



BNL-94219-2011-CP

***A deterministic, gigabit serial timing,
synchronization and data link for the RHIC LLRF***

T. Hayes, K.S. Smith, F. Severino

Presented at the 2011 Particle Accelerator Conference (PAC'11)
New York, N.Y.
March 28 – April 1, 2011

Collider-Accelerator Department

Brookhaven National Laboratory

**U.S. Department of Energy
Office of Science**

Notice: This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

This preprint is intended for publication in a journal or proceedings. Since changes may be made before publication, it may not be cited or reproduced without the author's permission.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

A DETERMINISTIC, GIGABIT SERIAL TIMING, SYNCHRONIZATION AND DATA LINK FOR THE RHIC LLRF *

T. Hayes[#], K.S. Smith, F. Severino
Brookhaven National Laboratory, Upton, NY 11973, USA

Abstract

A critical capability of the new RHIC low level rf (LLRF) system is the ability to synchronize signals across multiple locations. The 'Update Link' provides this functionality. The 'Update Link' is a deterministic serial data link based on the Xilinx RocketIO protocol that is broadcast over fiber optic cable at 1 gigabit per second (Gbps). The link provides timing events and data packets as well as time stamp information for synchronizing diagnostic data from multiple sources.

INTRODUCTION

The new RHIC LLRF was designed to be a flexible, modular system [1]. The system is constructed of numerous independent RF Controller chassis. To provide synchronization among all of these chassis, the Update Link system was designed.

The Update Link system provides a low latency, deterministic data path to broadcast information to all receivers in the system. The Update Link system is based on a central hub, the Update Link Master (ULM), which generates the data stream that is distributed via fiber optic links. Downstream chassis have non-deterministic connections back to the ULM that allow any chassis to provide data that is broadcast globally.

UPDATE LINK MASTER

The ULM is 16 layer printed circuit board mounted in a 3U, 19" enclosure. At the heart of the design is a Xilinx V5FX100T field programmable gate array (FPGA) with two embedded Power PC 440 processors (see figure 1). One of the processors functions as a front end computer (FEC) that interfaces to the RHIC control system. All the basic functionality of an RF carrier board is provided on the ULM [2].

The V5FX100T FPGA has 16 GTX tiles. GTX tiles are highly configurable serial transceivers supporting line rates up to 6.5Gbps in dedicated silicon blocks tightly coupled to FPGA resources. The GTX tiles provide 8B/10 encoding and decoding in hardware. The GTX receiver provides comma alignment and detection, has a user configurable equalization circuit to compensate for high frequency attenuation and optimizes data sampling time within the eye.

In this design, all 16 GTX transceivers are connected to one of a pair of Analog Devices AD8152 crosspoint switches. The AD8152 is a 34 x 34 channel, fully differential digital crosspoint switch supporting line rates up to 3.2 Gbps. The AD8152 also has the ability to connect an input to more than one output. Thus the

crosspoint switch becomes the buffer that allows a single GTX transmitter to be broadcast to all outputs. The crosspoint switches are connected to 34 small form factor pluggable (SFP) fiber optic transceivers. Connections between the two AD8152s allow maximum flexibility for routing signals from any GTX to any SFP. The use of SFPs allows the option of using single mode or multimode fiber depending on the distances to be covered.

The control signals for the crosspoint switches are connected to the FPGA so the switch channel assignments can be reconfigured in real time. Unused channels can be powered down to dramatically reduce power usage.

The ULM has one XMC daughter site. The high speed serial connections to this site are routed through the crosspoint switch. In this way, GTX resources are only used if the XMC is populated and otherwise can be connected to external signals.

LINK DETAILS

The serial outputs from the ULM use Xilinx RocketIO links running at 1 Gbps [3]. RocketIO is a very low latency link with no overhead. The top level interface of the transmitter presents a 32 bit bus but at the actual physical transceiver, the data is 8B/10B encoded and sent as two sequential 20 bit transmissions.

A common 100 MHz clock is supplied to ULM and all RF Controller chassis in the system in order allow the links between them to be deterministic. This is required to allow disabling the internal clock correction block and bypassing an elastic buffer that is used to pass data between two different clock domains within the GTX receiver.

