



XAPP870 (v1.0) January 3, 2008

Serial ATA Physical Link Initialization with the GTP Transceiver of Virtex-5 LXT FPGAs

Author: Matt DiPaolo and Simon Tam

Summary

Serial ATA (SATA) is a high-speed serial link replacement for the parallel ATA (PATA) physical storage interface. The serial link employs a high-speed differential link that utilizes Gigabit technology and 8B/10B encoding. The advantages of the SATA interface over the PATA interface include greater speed, simpler upgradeable storage devices, and easier configuration.

SATA has several unique features uncommon in most serial protocols. For example, the SATA physical link initialization involves reset, synchronization, and speed negotiation through the use of Out-of-Band (OOB) signaling. This application note explains the techniques to support SATA initialization in the GTP transceiver of the Virtex™-5 LXT platform. A reference design is also provided to demonstrate the linkup process.

Introduction

All communications between SATA host (FPGA) and device (hard disk drive) go through two pairs of differential wires. Providing reset and speed negotiation services between the two link partners through the uninitialized and synchronized link presents a unique challenge. This is addressed by using OOB signals. OOB signals are sequences of signal bursts and common mode idles. SATA-compatible devices can detect differential voltage levels to discern the presence and absence of the data signal. By evaluating the time duration of these patterns, the SATA host and device can determine the meaning of the received OOB sequence without having an established serial link.


The SATA specification [Ref 1] defines several generations of physical layer link speeds, namely, Generation 1 and Generation 2, corresponding to 1.5 Gb/s and 3.0 Gb/s, respectively. It also allows the same physical link to operate at both Generation 1 and Generation 2 speeds. The speed negotiation protocol enables the host and the device to determine each others' highest operating link speed during the initialization phase of host to device communications.

Out-of-Band Signals

OOB signals are used to complete the initialization sequence to reset and negotiate link speed before establishing a data communication link. OOB signals are not part of the data character set, hence, the name OOB. OOB signals are essentially sequences of data bursts interspersed with idles, each of which is a differential signal level below 50 mV_{ppd} (mV peak-to-peak differential). The time durations of the data bursts are fixed, but the time durations of the idles depend on the types of OOB signals defined. Table 1 contains definitions of OOB signals, and Figure 1 shows various OOB signals.

Table 1: OOB Signal Definitions

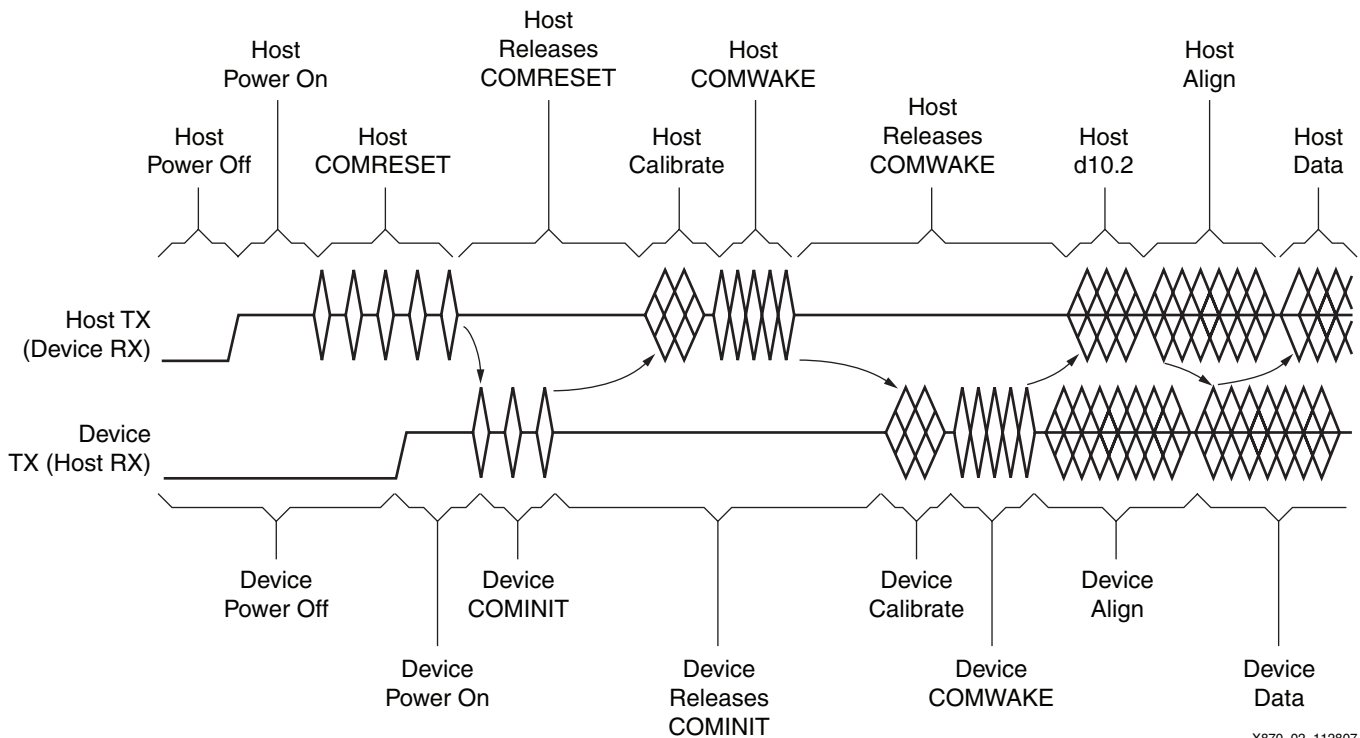
OOB Signal	Definition
COMINIT	Bursts of ALIGN primitives 160 UI (106.7 ns) long separated by idle periods at common mode level for 480 UI (320 ns). Always sent by the device.
COMRESET	Bursts of ALIGN primitives 160 UI (106.7 ns) long separated by idle periods at common mode level for 480 UI (320 ns). Always sent by the host.
COMWAKE	Bursts of ALIGN primitives 160 UI (106.7 ns) long separated by idle periods at common mode level for 160 UI (106.7 ns). Sent by both host and device.



X870_01_121107

Figure 1: OOB Signals on Link

SATA-compatible devices such as the GTP transceiver have a squelch detector that can sense the voltage level below a predetermined threshold. Therefore, the squelch detector can sense the presence and absence of the incoming signals above and below the threshold. This profile can be further deduced by logic to determine the type of OOB signals received. With this information available, control logic within the FPGA attempts to complete the initialization sequence as shown in Figure 2.



X870_02_112807

Figure 2: SATA Link Startup Sequence

The COMINIT, COMRESET, and COMWAKE signals are identical at both SATA Generation 1 and Generation 2 speeds.

