

P1394b Start-up

Colin Whitby-Strevens and friends

1.0 Introduction

This document outlines a start-up procedure for P1394b links. It is independent of the coding scheme used, though it indicates some requirements on such coding schemes.

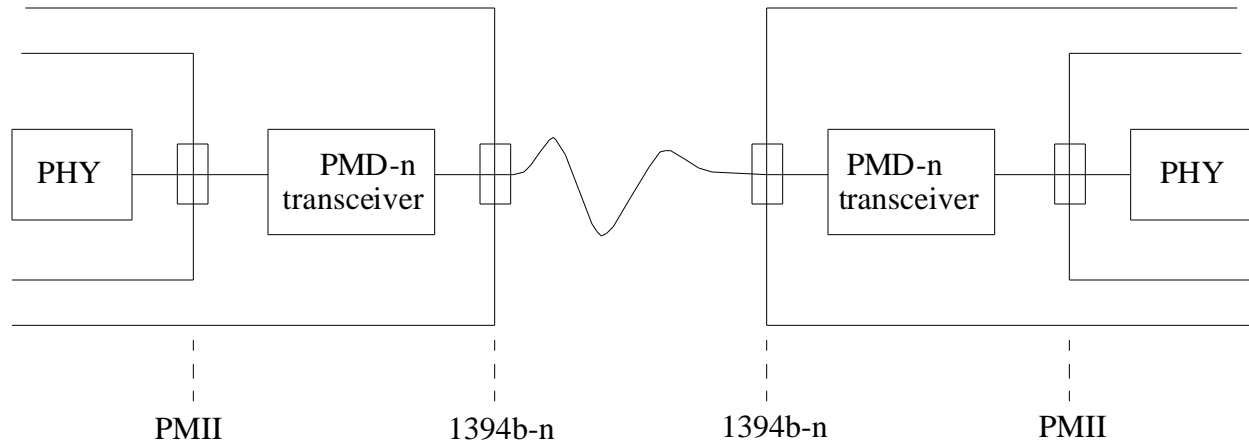
The goal is to start up in Beta mode if both ends of a connected link are Beta capable, otherwise start up in DS-Link mode. Normal bus reset, tree-ID, self-ID and normal operation will follow AFTER this start-up procedure. The technique aims to be as sympathetic with the current 1394 philosophy as possible.

2.0 Requirements

1. To start up as a DS-Link if either end of the link is not Beta capable
2. To start up in Beta mode if both ends of the link are Beta capable
3. To start up in Beta mode at the maximum Baud rate which is acceptable to both ends (this assumes a single speed of operation in terms of Baud rate, but with some form of padding to support lower speed data rates.)
4. To support the higher speed copper links being defined with P1394b, including speed determination/negotiation
5. To support a range of transceivers, which may have lower maximum capable speed than the PHY device (see model below)
6. To start up with no user intervention or configuration
7. To allow existing industry-standard transceiver interfaces to be exploited.

3.0 Model

The following reference model is proposed, to define the interoperability interfaces and as a basis for understanding the relationships between the PMD and P1394b specifications.



This model is intended to carry the following implications and non-implications:-

1. A new reference point is defined, called PMII (Physical Media Independent Interface). The current P1394b work defines everything towards the PHY from this interface, the new PMD subgroup defines everything towards the medium from this interface.
2. There is one specification for PMII, but a range of specifications for 1394b-n (n is used to distinguish one from another).
3. The PMD-n transceivers are unintelligent and unconfigurable. The role of a transceiver is to convert between the signals provided by the PMII and a suitable format for long distance transmission.
4. The PMII electrical specification is the specification currently being defined within P1394b for S800 and upwards short copper links. In this case, the PMD transceiver is null, and the PMII = 1394b-N, (where N is the value of n which identifies this interface).
5. The PMII may or may not be connectorised. When connectorised, it uses the standard 1394 connector.
6. The following packaging options are permitted (non-exhaustive list):-
 - The external interface of the device is the current 1394 connector. A short 1394 copper cable leads to a dongle, which converts to an appropriate optical interface and a short optical cable is terminated in an optical connector.
 - As above, with the optical cable replaced by a long-haul copper cable (e.g. UTP-5).
 - The external interface of the device is the current 1394 connector. A 1394 copper cable leads to an active wall plate with another 1394 connector, the transceiver, and connections to installed fibre (or long-haul copper)
 - The external interface of the device is an appropriate optical or long haul copper connector. The long reach transceiver is internal to the device, the PMII is not connectorised and may be subsumed within an integrated implementation.
7. The speed capabilities of the PHY may exceed the speed capabilities of the PMD transceiver. The PHY needs to discover the bandwidth capabilities of the channel with no user involvement.
8. The start-up procedure is therefore defined across the PMII.

