication" between frames. Can we live with this?

[Note for Tam: please check the exact requirements for RX_DV per 802.3 Figure 22-6. Must RX_DV rise immediately after SSD? Clause 22 does not require, but clause 24 (TX) state diagram definitely does!]

Data_nibble is a 1-nibble value containing the data received from the medium. Data_nibble is the MII signal RXD<3:0>. Rcv_symbol has the state Data_nibble after Start_delimiter has been received up until End_delimiter has been received.

End_Delimiter indicates the end of a valid frame. The End_Delimiter is equivalent to the {802.3-24.2.2.1.5 End-of-stream delimiter (ESD)}. End_Delimiters not on nibble boundards may not be detected.

Idle is a 1-nibble value used only between frames. Rcv_symbol has the state Idle after End_delimiter has been received up until tart_delimiter has been received.

Alignment between the Frame boundaries and the MII control signals.

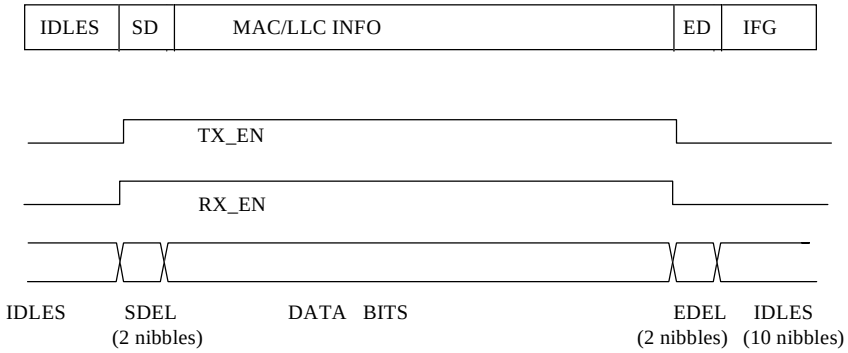

Figure X.XX:

x.2.2 PS_UNITDATA.request

This primitive defines the transfer of data from the MAC to the PSC.

PS_UNITDATA.request[Tx_symbol (ref)]

The Tx_symbol specified is one of the following:

- Start_delimiter
- Data_nibble
- End_delimiter
- Idle

Start_Delimiter is a 2-nibble field which indicates the reception of the beginning of a valid frame. The Start_Delimiter is equivalent to the {802.3-24.2.2.1.4 Start-of-stream Delimiter (SSD)}.

The analogous 802.3 primitive is PLS_DATA.request(OUTPUT_UNIT) which has values 0,1, or DATA_COMPLETE. Since there is a requirement on the reconciliation sublayer to generate TX_EN with the first nibble of data, there is no need for a corresponding primitive. Can we life with that?

Data_nibble is a 1-nibble value containing the data to be transmitted on the medium. Data_nibble is the MII signal TXD<3:0>. Tx_symbol has the state Data_nibble after Start_delimiter has been received up until End_delimiter has been received. At all other times, the TXD<3:0> is ignored by the PHY.

End_Delimiter indicates the end of a valid frame. The End_Delimiter is equivalent to the {802.3-24.2.2.1.5 End-of-stream delimiter (ESD)}.

Idle is a 1-nibble value used only between frames. Tx_symbol has the state Idle after End_delimiter has been received up until Start_delimiter has been received.

x.2.3 PS_STATUS.indication

This primitive is used by the PSC to inform the MAC of errors and significant status changes.

PS_STATUS.indication [Receive_Clock (ref), ???
Transmit_Clock (ref), ???
Signal_Acquired (ref),
Frame_Violation (ref),
Burst_Error (ref)]

Frame_Violation is the logical AND of {MII Receive_Error (RX_ER)} with {MII Receive_Data_Valid (RX_DV)}..

Burst_Error is the logical AND of {MII Receive_Error (RX_ER)} with NOT {MII Receive_Data_Valid (RX_DV)}..

Signal_Acquired is equivalent to {MII Carrier_Sense (CRS)}.

x.2.4 PS_CONTROL.request

This primitive is used by the MAC to request certain actions of the PSC.

PS_CONTROL.request [Initialize (ref),
Reset (ref),
Transmit_Error (ref) ???]

Initialize shall be defined as setting bit 0.9 of the PSMT_CONTROL.request (Restart Auto-Negotion), and setting all other bits to the normal operation state (if defined). If the PHY reports via bit 1.3 of the PSMT_STATUS.indication that it lacks the ability to perform Auto-Negotiation then Initialize shall be defined as setting all bits to the normal operation state (if defined). The PHYs shall be able to correctly exchange data with ?? s of the assertion of Initialize.