Table 2 lists the initialization steps defined in the SATA specification.

Table 2: **SATA Physical Link Initialization Procedure**

Step	Description
1	Host (FPGA) and device (hard disk) power off.
2	Power is applied. Host side signal conditioning pulls TX and RX pairs to common mode voltage.
3	Host issues a COMRESET signal.
4	After the power-on reset is released, the host releases the COMRESET signal and puts the bus in a quiescent condition.
5	When the device (hard disk) detects the release of the COMRESET signal, it responds by issuing a COMINIT signal. The device can initiate communications at any time by issuing a COMINIT signal.
6	Host calibrates and issues a COMWAKE signal.
7	The device detects the COMWAKE signal on its RX pair and calibrates its transmitter (optional). Following calibration, the device sends a six-burst COMWAKE signal and then a continuous stream of ALIGN primitives. After ALIGN primitives have been sent for 54.6 μ s without a valid response from the host (i.e., no ALIGN primitives have been received), the device enters an error state.
8	After detecting the COMWAKE signal, the host starts transmitting D10.2 characters at the lowest supported rate. Meanwhile, the host receiver locks to the ALIGN primitive and, when ready, returns the ALIGN primitive to the device at the same speed as it is received. A host must be designed such that it can acquire lock within 54.6 μ s (2048 Generation 1 ALIGN DWORDs). If no ALIGN primitive is received within 880 μ s (32768 Generation 1 DWORDs), the host restarts the power-on sequence and repeats indefinitely until stopped by the application layer.
9	The device locks to the ALIGN primitive and, when ready, sends the SYNC primitive to the host, indicating that it is ready to start normal operation.
10	Upon receipt of three back-to-back non-ALIGN primitives, the communication link is established and normal operation begins.

Each GTP transceiver has built-in support to generate and detect OOB signals. This simplifies the control logic for completing the initialization. Table 3 lists the OOB ports of the GTP transceiver.

Table 3: **GTP Transceiver OOB Ports**

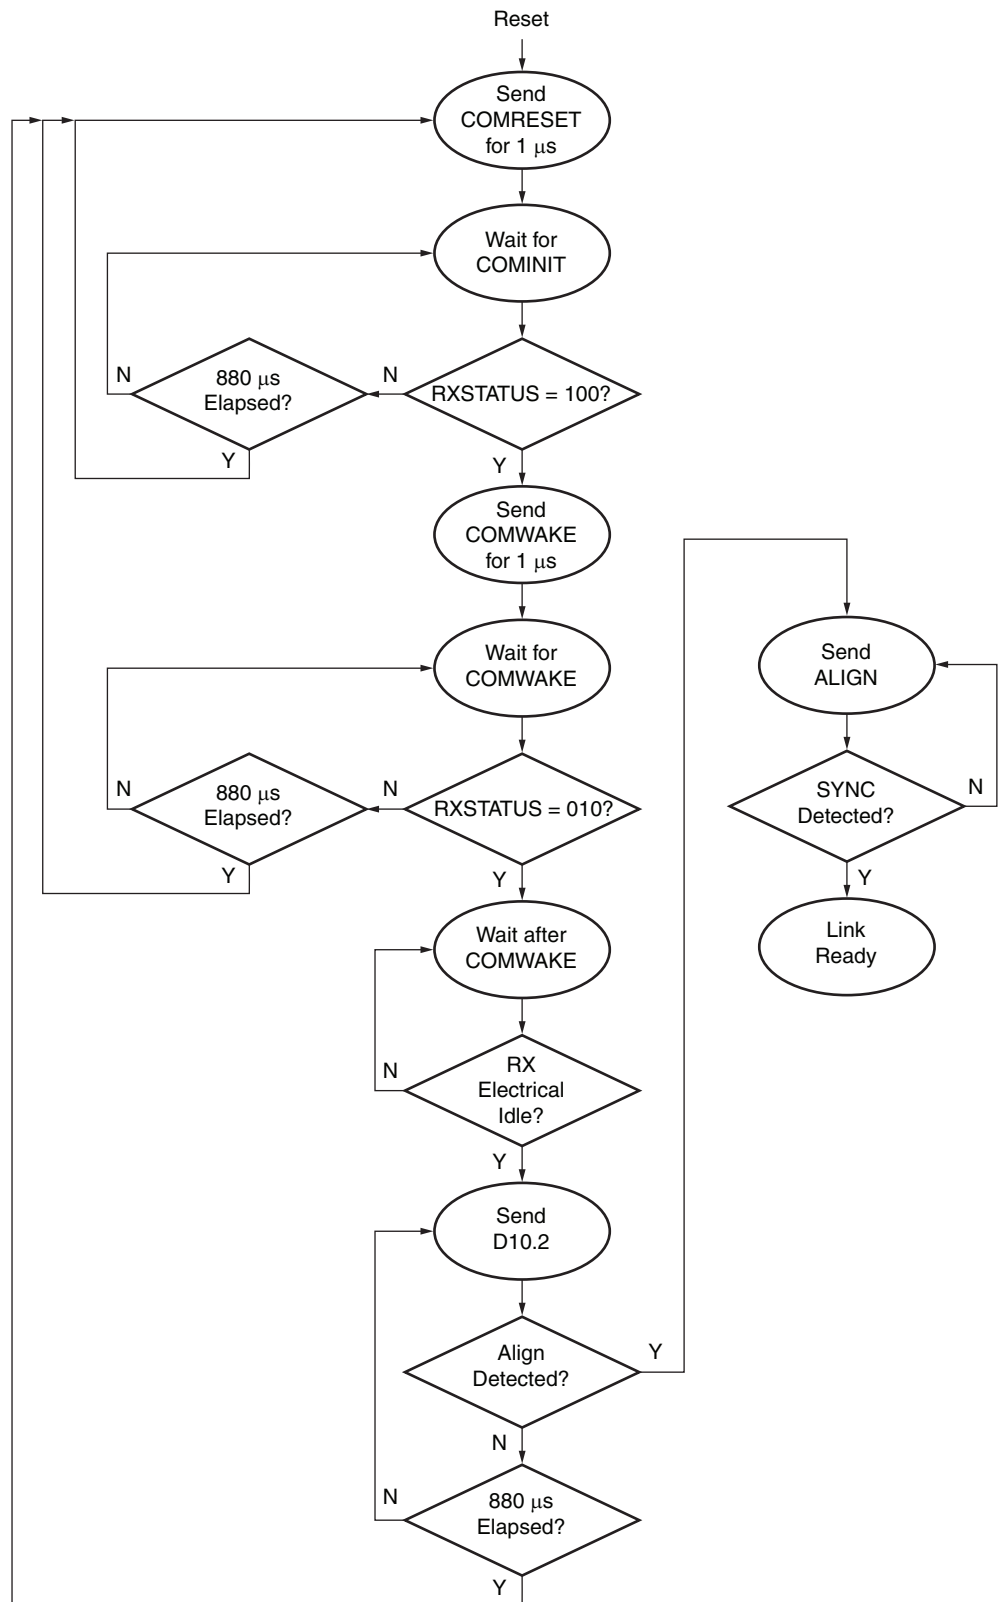
Port	I/O	Domain	Description
TXCOMTYPE0 TXCOMTYPE1	I	TXUSRCLK2	Selects the type of COM signal to send: 0: COMRESET/COMINIT 1: COMWAKE
TXCOMSTART0 TXCOMSTART1	I	TXUSRCLK2	Initiates the transmission of the COM sequence selected by TXCOMTYPE
RXSTATUS0[2:0] RXSTATUS1[2:0]	O	RXUSRCLK2	The decoding of RXSTATUS[2:0] depends on the setting of RX_STATUS_FMT. When RX_STATUS_FMT = SATA: RXSTATUS[0]: TXCOMSTART operation complete RXSTATUS[1]: COMWAKE signal received RXSTATUS[2]: COMRESET/COMINIT signal received

