

Point-to-Point Fiber Optic Links

National Semiconductor
Application Note 679
Filipe Sanna
Louise Yeung
April 1990



TABLE OF CONTENTS

1.0 POINT-TO-POINT APPLICATIONS

2.0 SYSTEM OVERVIEW

3.0 CHANNEL SYNCHRONIZATION

3.1 Synchronization Timing Examples

4.0 PHY LAYER COMPONENTS

4.1 System Block Diagram

4.2 Channel Block Diagram

5.0 HOST INTERFACE CONSIDERATIONS

5.1 Data Interface

5.2 Control Bus Interface

6.0 TIMING BUDGET FOR WIRED AND STRUCTURE

6.1 Pull-Up Scheme

6.2 GAL Scheme

INTRODUCTION

A station using the ANSI X3T9.5 (FDDI) physical layer standard can transmit and receive data at 100 Mbits/sec through a fiber optic cable. However, with several physical layers connected together in parallel, each with its own fiber optic cables for transmission and reception, the station can transmit and receive data at much higher speeds. National Semiconductor's physical (PHY) layer devices can be connected together in parallel to achieve such high bandwidth point-to-point links. The National devices required to implement a PHY layer are the DP83231 Clock Recovery Device (CRD™ device), the DP83241 Clock Distribution Device (CDD™ device), and the DP83251/55 Physical Layer Controller Device (PLAYER™ device). The bandwidth of the system depends on the number of PHY layers used—each set of PHY layer devices contributes 100 Mbits/sec to the system.


1.0 POINT-TO-POINT APPLICATIONS

The use of parallel FDDI PHY layers is a cost effective method to increasing the data throughput in multiples of 100 Mb/s. This task can be accomplished utilizing an existing FDDI fiber plant.

FDDI PHY layers in parallel require only one pair of fiber optic cables, one pair of transceivers, and one set of PHY layer chips per channel, while yielding a typical data throughput of 800 Mbits/sec (for a system with eight channels).


Any application where data throughput is the limiting factor to system performance is a candidate for a high-speed point-to-point link. For example, a point-to-point link can be installed between a CPU and a disk controller to speed up information storage and retrieval times. Another application area could be in the display capabilities of a graphics workstation which can be combined with the data processing power of a supercomputer to achieve visualization for data intensive simulations (*Figure 1*). As a third example, a high-speed networking backbone using a point-to-point link is depicted in *Figure 2*. Here, separate FDDI rings are connected together with a high speed link which provides bridging between rings without loss of performance.

Fiber optics afford a greater physical separation between stations than electrical signals. Hence, a fiber point-to-point link can be used to extend SCSI or IPI transmissions up to one kilometer.



TL/F/10797-1

FIGURE 1. Point-To-Point Link between a Supercomputer and a Graphics Workstation



TL/F/10797-2

FIGURE 2. High Speed Networking Backbone Using a Fiber Optic Point-To-Point Link

BMACTM, PLAYERTM, CDDTM, and CRDTM are trademarks of National Semiconductor Corporation.

2.0 SYSTEM OVERVIEW

A system using four PHY layers in parallel is shown in *Figure 3*. The diagram demonstrates conceptually how data is passed from Station I to Station II at 400 Mbits/sec. Each of the four PHY layers in Station I is connected to a PHY layer in Station II with fiber optic cables. Since National Semiconductor PHY layers are full duplex, each pair of PHY layers is linked by two fibers, one for transmission in each direction. The system interface shown contains two parts. SYS_REQ, which handles data transmission, and SYS_IND, which handles data reception.

Suppose that Station I wants to transmit a 32-bit word of data to Station II. SYS_REQ in Station I takes the data and splits it into four bytes, one for each PHY channel. Each PHY layer reads its byte from the PHYn_REQ ports of SYS_REQ (at 12.5 MHz) and sends the data out across the fiber as an 8-bit serial stream at 100 Mbits/sec. (Note that due to the 4B/5B encoding scheme used by the PLAYER devices the data actually passes through the fiber as a 10-bit serial stream at 125 MHz.)