There was one additional major hurdle to overcome to make the links deterministic. The receiver circuit in the GTX uses a 20 bit barrel shifter as part of the serial to parallel conversion process. Data is latched into the barrel shifter on the word clock and then shifted 'n' positions until data is lined up properly on a word boundary. Since the position of the barrel shifter is arbitrary when the link is initialized, this constitutes an unknown delay. To resolve this issue, special code was created to control the initialization of the GTX tile. The code waits until the GTX tile internal initialization is complete and then reads the barrel shifter using the Dynamic Reconfiguration Port (DRP). The DRP is a Xilinx interface that allows the FPGA to access and control certain portions of the dedicated hardware in the GTX tile. If the barrel shifter is not at the desired position, the DRP is used to temporarily unlock the internal PLL that generates the parallel data clock from the bit rate clock derived in the Clock Data Recovery

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. Department of Energy
[#]hayes@bnl.gov

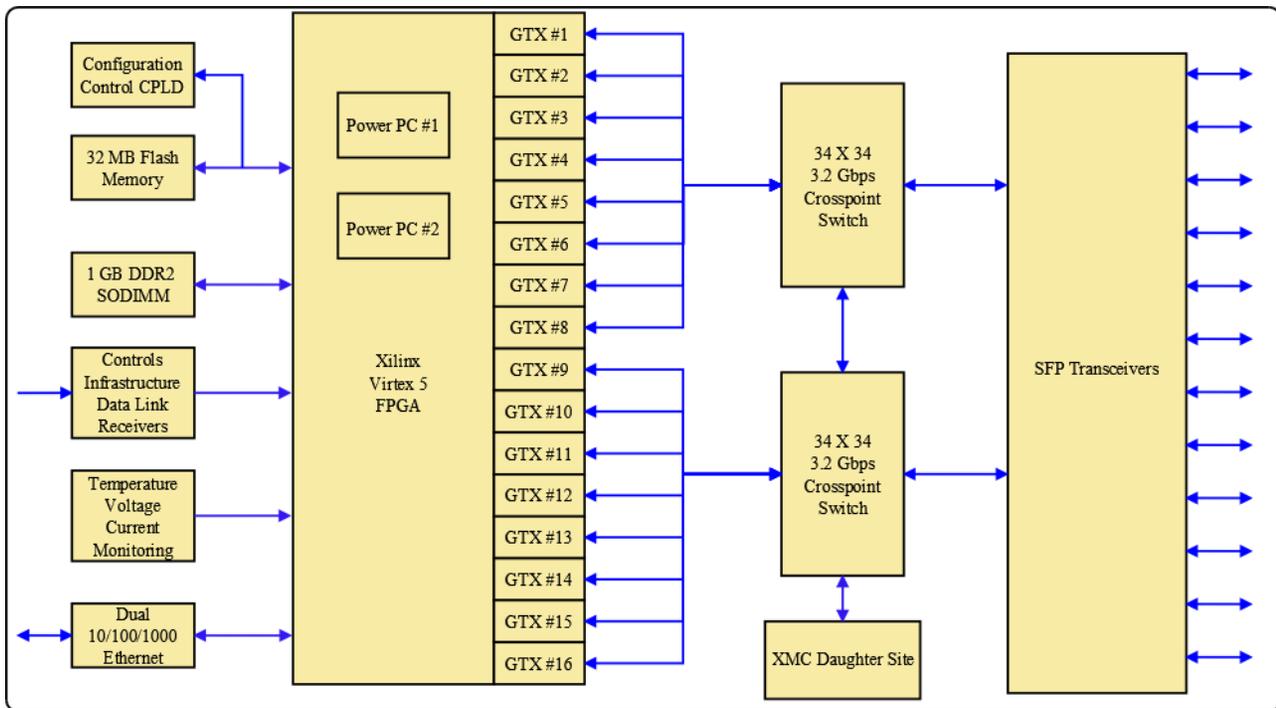


Figure 1: Update Link Master block diagram

circuit. Unlocking this PLL causes the barrel shifter to relock at a new position. This process is repeated until the desired barrel shifter position is achieved.