In addition, OOB functions in the GTP transceiver can be customized through various attributes. [Table 4](#) shows the default values of different OOB attributes.

Table 4: GTP Transceiver OOB Attributes

OOB Attributes	Setting	Description
COM_BURST_VAL	0101	Number of COM sequence bursts transmitted.
OOB_CLK_DIVIDER	6	Squelch clock divider setting. The quotient of reference clock/OOB_CLK_DIVIDER should be close to 25 MHz.
OOBDETECT_THRESHOLD	111	Optimal OOB detect threshold level. Refer to <i>the Virtex-5 FPGA Serial ATA Generation 2 Protocol Standard Characterization Test Report</i> . [Ref 2]
RX_STATUS_FMT	SATA	RX status encoding.
SATA_BURST_VAL	100	Number of bursts required to declare a COM match.
SATA_IDLE_VAL	100	Number of idles required to declare a COM match.
SATA_MAX_BURST	7	Maximum OOB burst reject threshold.
SATA_MAX_INIT	22	Maximum COMRESET idle time allowed.
SATA_MAX_WAKE	7	Maximum COMWAKE idle time allowed.
SATA_MIN_BURST	4	Minimum OOB burst reject threshold.
SATA_MIN_INIT	12	Minimum COMRESET idle time allowed.
SATA_MIN_WAKE	4	Minimum COMWAKE idle time allowed.

[Figure 3](#) shows a diagram of the initialization state machine implemented in the initialization module (OOB_Control) as part of the reference design. It completes the OOB handshake described in [Figure 2](#) and enforces time-outs according to the SATA protocol.



X870_03_112807

Figure 3: Initialization Controller State Diagram

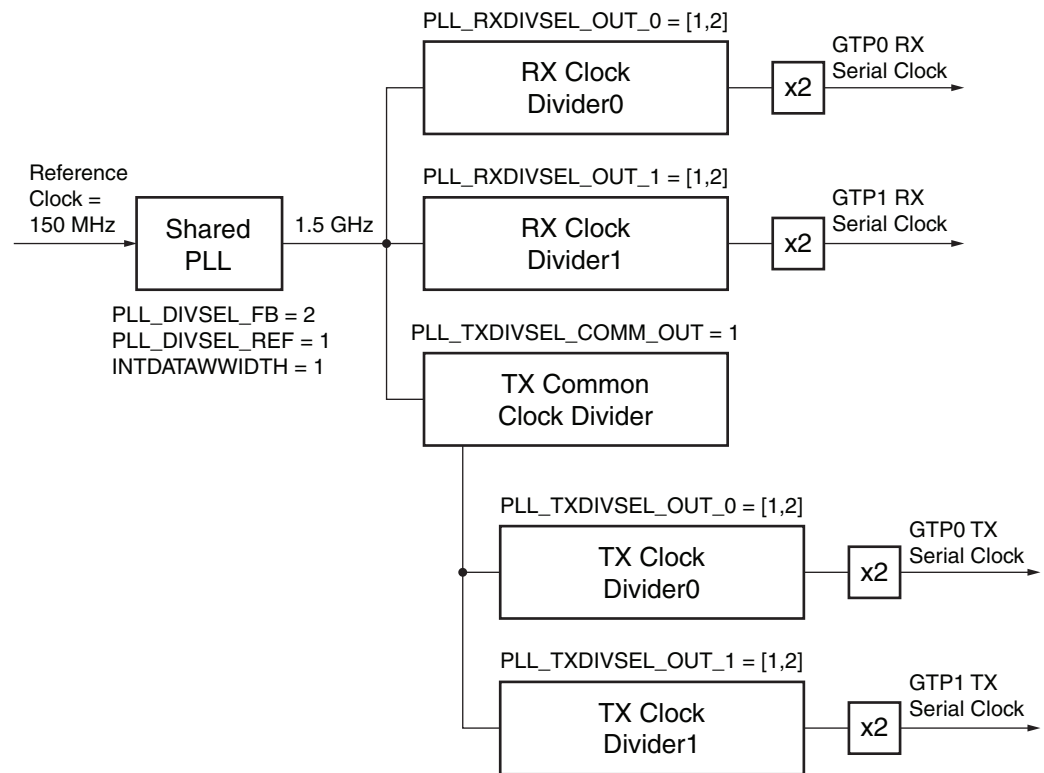
The initialization module works closely with the speed negotiation control module described in “Speed Negotiation.” After a link speed is configured by the speed negotiation control module, the initialization module attempts to establish a link with the device. The time allowed is about

3.5 ms or 4 tries. If communication cannot be established after the elapsed time, the speed negotiation control module assumes that the link rate is not appropriate and proceeds to different speeds through dynamic reconfiguration. The linkup process then repeats with the initialization module trying to link up once again with the hard disk.

If the hard disk powers up before the FPGA is ready, the hard disk sends COMINIT signals, expecting to receive COMWAKE signals from the FPGA. Powering up the hard disk before the FPGA does not cause miscommunication between the hard disk and the FPGA because the OOB controller always starts up by sending COMRESET signals. Thus, the FPGA always returns an early-started hard disk to the same reset state to ensure successful linkup.

Speed Negotiation

Speed negotiation in the GTP transceiver is accomplished by changing the internal shared phase-locked loop (PLL) divider settings of the GTP transceiver during runtime to eliminate the need for reconfiguration. The dynamic reconfiguration port (DRP) allows access to the internal attributes of the GTP transceiver through a simple bus interface. Figure 4 shows a block diagram of the shared physical medium attachment (PMA) PLL along with the dividers responsible for setting the line rate.