4.0 Proposed procedure

The 1394-1995 interface uses both the TPA and TPB pair for transmission and reception of data, as well as for digital DC arbitration signalling. Common mode signalling is performed with respect to a bias voltage generated on TPA. This voltage is sensed on TPB as a connection indication. Current is also drawn at TPB which is sensed as a voltage drop at TPA for speed signalling.

It is required that the PMII supports these facilities unchanged when the PHY is operating as a DS link.

In Beta mode, data is transmitted on TPB and received on TPA. In addition, the potential use for optical transmission implies either (i) there is no DC signalling across the PMII, or (ii) more intelligence has to be built into the transceiver. The second option can be held in reserve, but the first option is to be preferred.

The power-on start-up procedure for a Beta-capable PHY proceeds through the following states

1. Determination of the connection status (the other end is a DS-only link, a Beta-capable link, or disabled or unconnected)
2. (If both ends are Beta capable) speed determination, bit synchronisation at the agreed speed and symbol alignment
3. Tree-ID etc.

After state 2, normal operation of the link is continuous. A procedure is defined for recovery from loss of synchronisation, a procedure is defined to maintain a connection in a low-power mode, where transmission is not continuous, and a procedure is defined to determine disconnection. After disconnection, the PHY enters an “await connection” state (see below).

4.1 Connection status

This state determines whether:-

1. the link is connected to an enabled port on a DS-Link only device, if so, start-up is as per P1394a
2. the link is connected to an enabled port on a Beta capable device, if so, enter Beta Mode speed determination, see below
3. the port is unconnected, or connected to a disabled port, so enter an “await connection” state

The use of TpBias is restricted to its use as specified for 1394-1995 DS Link signalling. A Beta capable link does not, initially, generate TpBias on TPA. In Beta Mode, TPA will always provide termination for reception of the high-speed data, and will not supply an electrical signal.

On power on, a Beta capable PHY should wait for the TpBias debounce time, and then check for TpBias (on TPB) on each of its ports. If TpBias is sensed, this implies that the port at the “other” end is a DS-Link only device and is enabled. The PHY generates TpBias on TPA and the port starts up in DS mode as per P1394a (reset, Tree-ID, Self-ID and normal operation).

If TpBias is not sensed, then a sensing tone is transmitted (e.g. 125 MHz), and listened for. If, at any time, TpBias is sensed, then this procedure is abandoned (this will occur if a DS-only device is switched

on and its port enabled just after the Beta-capable device is turned on). The tone is transmitted for a defined period of time - long enough to overcome the time constants in the PMD. 5.0 usec is proposed, but needs to be validated by the PMD group as being sufficient to overcome the time constants in all anticipated PMDs.

If a tone is received, then this confirms that the “other” end is Beta capable. The PHYs at both ends enter speed determination state.

Care must be taken when specifying the frequency of the tone that it will be low enough in frequency to pass through the slowest PMD, but not so low as to be attenuated by the pass-band of a very high frequency PMD.

4.2 Speed determination, bit synchronisation and symbol alignment

This is largely still open, and requires further study. The PHYs can exchange information giving their maximum operating speed, but this does not deal with the problem of a slower transceiver. A possible method which includes a crude transceiver test is described below.

Three symbol sequences are defined - S1, S2 and S3. For ease of understanding, consider each sequence as four 10-bit symbols. The choice of symbols depends on the coding scheme to be adopted. Each sequence has the property that the PHY can perform symbol alignment by examination of the bit-patterns within the symbols.

The ability to transmit and receive at a lowest common denominator speed is established first.

The sequence S1 is transmitted repeatedly at the “S100-equivalent” rate (122.88 MBaud). Simultaneously, the sequence is listened for. The receiver performs symbol alignment until it can receive the sequence successfully.

The exchange of S1 continues for long enough to provide a sufficient level of confidence in the link signalling integrity. The PHY then transmits S2 repeatedly to confirm “good” reception, and awaits reception of the S2 sequence. If the PHY is capable of a higher speed, it then transmits S3 repeatedly, otherwise it continues to transmit S2. If it is capable of a higher speed, and it receives the S3 sequence, then it starts to transmit the S1 sequence at the S200 equivalent rate, and attempts to lock onto reception at this rate. Upon successful reception for a defined period of time, the PHY confirms reception by changing the transmitted pattern to S2. The process repeats to the next highest speed.