(This keeps getting worse . . .
come alive from reset or power-down: (no, this is included below)
500 ms
AN waits in TRANSMIT DISABLE before starting
1500 ms
exchange base pages (I found this before, but can't now!)
1000 ms
bring up TX link
1000 ms
total:
3500 ms

The time needed to negotiate extended pages is open ended, but as of now, I see no reason for us to use extended pages. The auto-negotiate does not start until a maximum of 500 ms after either Reset or Power-down are cleared to zero.)

Reset is intended to halt operation of the PHY. The exact state of the PHY after Reset is asserted is undefined. Suggested actions are setting the Reset (0.15) or Power Down (0.11) bits of the PSMT_CONTROL.request. *(What happens if a REMOVED station receives a frame???)*

Transmit_Error (do we need this?)

The 802.5 concepts, INSERT and REMOVE, each have two applications: the transition into and out of the BYPASS state (i.e. opening and closing the adapter); and the transition into and out of the LOBE TEST state (i.e. preparing a wrap path for transmission testing). Here, the two applications have SEPARATE primitives. Initialize and Reset are used to transition from and to the BYPASS state. The LOBE TEST transitions are performed with MAC frames, and are not included here, because they are no longer PHY related.

x.2.5 PSMT_CONTROL.request

PSMT_CONTROL.request [Req_function {802.3-22.2.4.1}]

Where Req_function is the assertion or clearing of bits in the {802.3-22.2.4} Management functions reigisters. Bit definitions changed for 802.5 are listed in section (ref) PHY Layer Definitions.

The most important PSMT_CONTROL is Link_Status. We may wish to use this as an indication to start the registration state. It may be instead of, or in conjunction with, a timer.

x.2.6 PSMT_STATUS.indication

PSMT_STATUS.indication [Status_function {802.3-22.2.4.2}]

Where Req_function is the reading of bits in the {802.3-22.2.4} Management functions reigisters. Bit definitions changed for 802.5 are listed in section (ref) PHY Layer Definitions.

9.a.2 Station-specific 100 Mbps Protocol Specification

(null set)

9.b.2 C-Port-specific 100 Mbps Protocol Specification

(null set)

Framing

(someplace, the IFG needs to be defined as at least 10 nibbles of Idles)

PHY Layer Definition

Overview

Scope

This physical layer (PHY) is divided into two sublayers: physical signaling components (PSC) and physical Medium components (PMC). This clause describes the operation of the PSC and the PMC. This clause replaces clause 5: Station sepcific components, and clause 7: Station attachment specificatins, *(and clause 8: concentrato specifications ?),* for 100 Mbps operation. Figure (ref) shows a comparison of the ISO model layers for 802.3 and 802.5