X870_04_112807

Figure 4: Shared PMA PLL Detail

A 150 MHz reference clock is fed to the shared PLL. The output of the PLL clock frequency is 1.5 GHz. Table 5 shows the attribute for setting the 1.5 GHz clock. Both SATA Generation 1 and Generation 2 speeds are based on this PLL clock frequency. This 1.5 GHz serial clock frequency is divided by two if the line rate is Generation 1 and divided by one if the line rate is Generation 2.

Table 5: PLL VCO Settings for 1.5 GHz

Attribute	Value
INTDATAWIDTH	1
PLL_DIVSEL_FB	2
PLL_DIVSEL_REF	1
PLL_TXDIVSEL_COMM_OUT	1

Several attributes are responsible for setting the clock dividers. The PLL dividers and their DRP addresses and values are shown in Table 6 and Table 7 for Generation 1 and Generation 2, respectively. The PLL_[T/R]XDIVSEL_OUT attribute has a two-bit value. However, bit 1 is always set to zero. Only bit 0 is manipulated to change speed. The reference design only has one of the two GTP transceivers (GTP0) configured for SATA operation. However, the technique to configure GTP1 is identical to configuring GTP0. Refer to the *Virtex-5 RocketIO GTP Transceiver User Guide* [Ref 3] for more information about the DRP and the architecture of the GTP.

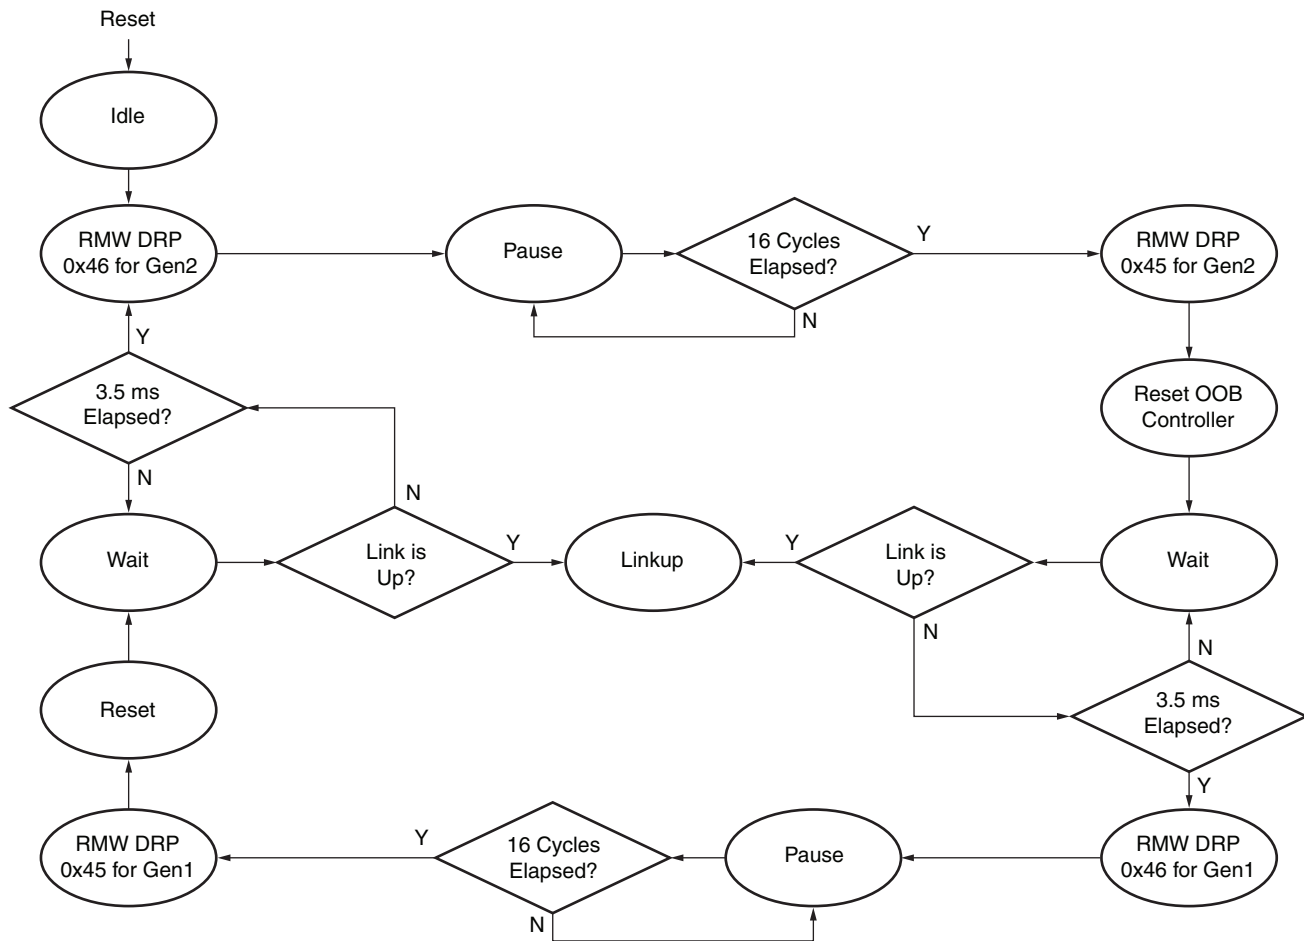
Table 6: SATA Generation 1 PLL Divider Settings

Attribute	DRP Address	Value
PLL_RXDIVSEL_OUT_0[0]	0X46[[2]	1
PLL_RXDIVSEL_OUT_1[0]	0X0A[0]	1
PLL_TXDIVSEL_OUT_0[0]	0X45[15]	1
PLL_TXDIVSEL_OUT_1[0]	0X05[4]	1

Table 7: SATA Generation 2 PLL Divider Settings

Attribute	DRP Address	Value
PLL_RXDIVSEL_OUT_0[0]	0X46[2]	0
PLL_RXDIVSEL_OUT_1[0]	0X0A[0]	0
PLL_TXDIVSEL_OUT_0[0]	0X45[15]	0
PLL_TXDIVSEL_OUT_1[0]	0X05[4]	0

The SATA link speed can be changed from Generation 1 to Generation 2 or vice versa by setting the dividers to the appropriate values through the DRP. A state machine in the logic fabric outside of the GTP transceiver needs to perform the necessary read-modify-write (RMW) transactions to the correct DRP addresses. This state machine, shown in Figure 5, is implemented in the speed negotiation control (SNC) module in the reference design.