After travelling through the fibers, the data arrives at Station II. Each PHY layer on the receiving end reads data from the fiber and presents its byte to the corresponding PHYn_IND port (at 12.5 MHz) in Station II. SYS_IND then rejoins the bytes back into the 32-bit word sent by Station I and can present the data to the host (at 12.5 MHz). This demonstrates how Station I can send 32 bits to Station II in 80 ns, giving an effective data throughput of 400 Mbits/sec.

Even though this is a non-FDDI application, the general rules for FDDI framing must be followed. In particular, each frame must start with a JK symbol and end with valid FDDI ending delimiter (*Figure 4*). Furthermore, the frame size must be between three and 4500 bytes long (see PLAYER device datasheet for more detail). At least four pairs of idle symbols should be inserted between the frames to allow for readjustment of the PLAYER device's elasticity buffer. However, to guarantee at least one opportunity to recenter the elasticity buffer between frames in the event of clock drift or a single line hit in the interframe gap, the user is advised to insert eight idle symbol pairs.




FIGURE 3. System Using Four PHY Layers in Parallel

TL/F/10797-3




FIGURE 4. Valid Frame Format

The FDDI standard specifies the maximum clock drift between two stations to be 100 parts per million (ppm) peak-to-peak and the maximum frame size to be 4500 bytes. However, it is possible to transmit more than 4500 bytes per packet if tighter clock tolerances are observed. The equation to determine the maximum allowable frame size is given below:

$$\text{Frame}_{\text{MAX}} = 40 \text{ ns} \frac{1}{\text{Clock Drift per Bit}}$$

The clock drift per bit is calculated by taking the maximum difference in the CDD device frequencies between the transmitting station and the receiving station multiplied by 8 ns. For example, if the 12.5 MHz crystal used by the CDD device in the transmitting station had a +25 ppm accuracy error and the 12.5 MHz crystal used by the CDD device in the receiving station had a -30 ppm accuracy error, then the clock drift per bit would be 55 ppm (peak-to-peak) x 8 ns or 4.4 ps. The maximum frame size that could be transmitted using these crystals is thus 40 ns/4.4 ps or 9090 bytes, significantly more than the 4500 specified.

3.0 CHANNEL SYNCHRONIZATION

Since the data traveling from Station I to Station II passes through different devices and different fibers, it may not arrive at each of the PHY channels on the receiving end at exactly the same time. Hence there must be some method of aligning the incoming bit streams so that data is passed to the PHYn_IND ports in the correct sequence. The JK symbol serves as the reference point for synchronizing the bit streams among channels. National Semiconductor PLAYER devices have an open drain output called Cascade Ready (CR-pin 48) which is released when a JK is received on that channel. ANDing the Cascade Ready (CR) pins of all the PLAYER devices creates a signal (called Cascade Start) which indicates that all of the channels have received a JK symbol. This signal is tied to the Cascade Start (CS-pin 47) input of each PLAYER device and indicates when all of the devices achieve synchronization.

The first PLAYER device that receives a JK symbol pair will present that pair to the host (through the A Indicate Port). Meanwhile, it will activate the open drain CR output. If needed, it will output a second JK to the host as it waits for synchronization from the other PLAYER devices. During this time, the incoming data can be temporarily stored in the elasticity buffer. The ability of the PLAYER devices to output two consecutive JK symbols yields an 80 ns synchronization window. Each PLAYER device that receives a JK symbol will present a JK symbol to the host and release its CR line.

Once all of the PLAYER devices have released their CR lines, the CS signal feeding each PLAYER device will go high. At this point the read pointers of all the PLAYER device's elasticity buffers will be aligned and all of them will output JK symbols to the host. Simultaneous reception of JK symbols on every channel informs the host that synchronization has occurred, and that the subsequent data bytes will be properly aligned.