Several problems were encountered while developing the initialization routine. First, the register address of the barrel shifter position read back is not documented by Xilinx. Next, the barrel shifter showed a propensity for coming back to a few select values even though all evidence in the documentation indicates the position should be truly random. By using sequence of different ways to unlock PLL, we were able to determine an algorithm that would quickly achieve the correct barrel shifter position when the link was initialized.

All Update Link transmissions are 32 bit words. The data format is defined with the upper 16 bits as the Packet Identifier (PID) and the lower 16 bits as the data payload. Each data word transmitted is assigned a unique PID. A special PID of 0x00 is assigned for timing events. For timing events, the lower 16 bits represent a particular event code.

SYSTEM OPERATION

The Update Link Master generates an Update event that is broadcast every 10 microseconds, defining an update period. Running the link at 1Gbps, 250 32 words can be sent per period, which is more than sufficient for our current needs. The line rate can be increased up to 3.125 Gbps to allow greater throughput if necessary.

After each Update pulse, three data words are sent that together represent a 48 bit time stamp. The time stamp is used to synchronize diagnostic data recorded in different RF controllers.

Data for the Update Link can come from one of several sources. The Update event and time stamps are generated

by the Update Link firmware itself. The ULM has receivers for the standard timing links that are part of the RHIC control system infrastructure. Any event from these links can be encoded and rebroadcast on the Update Link. Data or events can also be sent manually through a user interface connected via the Power PC. The majority of the data delivered over the update link is generated in the downstream distributed system as described below.

The path of data flow in the system is shown in figure 2. Data is generated in a single GTX transmitter that is broadcast, by means of the crosspoint switches, to all RF Controllers in the system. Each controller has an SFP transceiver that receives the signal. A distribution chip then sends the received signal to the main FPGA on each RF Controller and to the FPGA on all 6 XMC daughter sites. An Update Link Receiver (ULR) firmware block in each target FPGA decodes the serial data stream.

The path for data to get sent back to the update is more involved. Each daughter card has a series of FIFOs that are used for holding data to be broadcast. A state machine scans the FIFOs for the presence of data and sends the data back to the carrier board via a 2.5 Gbps Aurora Link [4]. The carrier FPGA has a separate FIFO dedicated to receiving outgoing data from each daughter site. Just like the on the daughters, a state machine in the carrier FPGA scans the FIFOs and sends data back to the ULM via the same RocketIO transceiver that delivers the Update Link. At the ULM, there is yet another set of FIFOs to store data waiting to be broadcast and a state machine to control the priority of access to them. At each level of system, there is a specific order in which the FIFOs are polled and a limit to how many pieces of data can be pulled from a FIFO at once to control the priority of data delivery.