X870_05_112807

Figure 5: Speed Negotiation State Machine

The speed negotiation state machine starts by setting the GTP transceiver in SATA Generation 2 speed first. This is accomplished by setting both the PLL_RXDIVSEL_OUT_0[0] and PLL_TXDIVSEL_OUT_0[0] to zero, assuming all the other bits are zero. The dividers are set to divide by one; therefore, the line rate is 3 Gb/s. After the dividers are set, the state machine asserts the GTP transceiver main reset to force the new divider settings to take effect and to reset the OOB linkup module. The state machine waits about 3.5 ms, after which the OOB controller tries to link up with the hard disk at 880 μ s. The SNC module allows the OOB module four attempts to link up with the partner. After the fourth attempt, the linkup is assumed to be unsuccessful. The state machine then proceeds to configure the GTP at Generation 1 speed. The method is identical to that already described for the Generation 2 speed. However, the PLL divider settings PLL_RXDIVSEL_OUT_0[0] and PLL_TXDIVSEL_OUT_0[0] are set to one. This sets the dividers to divide by two, making the link rate 1.5 Gb/s. The SNC state machine again waits about 3.5 ms for the OOB controller to link up. If the OOB controller does not link up after the elapsed time, the state machine configures the GTP transceiver to Generation 2 speed again as described earlier. This procedure loops forever until a successful link is established or the system is reset.

The SATA specification states that after sending ALIGN primitives to the host, the device expects to receive ALIGN primitives from the host within 54.6 μ s (as described in Table 2, step 8). When a GTP transceiver configured to Generation 2 speed tries to link up with a Generation 1 hard disk, the communication is not established and the device enters an error state. In this case, speed negotiation is still possible. After the SNC reconfigures the link speed to Generation 1, it resets the OOB controller to perform linkup once again and issues a COMRESET signal. According to the specification, the hard disk must enter the hardware reset

state regardless of its current state. The GTP transceiver is now at the same communication speed with the hard disk; therefore, the communication link is successful. This method of speed negotiation is allowed in the SATA specification.

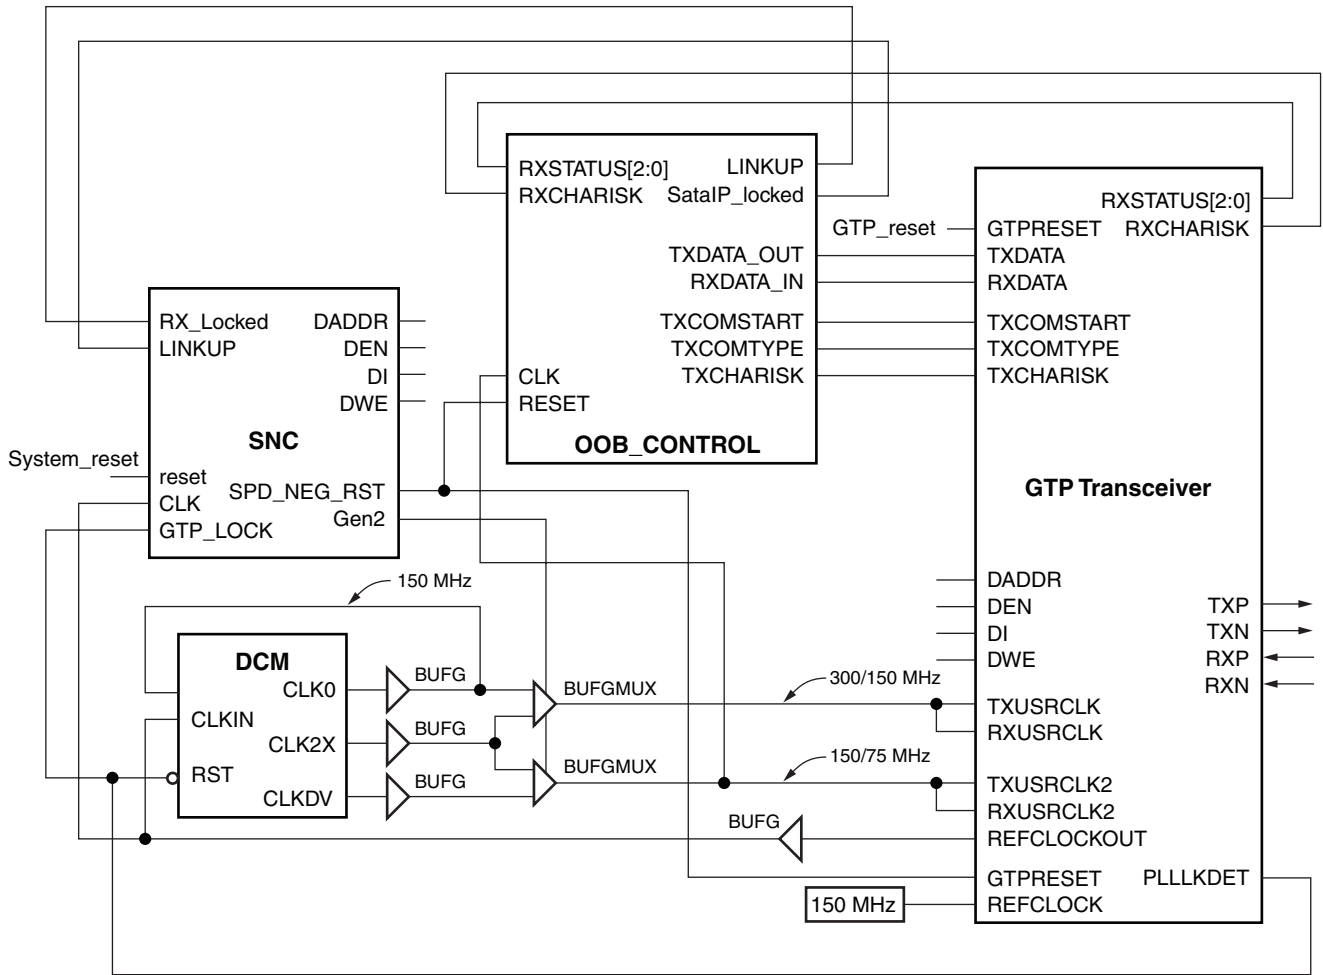
Reference Design

The reference design demonstrates how to use the GTP transceiver to complete the SATA initialization and speed negotiation. It takes the SATA link from power-up or reset through OOB handshake and speed negotiation to receive a device signature from the connected hard disk. The target board is the Xilinx ML505 demonstration board, which has two SATA connectors. For more information on the ML505 board, refer to the *ML505/ML506 Evaluation Platform User Guide*. [Ref 4] The reference design files can be downloaded from:

www.xilinx.com/support/documentation/application_notes/xapp870.zip.