The synchronization process is repeated with each new frame received, and may not always proceed exactly as described above. Depending on the skew between the fastest and slowest channels, the PLAYER devices will either synchronize the bit streams or generate an error. Figure 5 shows four different scenarios for the synchronization of several PHY channels. Each is described in the following section.

3.1 Synchronization Timing Examples

Figure 5a presents the timing waveforms for a single PLAYER device which was the first device to receive a JK at the Elasticity Buffer (EB) in a situation where the last JK received by another PLAYER device is 60 ns behind it. PMD_IND shows a simplified version of the data coming into the PLAYER device from the CRD device. The Local Byte Clock (LBC) is a 12.5 MHz TTL signal from the CDD device which is used by the PLAYER device. After a propagation delay, data appears at the A Indicate Port of the PLAYER device on each rising edge of LBC. The CR line is pulled high on the first falling edge of LBC after the PLAYER device completely receives the JK symbol. The delay in this signal (t_1) is due primarily to the external propagation delay of the CR line pullup structure and some delay time within the PLAYER device itself.

The important point to note in this scenario is that the CS signal does not go high for some time (t_2) after the CR signal goes high. This indicates that this PLAYER device was one of the first PLAYER devices in the system to receive its JK symbol, since CS = 1 only when all of the PLAYER devices have received a JK symbol. Since the CS signal failed to go high before the first falling edge of LBC after the CR line is released, this PLAYER device outputs a second JK symbol to the host through the A Indicate Port. By the next falling edge of LBC, the CS signal has gone high and all of the PLAYER devices outputting the first byte of data after the JK symbol to the host.

Figure 5b depicts the timing waveforms for a single PLAYER device which was the first device to receive a JK symbol in a situation where the last JK symbol received from another PLAYER device is less than 20 ns after the first JK symbol. Again, PMD_IND shows the beginning of a frame coming from the CRD device. The PLAYER device releases the CR line upon reception of the JK symbol and after a short delay (t_3), the CR signal goes high. In this scenario, however, the CS signal goes high within one LBC period, so that the PLAYER device shown only has to report one JK symbol to the host before outputting data. This indicates that all of the symbols are coming in with a skew of less than 40 ns between the slowest and fastest channels. The delay between the release of CR and the assertion of CS (t_4) depends on skews in LBC between the channels, the reaction time of the wired-AND structure used to create the CS signal, and the skew between the data coming in on the different channels.

Figure 5c demonstrates the timing waveforms for a channel which receives the JK symbol pair last. Here, the CS signal goes high immediately after the PLAYER device releases the CR line. The only delay (t_6) is due to skews in LBC between the channels and the reaction time of the external wired-AND structure. Since this PLAYER device senses the CS signal within the first falling edge of LBC, it only needs to output one JK symbol to the host before outputting the data stream.

Figure 5d shows an error situation where one or more of the channels never receives a JK symbol. In this case, the

PLAYER device shown recognizes its JK symbol and releases the CR line. Since one of the PLAYER devices never gets a JK symbol and hence never releases the CR pin, the CS signal fails to go high. The PLAYER device shown in this diagram attempts to compensate for skew by outputting two

JK symbols to the host, but when the CS signal fails to rise the device is forced to output the incoming data and report a Cascade Sync Error (CSE) in the Receive Condition Register B (RCRB).




FIGURE 5a. First PLAYER device to receive a JK symbol in a large skew scenario. The CS signal goes high about 60 ns after the CR line is asserted. This indicates that some of the channels received JK symbols more than one LBC period after the channel shown.

TL/F/10797-5




FIGURE 5b. First PLAYER device to receive a JK symbol in a small skew scenario. The CS signal goes high shortly after assertion of the CR signal, indicating that all of the channels received JK symbols within t_4 .

TL/F/10797-6




FIGURE 5c. Last PLAYER device to receive JK symbol (in any scenario). Since this device was the last to assert its CR line, t_6 is due only to the delay of the wired-AND structure.