In the event that transmission is not successfully received, then the PHY reverts to the lower speed. Care needs to be taken to deal with the case when the transmission is OK in one direction but not the other.

The result of this procedure is bi-directional transmission at the maximum rate acceptable to the PHYs, the transceivers and the medium. Bit alignment and symbol alignment is achieved.

Upon completion of the procedure, the PHY commences to transmit IDLE symbols according to the coding scheme being used. The PHY also sets the Bias bit (simulating TpBias for the PHY-LINK interface specification), with the consequent implications for the Con bit.

4.2.1 Choice of sequences

The properties required of a sequence are

1. Each sequence should be DC balanced
2. Each sequence must provide suitable patterns for symbol alignment

We have not examined the use of 4B5B coding, but anticipate that a suitable set of sequences can be defined.

In the case of the IBM 8B10B code, then the Fibre Channel primitive sequences OLS (for S1), LR (for S2) and LRR (for S3) are suitable. Note that by using these sequences, the procedure is similar to the Fibre Channel start-up procedure. However, it differs as Fibre Channel does not have to provide a speed determination function.

In the case of the Coles 8B10B code, then S1 can have the form SC1, SC2, SC1, SC2; S2 can have the form SC1, SC1, SC2, SC2; and S3 can have the form SC1, SC2, SC2, SC2. (SC1 is the bit pattern 1010101010, and SC2 is the bit pattern 0101010101). Note that SC1 is also suitable as the tone used during Connection Status.

4.3 Await connection

The Beta capable PHY has received neither a TpBias signal nor a tone.

It transmits a tone for a period of n microseconds (we propose n=5). This is repeated at intervals of 60n (say) microsecond until either a tone is heard or TpBias is detected on TPB. If a tone is heard, then the speed determination procedure is entered (see above). If TpBias is detected, then TpBias is generated on TPA and start-up is as per 1394-1995 for DS-Links.

This scheme has been chosen done to avoid excessive energy being transmitted into an unterminated cable. This would be necessary for a UTP 1394 link, but could also be a sensible move for any 1394b copper link.

A further consideration is the use of “sleep” modes, and the need to be able to detect a wake-up signal without requiring the PHY to be powered all the time. This may be achievable either by (i) a “signal detect” signal from the transceiver (communicated to the PHY by using some form of TpBias-type signal) or (ii) by having a low level active circuit which can respond to the reception of the tone and wake up the rest of the PHY. The second option is preferable.

5.0 Loss of synchronisation

This is a topic for further study. However, a good starting point is the Fibre Channel state machine, which provides hysteresis on the detection of an error (e.g. an invalid symbol, or excessive running disparity), such that synchronisation is maintained unless a rapid sequence of errors is detected. The PHY transmits an indication of loss of synchronisation (with the IBM 8B10B code this could be the NOS primitive sequence, with the Coles 8B10B code this could be the S1 sequence). The PHY enters

the speed determination state on loss of synchronisation and on receipt of the loss of synchronisation indication.

6.0 Analysis

We analyse various scenarios, to ensure that the system will work with DS-Link only PHYs and Beta capable PHYs.

6.1 Local Beta Capable to Remote DS-only - local on first

The local PHY sees no TpBias, and receives no tone. It drops into Await Connection status. When the remote PHY is powered on/enabled, it transmits TpBias. The local PHY enters DS link mode. The remote PHY may see a tone on TPA when it powers on. This could appear as RX_REQUEST or RX_CHILD_HANDSHAKE. However, the remote PHY should ignore these as noise, as the tone will be removed by the local PHY before it asserts TpBias.

6.2 Local Beta Capable to Remote DS-only - remote on first

The local PHY immediately (after the TpBias debounce time) sees TpBias when turned on, enters DS-Link mode and turns on TpBias. The remote PHY senses the TpBias generated by the local PHY.

6.3 Local Beta Capable to Remote DS-only - both on, connection made last

The local PHY will be in the Await Connection state. As soon as the connection is made, the local PHY immediately (after the TpBias debounce time) sees TpBias, ceases transmission of the tone (if necessary), and enters DS-Link mode. The remote PHY senses the TpBias generated by the local PHY. Prior to this, the remote PHY may first see a tone on TPA. This could appear as RX_REQUEST or RX_CHILD_HANDSHAKE. However, the remote PHY should ignore these as noise, as the tone will be removed by the local PHY before it asserts TpBias.