Figure 6 shows the clocking scheme of the reference design. A 150 MHz reference clock is input to the GTP transceiver. The GTP transceiver user data interface is configured to two bytes wide. Therefore, the clock ratio between the USRCLK and USRCLK2 domain is 2:1. The actual clock frequencies depend on the established SATA link speed. For example, if the SATA link speed is 1.5 Gb/s, USRCLK/USRCLK2 should be 300/150 MHz. If the SATA link speed is 3.0 Gb/s, USRCLK/USRCLK2 should be 150/75 MHz. The clock multiplexing is made possible by a pair of BUFGMUXs. The SNC module controls the select pin of the BUFGMUXs and switches between the two SATA generation speeds. The 75, 150, and 300 MHz clocks are generated from a DCM. The DCM input clock is 150 MHz and comes from the reference clock output of the GTP transceiver.

The initialization control module OOB_Control generates and detects OOB signals through the OOB ports on the GTP transceiver. It follows the OOB exchange protocol described in Table 2 to correctly synchronize both the GTP transceiver and the hard disk.



X870_06_121707

Figure 6: Reference Design Clocking Scheme

According to the PATA specification [Ref 5], all devices must send a signature to the host immediately after initialization to inform the feature set supported by the host. For hard disks, the signature should be written to the command block registers as shown in Table 8.

Table 8: Hard Disk Signature

Command Block Register	Hard Disk	Non-Hard Disk
Device	00h	000x0000b
LBA High	00h	EBh
LBA Low	01h	01h
LBA Mid	00h	14h
Sector Count	01h	01h

In SATA, the device signature is encapsulated in frame information structure (FIS) format. A representative device-to-host FIS format is shown in [Table 9](#).

Table 9: Frame Information Structure Format

DWORD	Byte3	Byte2	Byte1	Byte0
0	Error	Status	Interrupt	FIS type (34h)
1	Dev/Head	LBA High	LBA Mid	LBA low
2	Features (exp)	LBA High (exp)	LBA Mid (exp)	LBA low (exp)
3	Control	Reserved	Sector Count (exp)	Sector Count
4	Reserved	Reserved	Reserved	Reserved

The hard disk signature shown in [Table 9](#) is observable in the ChipScope™ analyzer after a successful linkup. Capturing this signature is a good way to indicate that the GTP transceiver is able to handle all SATA requirements to establish a successful link.

FPGA Design Resources

The reference design was implemented on an XC5VLX50T FPGA. [Table 10](#) shows the design resource utilization. The results reflect the inclusion of the ChipScope analyzer module in the design.

Table 10: Reference Design Resource Utilization

Resource	Count
BUFG	5
IOB	18
Slice LUT	545
Slice Register	445

Reference Design Matrix

The reference design matrix for this application note is shown in [Table 11](#).

Table 11: Reference Design Matrix

Parameter	Description
General	
Developer Name	Matt DiPaolo, Simon Tam
Target Devices (Stepping Level, ES, Production, Speed Grades)	Virtex-5 LXT platform
Source Code Provided?	Y
Source Code Format	Verilog
Design Uses Code or IP from Existing Reference Design, Application Note, 3rd party, or CORE Generator™ Software?	N
Simulation	
Functional Simulation Performed?	Y
Timing Simulation Performed?	N
Testbench Provided for Functional and Timing Simulations?	Y
Testbench Format	Verilog
Simulator Software and Version	ModelSim SE 6.2g

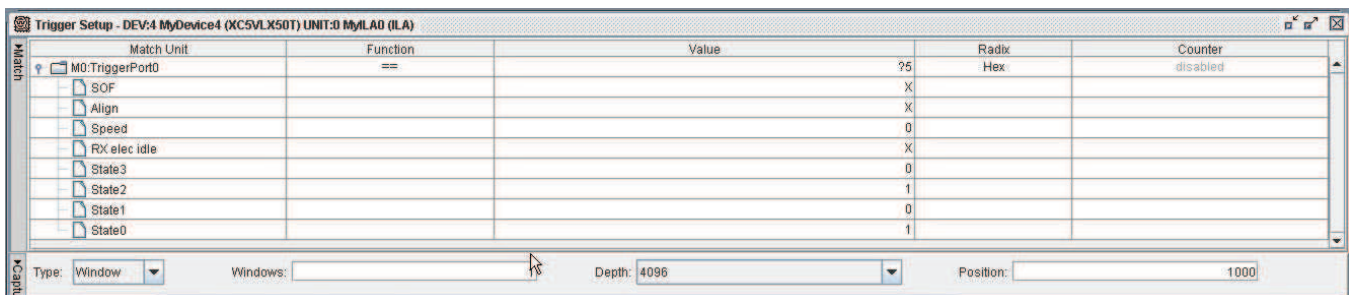
Table 11: Reference Design Matrix (Cont'd)

Parameter	Description
SPICE/IBIS Simulations?	N/A
Implementation	
Synthesis Software Tools and Version	XST 9.2i
Implementation Software Tools and Version	ISE™ software, version 9.2i
Static Timing Analysis Performed?	Y
Hardware Verification	
Hardware Verified?	Y
Hardware Platform Used for Verification	ML505

Running the Reference Design

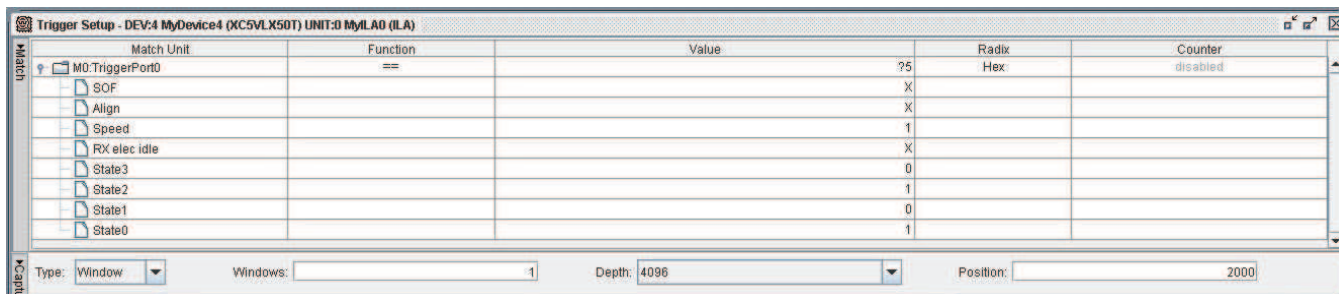
Follow these steps to inspect the linkup using the ChipScope analyzer:

1. Connect the SATA Host1 connector on the ML505 board to a SATA Generation 1 or Generation 2 hard disk through a regular SATA cable. Power up the demonstration board and the hard disk.
2. Launch the ChipScope analyzer and initialize the JTAG chain. Open the project file XAPP870.cpj in the /Chipscope directory.
3. Download the bitstream using the ChipScope analyzer. Upon completion, three GPIO LEDs at the bottom edge light up beside the DONE LED. These LEDs indicate:
 - ◆ LED2: SATA link is up
 - ◆ LED1: DCM is locked
 - ◆ LED0: GTP PLL is locked
4. Set the state trigger to 6.
5. If the hard disk is a Generation 1 device, set the speed trigger to 0. If the hard disk is a SATA Generation 2 device, set the speed trigger to 1.
6. Set the rest of the triggers to X (unknown). The trigger settings for Generation 1 and Generation 2 are shown in [Figure 7](#) and [Figure 8](#), respectively.