TL/F/10797-7




FIGURE 5d. PLAYER device which receives JK symbol in scenario where one or more channels never receive JK symbols.

TL/F/10797-8

4.0 PHY LAYER COMPONENTS

4.1 System Block Diagram

The number of PHY layers connected together in parallel is limited only by the timing budget of the CS line (explained in Section 5.2) and the timing skews between channels. As an example, a system level block diagram using four PHY layers connected together in parallel is presented in *Figure 6*. All of the PHY layers within a given station are driven with a single set of clock signals, and all are controlled and monitored by the host system through the Control Bus interface. Each channel has two dedicated fibers, one for transmission and one for reception. The full duplex architecture eliminates the need for complex handshaking between the two stations. The four channels communicate through the CR and CS signals. For simplicity, the CR lines are shown connected to a pullup resistor—a more detailed look at the connection of these pins is given in Section 6.

The global clock scheme should be arranged to minimize the skews in the clock signals between PHY layers. Smaller clock skews between channels will leave more tolerance for device skews and fiber optic variations. For further recommendations concerning the CDD device in a multiple PLAYER device environment, see the CDD device datasheet (CDD Device Driving Multiple PLAYER Devices).

Figure 7 illustrates the source of timing deviations between the channels and demonstrates the need to minimize timing skews between the channels wherever possible. In Station I, we are concerned with TXC and TBC while in Station II we examine LBC, since Station I is transmitting to Station II. The time parameters shown in the figure represent the maximum deviations in propagation delay between channels. For example, if t_1 were 10 ns, this would mean that TXC/TBC could arrive at PHY1 up to 10 ns before arriving at PHY2.

t_1 represents the skews in TXC/TBC between the channels, t_2 encompasses the skews in the PHY layer's transmitting path, t_3 represents the differential skews amongst the fibers, t_4 includes all of the skews inherent in the PHY layer's receiving path, and t_5 represents the skews in LBC between the channels in the receiving station. All of the skews together must not exceed 80 ns in order to prevent synchronization errors, and smaller total skews will provide greater stability across temperature and power fluctuations.

In a worst case scenario where all devices were badly skewed, t_2 together with t_4 yields a base 30 ns of skew between the channels. This leaves 50 ns available for differential skews in the clock signals and the fiber. It is recommended that t_1 be held to 4 ns and t_5 kept under 8 ns to prevent misclocking of the data. Hence, the maximum skew among fibers should be less than 38 ns.




FIGURE 6. PHY Layer Block Diagram for an Example System Using Four PHY Layers in Parallel

TL/F/10797-9

The following equation summarizes the tradeoff between cable length and variance:

$$\frac{l(1+v)}{s(1-w)} - \frac{l(1-v)}{s(1+w)} < 38 \text{ ns}$$


where: **s** is the speed of the signal in the cable

l is the average length of the fiber in meters

v is the variance in the length of the cable

w is the variance in the speed of the signal through the fiber

For example, for a typical FDDI fiber optic cable, $s = 1.9 \times 10^8$ m/s, $v = 0.005$, and $w = 0.001$. Solving for **l** with these parameters results in a length of 667 meters.



TL/F/10797-10

Where:

t_1 = Worst case clock skew between two PHYs

t_2 = Worst case skew between PLAYER device propagation delays

t_3 = Worst case skew between Fibers

t_4 = Worst case skew between PLAYER device propagation delays

t_5 = Worst case clock skew between two PHYs

Total skew = $t_1 + t_2 + t_3 + t_4 + t_5$

Note: Total skew must not exceed 80 ns in order to prevent synchronization errors.

FIGURE 7. Origin of skew between channels. Adding all of the skews (t_1 through t_5) gives the total possible skew for the system.