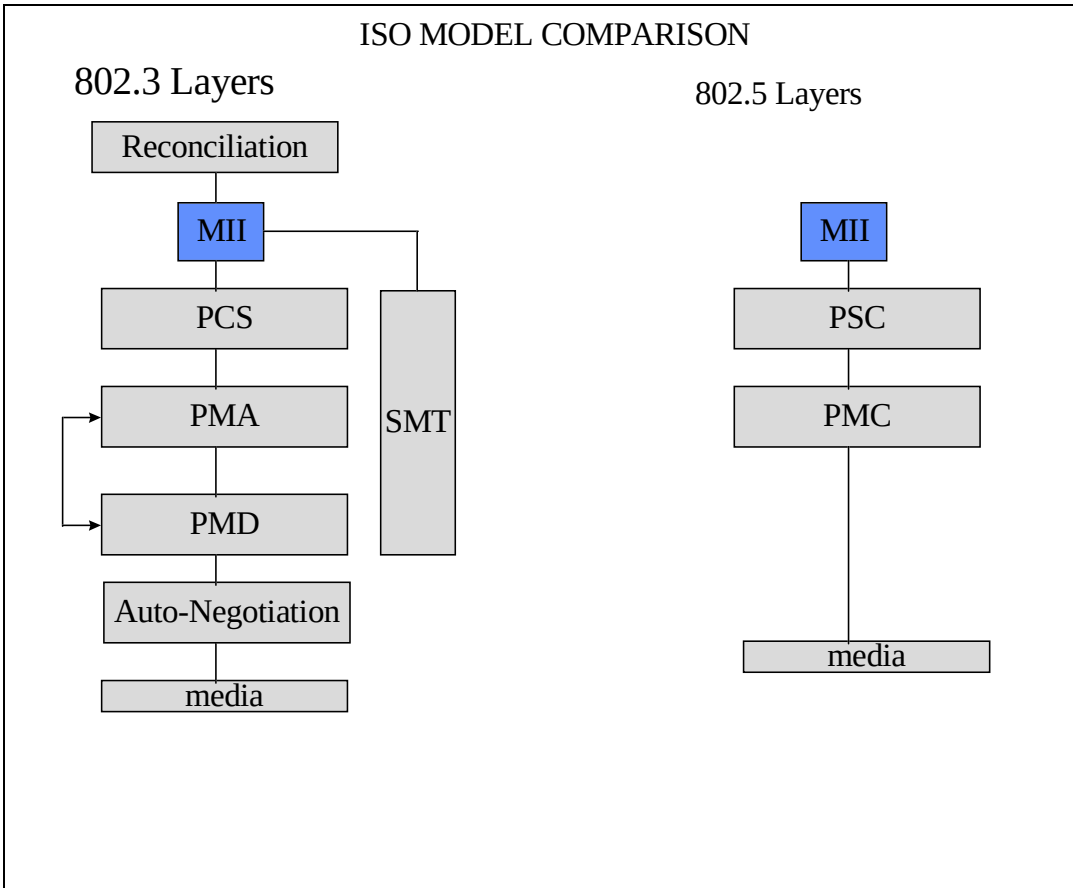

Objectives

- Provide a low-cost physical connection for switched token ring by using a pre-defined transceiver.
- Provide a standard interface, capable of being extended to additional media types, such as optical fiber.

Functional specifications

The 100 Mbps PHY is specified by incorporating portions of the ANSI/IEEE 802.3 standard, by reference, with the modifications noted below. The PSC shall meet all requirement of 802.3 Clause 24: Physical coding sublayer (PCS) and physical medium attachment (PMA), type 100BASE-X. The PMC shall meet all requirements of 802.3 Clause 25: Physical layer entity (PHY) and baseband medium, type 100BASE-TX.

These two clauses reference other portions of the 802.3 standard, notably Clause 22: Reconciliation sublayer and media independent interface, and clause 28: Physical Layer link signaling for 10 Mb/s and 100 Mb/s Auto-Negotiation on twisted-pair. Requirements placed by reference shall be met by the 100 Mbps PHY wherever applicable.

Do we want to include a version and date of the above? This would prevent future changes from affecting us, but this might also destroy our ability to use the next generation of TX parts.

General exceptions

1. The definitions of the bits in the auto-negotiation registers will be changed to reflect the requirements of 802.5, as selected by the Selector Field.

2. The physical layer device shall support Full Duplex transmission.

Specific exceptions

somewhere, Full duplex needs to be specified as a requirement. But where? I can't find any mention of full or half duplex in clauses 24 and/or 25.

Change to 25.4.3

Clause 25.4.3 is replaced by the following:

25.4.3 Contact Assignments

100 Mbps Token Ring for unshielded twisted pair (UTP) adopts the contact assignments of 4/16 UMIC. Therefore the contact assignments shown in [TP-PMD] table8-1 shall instead be depicted in table 25-2.

Table 25-2, "UTP MIC contact assignments"

CONTACT	PHY without internal crossover MID SIGNAL	PHY with internal crossover MDI SIGNAL
1		
2		
3	Transmit +	Receive+
4	Receive+	Transmit+
5	Receive-	Transmit-
6	Transmit-	Receive-
7		
8		

100 Mbps Token Ring for Shielded twisted pair (STP) adopts the contact assignments of the 4/16 MIC_S. Therefore the contact assignments shall meet the mapping shown in [802.5] clause 7.2, Table 15 - "Station port signal to MIC_S and MIC_U contact mapping", and the MIC_S shall meet the requirements of 7.2.5.1 "MIC-S (STP medium interface connector)". *The STP-MIC defined in [TP-PMD] 8.2.1 shall be an acceptable alternative for the MIC_S.*

Change to Annex 28A:

The following 5 bit word is added to the selector field mappings in Table 28A-1:

V	W	X	Y	Z	IEEE 802.5

Change to Annexes 28B: Selector base page definition, 28C: Next page message code field definitions.

The technology field definitions in these annexes are replaced by the following.

28B.2 Technology ability field bit assignments

(Check out how 802.9 did this. Since we are defining this for 802.5, it is not a change to the 802.3 standard, but I don't wish to copy all of clause 28 (or even all of Annex 28B) if I can help it.)