X870_07_121807

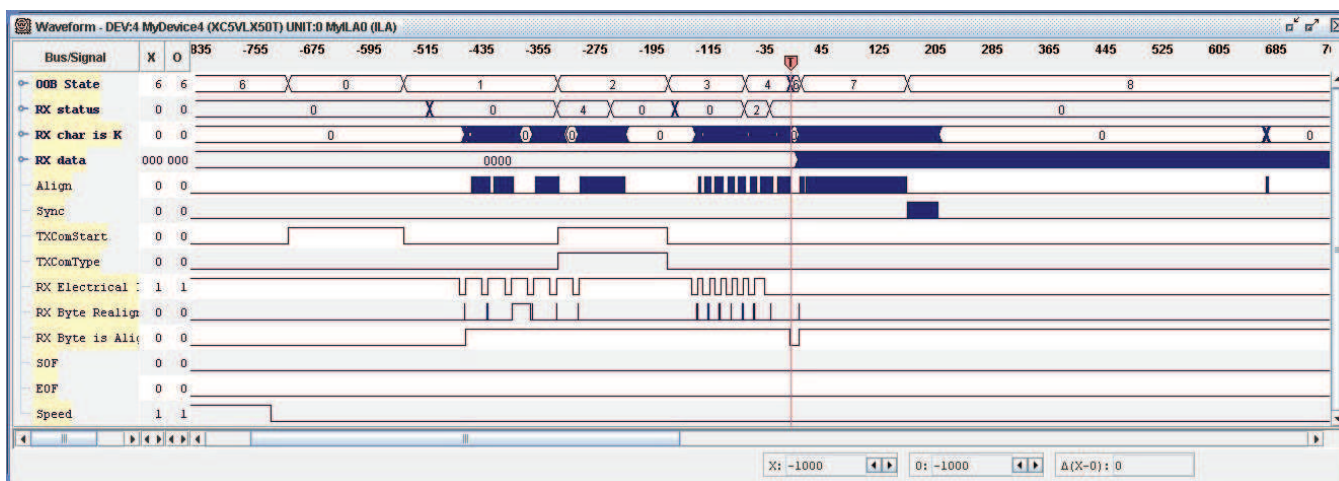
Figure 7: ChipScope Analyzer Trigger Settings for Generation 1



X870_06_121807

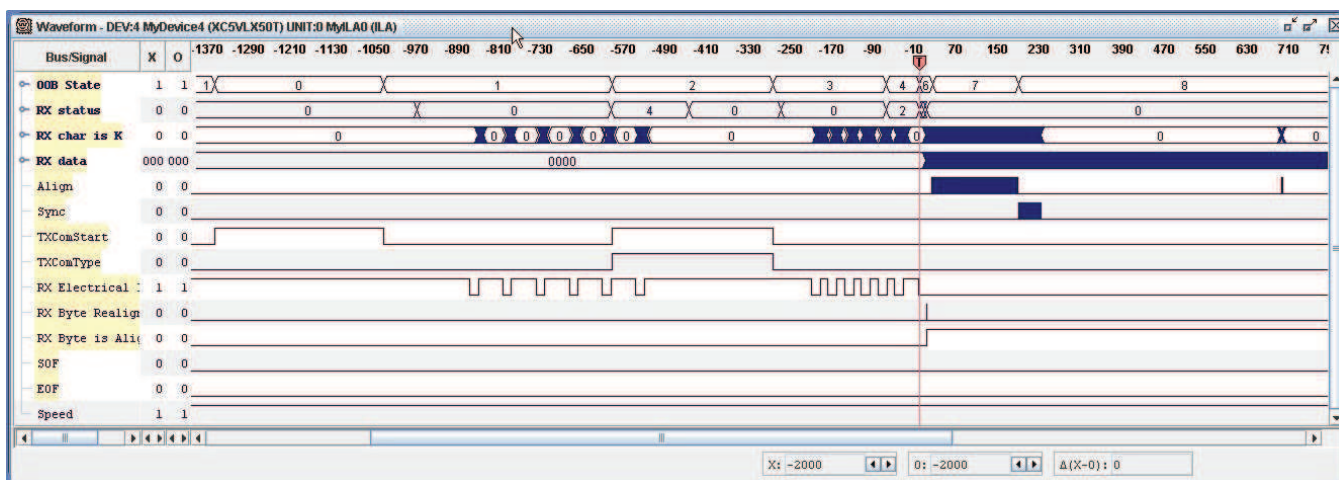
Figure 8: ChipScope Analyzer Trigger Settings for Generation 2

- Hit the center pushbutton (SW14) in the lower-right corner to start the initialization.
- For Generation 1, the screenshot shown in Figure 9 appears. For Generation 2, the screenshot shown in Figure 10 appears. The transitions on the RX Elec Idle signal correspond to the COMINIT and COMWAKE signals sent back from the hard disk. The SATA link is established when the state is 8.