4.2 Channel Block Diagram

Figure 8 shows the components which constitute a single PHY channel (the CDD device is common to all channels so it is not shown here). The fiber optic transceivers are standard FDDI devices which translate electric signals to light pulses and vice versa. The fiber optic receiver accepts data from the fiber and sends two pairs of differential ECL signals to the CRD device, namely signal detect and data. The CRD device extracts a clock signal from the incoming data and passes a resynchronized equivalent of this data and a recovered clock signal to the PLAYER device, as well as signal detect and clock detect signals.

Within the PLAYER device, the incoming data stream is decoded (from 5B to 4B) and placed in the elasticity buffer. When in cascade mode, the elasticity buffer is used not only to absorb variations between the received clock and the local clock, but also to smooth out skews between incoming data presented to the different PHY channels. If all of the PLAYER devices receive JK symbols within 80 ns of each other and release their CR pins, then the CS pin will go high and all of the PLAYER devices will read from the first data location of the elasticity buffer. This cell contains the first byte of data received after the JK symbol. Hence, the elasticity buffers facilitate the coordination of data output between the different PHY channels. If the last PLAYER device receiving a JK does so more than 80 ns after the first


PLAYER device, then the Cascade Sync Error (CSE) bit will be set in the Control Bus register RCRB by the first PLAYER device to have recognized a JK symbol.

5.0 HOST INTERFACE CONSIDERATIONS

5.1 Data Interface

The system interface should consist of a transmit holding register and buffer for transmission and a buffer and receive register for incoming data. A state machine is required to decode the symbols coming from the PLAYER device so that only data is stored. Furthermore, a controller will be required to monitor and manage the PLAYER device through the Control Bus interface. This controller must handle the initialization of the PLAYER device and report error conditions to the host.

Each PLAYER device takes ten bits of data at the A Request Port, a pair of 4-bit data symbols plus a parity and control bit. (See the PLAYER device datasheet for the PHY__MAC byte wide interface table.) The system interface can thus generate parity and control for each PLAYER device separately and check control and parity coming from each channel. To simplify the system interface, however, the parity pins can be tied to ground and parity checking can be disabled in the Current Transmit State Register (CTSR). Parity information coming from the PLAYER device can similarly be ignored.



TL/F/10797-11

FIGURE 8. Block diagram of one PHY channel. The components shown here are repeated for each separate channel in the system.

In error situations, one or more PLAYER devices may report a Cascade Sync Error, but they may not do so simultaneously depending on when they receive JK symbols. The Cascade Sync Error (CSE) bit of the Receive Condition Register B (RCRB) will be set by each PLAYER device which receives a JK but does not sense the CS pin go high before the second falling edge of LBC from when CR was released. CS has to be set approximately within 80 ns of CR release. If a JK symbol is completely corrupted from a line hit or bad connection, the PLAYER device on that channel will not report a CSE. Only the data on the channel(s) which did not report a CSE will be corrupt, however, these are the channels which were unable to synchronize with the rest of the group. All of the PLAYER devices which receive a JK symbol (and release the CR pin) will read data from the first cell of the Elasticity Buffer. Therefore, a line hit on a single fiber will not wipe out the entire frame. The rest of the channels may still output synchronized data. This is particularly important in applications where partial data reception is still useful. For example, during screen updates in high resolution graphics systems, only one line of pixels would be lost instead of an entire block of the screen blanking out.

5.2 Control Bus Interface

If no JK symbols are corrupted, but they arrive with more than 80 ns of skew, all of the PLAYER devices will eventually report a CSE error. Hence the control microprocessor has the ability to pin point the corrupted channel or determine if the problem is due to excessive skew between the channels. Note that the Control Bus registers can be programmed to assert the interrupt (INT) pin upon detection of the CSE flag.

To place the PLAYER devices in Cascade Mode, the Mode Register (MR) must have the Cascade Mode (CM) bit set to one. The Cascade Synchronization Error (CSE) of the Receive Condition Register B (RCRB) is set to one if the CS signal fails to go high within 80 ns of recognizing the JK symbol. The RCRB also reports Elasticity Buffer errors through the EBOU bit, signaling a loss of data from the fiber. These bits must be cleared by the Control Bus controller.