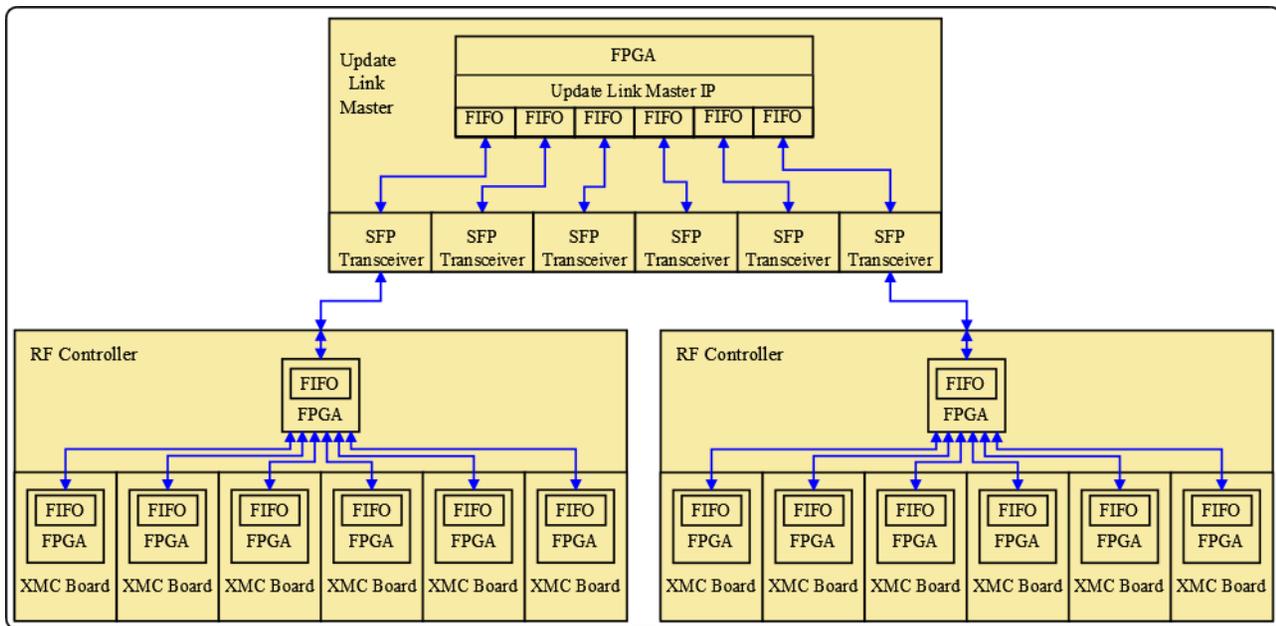


Figure 2: Update Link System data flow

The ULM master has a total of 34 SFPs with which to broadcast the Update Link. Since the FPGA on the ULM has only 16 GTX tiles, only 16 RF Controllers can be routed back through the crosspoint switches to send data to the ULM. In the event that more than 16 RF Controllers need to send data, a second ULM chassis can be configured not to be the master but to consolidate data from a number of SFPs onto a single data stream that is then connected back to the master.

Examples

In the RHIC LLRF system, a processor on one of the XMC daughter cards receives data on the magnetic field and the desired radial position from the control system and uses this information to calculate the nominal revolution frequency of the beam. The LLRF beam control loops calculate a correction to this at a 100 kHz rate. The 32 bit revolution frequency word is split into two 16 bits words that are each combined with their appropriate PID and placed in the high priority FIFO. This data moves up through the RF Controller to the ULM master and is broadcast to all receivers. The revolution frequency data is used by numerous Direct Digital Synthesizers (DDS) distributed throughout the system and applied as the new input to the DDS on the next Update pulse. Since all the DDSs are driven by the common 100 MHz master clock and apply new data to the synthesizers at the same time, all synthesizers in the system remain locked in phase [5].

Another example of the use of this system is in making digital phase measurements between various synthesizers. An event is broadcast that causes all synthesizers to latch the current value of their phase accumulators. The latched phase from a reference synthesizer is then pushed

up to the ULM and broadcast so other synthesizers can measure their latched phase with respect to the reference.

RESULTS

To test the quality of the physical link, we performed a loop back test using the ChipScope™ Integrated Bit Error Rate Tester (IBERT) core, an extremely useful tool from Xilinx. The test sent Pseudo Random Bit Stream (PRBS) data through the link at 1 Gbps. The test was terminated after three days with no bit errors.

The Update Link system was successfully commissioned along with the rest of the new RHIC LLRF system during the 2010 run. There continue to be some minor issues related to the link initialization firmware in the Update Link Receivers but overall performance has been exceptional. The ability to make direct digital phase measurements between physically separate synthesizers has opened up a number of new operational capabilities in the LLRF.

REFERENCES

- [1] K.S. Smith et al, "Concept and Architecture of the RHIC LLRF Upgrade Platform," these proceedings
- [2] T. Hayes, K.S. Smith, "A Hardware Overview of the RHIC LLRF Platform," these proceedings
- [3] "Virtex-5 FPGA RocketIO GTX Transceiver User Guide" (2009), <http://www.xilinx.com>
- [4] "Aurora 8B/10B Protocol Specification" (2010), <http://www.xilinx.com>
- [5] T. Hayes et al, "A High Performance DAC/DDS Daughter Module for the RHIC LLRF Platform," proceedings of PAC 2011, New York