X870_09_121807

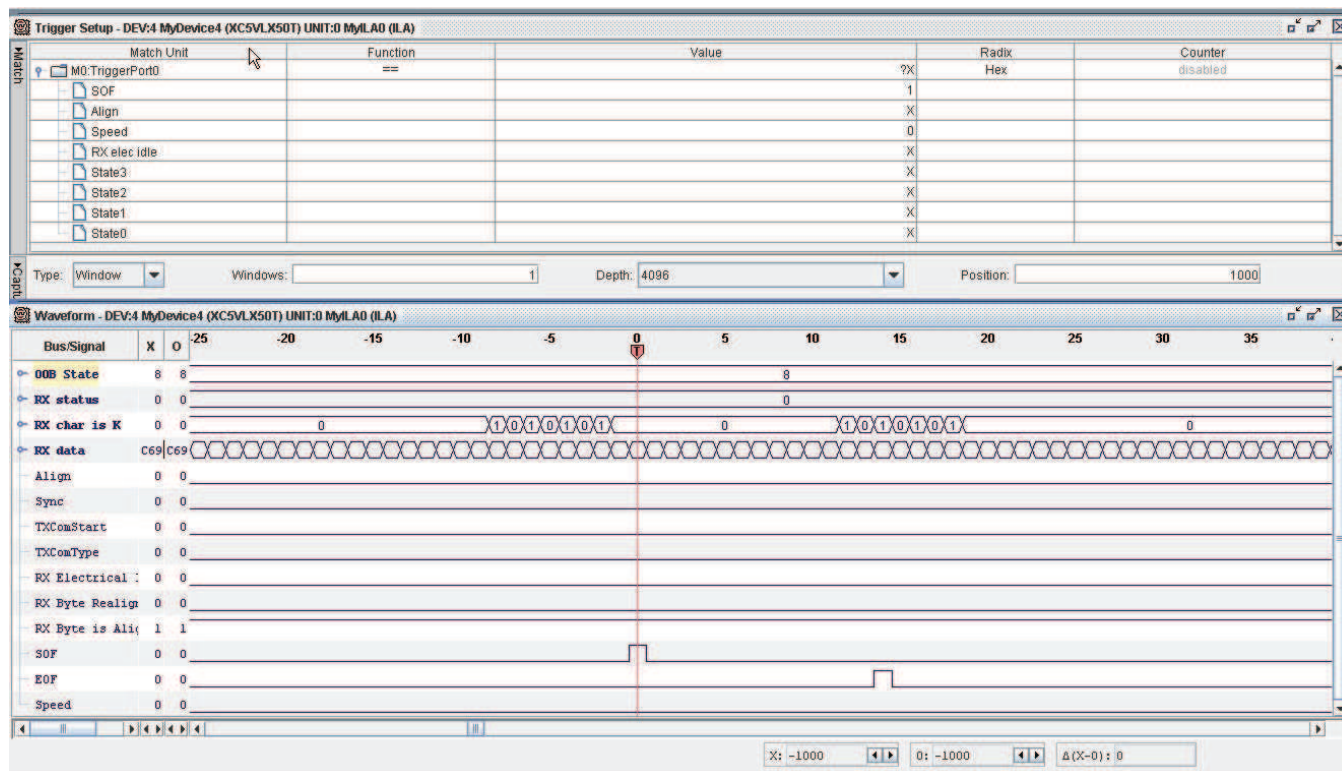
Figure 9: SATA Hardware Reset Captured in the ChipScope Analyzer for Generation 1



X870_10_121807

Figure 10: SATA Hardware Reset Captured in the ChipScope Analyzer for Generation 2

9. In the trigger window, set SOF to 1 and the rest of the signals to X.
10. Arm the ChipScope analyzer and hit the center pushbutton (SW14) of the demonstration board again. The result should look like [Figure 11](#), regardless of the link speed. The SOF and EOF assertions can be observed, indicating the start of a frame and the end of a frame, respectively. These signals correspond to the SOF (0x7C-0xB5-0x37-0x37) and EOF (0x7C-0xB5-0xD5-0xD5) primitives. Everything between the SOF and EOF primitives is an FIS.



X870_11_121807

Figure 11: Frame Information Structure

This particular FIS type is called a device-to-host register FIS and is shown in [Table 9](#). The device sends this FIS after completing diagnostics to the host to update the shadow registers. This FIS is known to contain the initialization signature. It identifies itself either as a hard disk or other storage device (CD, DVD, CompactFlash memory, etc.) and reports its operating status.

The data is scrambled before transmission. Therefore, the display in the ChipScope analyzer is not the actual data. The FIS has 20 bytes of data. The last four bytes are a cyclic redundancy check (CRC) checksum. [Table 12](#) shows the scrambled FIS and the unscrambled content. The FIS captured by the ChipScope analyzer should resemble the values in the Scrambled Data column in [Table 12](#).

Table 12: Initialization Signature

DWORD	Scrambled Data (Hex)	Unscrambled Data (Hex)
0	C28276B9	00500034
1	1F26B369	00000001
2	A508436C	00000000
3	3452D355	00000001
4	8A559502	00000000
CRC	671F9A8E	DC052495

Comparing the content with the FIS format, the signature clearly indicates that the device is a regular hard disk.

Conclusion

This application note highlights some of the important aspects of using the GTP transceiver as a SATA Generation 1 and Generation 2 physical layer device. The reference design also demonstrates that the GTP transceiver of the Virtex-5 FPGA can support these SATA features:

- 1.5 Gb/s and 3.0 Gb/s link speeds
- Speed negotiation
- Out-of-band signals
- Spread spectrum clocking

Combined with high performance programmable FPGA logic and multiple GTP transceivers, Virtex-5 LXT devices provide an ideal platform to support many SATA applications.

References

This application note uses the following references:

1. Serial ATA Working Group: *Serial ATA: High Speed Serialized AT Attachment Revision 1.0a*
www.serialata.org.
2. [RPT087](#), *Virtex-5 FPGA Serial ATA Generation 2 Protocol Standard Characterization Test Report*.
3. [UG196](#), *Virtex-5 FPGA RocketIO GTP Transceiver User Guide*.
4. [UG347](#), *ML505/ML506 Evaluation Platform User Guide*.
5. AT Attachment with Packet Interface -7 Volume 1, Revision 4b, T13 American National Standard
www.t13.org.

Appendix

Although every effort has been made to ensure the reliability of the reference design, it has not been verified or certified by any SATA governing organization. The reference design has only been tested with a limited number of SATA hard disk drives as listed in [Table 13](#).

Table 13: Disk Drives Tested with ML505 SATA Link Initialization Reference Design

Manufacturer and Brand	Speed	Density (GB)	Model Number
Hitachi Deskstar	Generation 1	250	HDS722525VLSA80
Maxtor DiamondMax 10	Generation 1	100	6L100M0
Maxtor DiamondMax 21	Generation 2	200	STM3200820AS
MTRON Solid State Drive	Generation 1	16	MSD-SATA6035
Seagate Barracuda 7200.10	Generation 2	300	ST3300620AS
Seagate Barracuda 7200.7	Generation 1	160	4MT10EES
Western Digital Caviar SE16	Generation 2	250	WD2500KS
Western Digital Raptor	Generation 1	36	WD360GD-00FLA2

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
01/03/08	1.0	Initial Xilinx release.

Notice of Disclaimer

Xilinx is disclosing this Application Note to you “AS-IS” with no warranty of any kind. This Application Note is one possible implementation of this feature, application, or standard, and is subject to change without further notice from Xilinx. You are responsible for obtaining any rights you may require in connection with your use or implementation of this Application Note. XILINX MAKES NO REPRESENTATIONS OR WARRANTIES, WHETHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL XILINX BE LIABLE FOR ANY LOSS OF DATA, LOST PROFITS, OR FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR INDIRECT DAMAGES ARISING FROM YOUR USE OF THIS APPLICATION NOTE.