The Technology bit field consists of bits D5 through D12, (A0-A8 respectively) in the IEEE 802.3 Selector Base Page. Table 28B-1 summarzies the bit assignments.

Note that the order of the bits within the Technolgy Ability Field has no relationship to the relative priority of the technologies.

A0		*(10BaseT Half Duplex)*
Bit	Technolgy	Notes
A1		*(10BaseT Full Duplex)*
A2	100Base-TX Half Duplex Capability	
A3	100Base-TX Full Duplex Capability	
A4	100Base-T4 Capability	*(ICS TX, National TX PHY cannot write this bit)*
A5	Switched Port	?
A6	RESERVED for Shared Port	?
A7		?

(Requested) Changes to ANSI/IEEE Std 802.3, 1996, and its approved supplements, Clause 28

Table 28A-1 – Selector field value mappings

S4	S3	S2	S1	S0	Selector description
0	0	0	0	0	RESERVED for Future Augo-Negotiatioin Development
0	0	0	0	1	IEEE 802.3
0	0	0	1	0	IEEE 802.9
V	W	X	Y	Z	IEEE 802.5
1	1	1	1	1	RESERVED for Future Auto-Negotiation Development

Glossary:

PCS Physical coding sublayer : the 802.3 layer equivalent to the PSC

PMA Physical Medium attachment: the 802.3 sublayer of the PHY responsible for interfacing with the transmission medium, independant of the type of medium. The PMA and PMD together are equivalent to the PMC.

PMD Physical Medium dependant: the 802.3 sublayer of the PHY responsible for interfacing to a specific type of transmission medium. The PMA and PMD together are equivalent to the PMC.

FDX This term is used is used to indicate the *"Full-Duplex"* operation supported by the 802.5 TXI Access Protocol. The TXI Access Protocol is used to support the 100/16/4 Mbit/s "Dedicated Token Ring" Station.

HDX This term is used is used to indicate the *"Half-Duplex"* operation supported by the 802.5 TKP Access Protocol. The TKP Access Protocol is used to support the 16/4 Mbit/s "Token Passing" Token Ring Station. This protocol can be modified to support the 100 Mbit/s "Token Passing" Station.

HFDX This term is used is used to indicate the *"Half Full-Duplex"* operation supported by the *modified* 802.5 TXI Access Protocol. The HFDX Access Protocol is used to support 100 Mbit/s "Token Passing" Station.

LFFDX This term is used is used to indicate the *"Low Function Full-Duplex"* operation supported by the *modified* 802.5 TXI Access Protocol. The LFTXI Access Protocol is used to support the 100/16/4 Mbit/s "Low Function Dedicated Token Ring" Station.

The following terms are used to describe the protocol defined by the ISO/IEC 8802-5 Standard.

TKP Access Protocol This protocol is specified by the ISO/IEC 8802-5:1995 and its update: ISO/IEC 8802-5:1997 REV (Revision) Standards. These Standards specify the original "Transmit on Token Protocol".

TXI Access Protocol This protocol is defined in the ISO/IEC 8802-5:1997 Supplement to the ISO/IEC 8802-5:1995 Standard titled "Dedicated Token Ring". This Supplement specifies the "Transmit Immediate Protocol".

The following terms are used to describe the protocols which will be proposed as an update to the ISO/IEC 8802-5.2 Revision to support 100 Mbit/s Station using "Token Passing".

HFDX Access Protocol This *shared* protocol is proposed to support the HSTR Splitter, HSTR Station, and the 100/16/4 Switch.

LFTXI Access Protocol This *switched* protocol is proposed to support the Low Function HSTR Station and the 100/16/4 Switch.

Annex XXA (informative)
Minimal PHY Design

Objective:

Provide information to allow implementers to design a minimum-function 100Mb/s transceiver.

100BaseTX Functions required by 100Mb/s Token Ring:

1) MII Signals:
- TXD
- TX_EN
- TX_CLK
- RXD
- RX_ER
- RX_CLK
- RX_DV
- MDC
- MDIO

2) Abilities:
- 100Base-X Full Duplex

100BaseTX Functions optional for 100Mbps Token Ring: (used if present)

1) Abilities:
- Auto-Negotiation Ability

100BaseTX Functions not used by 100Mb/s Token Ring:

1) MII Signals:
- COL
- CRS
- TX_ER

2) Abilities:
- 100Base-T4
- 100Base-X Half Duplex
- 10 Mb/s Full or Half Duplex
- 100Base-T2 Full or Half Duplex
- Preamble suppression
- Remote Fault
- Extended (register set) capabilities.