When the number of PLAYER devices and the total capacitance is small, it may be possible to tie all CR pins and CS pins together and use a single pullup resistor. The lower limit of the pullup resistor is calculated as follows. The CR pins typically sink a 13 mA maximum, so the equation for the smallest resistor which should be used is:

$$R_{MIN} = V_{CC}/0.013\Omega$$

Hence for a voltage supply of 5V, the resistor value is $5/0.013 = 385\Omega$. The upper limit of the pullup resistor depends on the capacitance of the system and the number of PLAYER devices used. Restricting the timing budget (t_b) to 20 ns (worst case) for the AND function, we arrive at the following equation:

$$R_{MAX} = t_b/(m \times n) \Omega$$


where: **m** is the capacitance associated with each PLAYER device's CR line (including the IC capacitance (4 pF), the socket capacitance, and the trace capacitance) measured in pF

n is the number of PHY channels (number of cascaded PLAYER devices).

Thus if **m** is 20 pF and **n** is 2, the maximum pullup resistor is 500Ω , which meets the specification for the minimum resistor size. However, it is apparent that for many PLAYER devices a passive pullup resistor is too slow.


6.2 GAL Scheme

As a result, we recommend using external logic devices for any system with three or more cascaded PLAYER devices. The choice of devices is limited by the propagation delay as well as fan in or fan out. Each CR pin should be pulled up with a 400Ω resistor and fed into the AND gate. A recommended device to perform the AND function is the GAL16V8A chip, which offers 8 inputs and supplies 8 outputs with a propagation delay of 10 ns. This chip will allow up to eight PLAYER devices to be cascaded together while still maintaining the necessary delay, fan in and fan out characteristics. The device should be placed in the electrical center of the cascaded devices to prevent excessive timing skews among the chips. Figure 9 depicts an eight channel system using the GAL16V8A to AND the CR signals together.



TL/F/10797-12

FIGURE 9. Example of an eight channel system using a GAL16V8A chip to perform the AND function on the CR lines. The resulting signal (\overline{CS}) is fed back into each of the PLAYER devices.



LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.



National Semiconductor Corporation
 2900 Semiconductor Drive
 P.O. Box 58090
 Santa Clara, CA 95052-8090
 Tel: 1(800) 272-9959
 TWX: (910) 339-9240

National Semiconductor GmbH
 Livny-Gargan-Str. 10
 D-82256 Fürstenfeldbruck
 Germany
 Tel: (81-41) 35-0
 Telex: 527849
 Fax: (81-41) 35-1

National Semiconductor Japan Ltd.
 Sumitomo Chemical
 Engineering Center
 Bldg. 7F
 1-7-1, Nakase, Mihama-Ku
 Chiba-City,
 Ciba Prefecture 261
 Tel: (043) 299-2300
 Fax: (043) 299-2500

National Semiconductor Hong Kong Ltd.
 13th Floor, Straight Block,
 Ocean Centre, 5 Canton Rd.
 Tsimshatsui, Kowloon
 Hong Kong
 Tel: (852) 2737-1600
 Fax: (852) 2736-9960

National Semicondutores Do Brazil Ltda.
 Rue Deputado Lacorda Franco
 120-3A
 Sao Paulo-SP
 Brazil 05418-000
 Tel: (55-11) 212-5066
 Telex: 391-1131931 NSBR BR
 Fax: (55-11) 212-1181

National Semiconductor (Australia) Pty, Ltd.
 Building 16
 Business Park Drive
 Monash Business Park
 Nottingham, Melbourne
 Victoria 3168 Australia
 Tel: (3) 558-9999
 Fax: (3) 558-9998

National does not assume any responsibility for use of any circuitry described, no circuit patent licenses are implied and National reserves the right at any time without notice to change said circuitry and specifications.