

PHENIX Level-1 Trigger Systems for RHIC Run-3

John Lajoie, John Hill, Sergei Belikov, Harold Skank, Gary Slegee, and Steve Skutnik

Abstract— The PHENIX detector at the Relativistic Heavy Ion Collider (RHIC) will study the dynamics of ultra-relativistic heavy ion collisions and search for exotic states of matter, most notably the Quark Gluon Plasma (QGP). Substantial event selectivity is needed at RHIC to enhance interesting events relative to more common ones and to satisfy the requirements of the data acquisition system. The first on-line screening is achieved by the Level-1 Trigger. The Level-1 Trigger is a beam-clock parallel-pipeline system that uses Local Level-1 (LL1) algorithms pertaining to the fastest PHENIX subdetectors, followed by a Global Level-1 (GL1) system that processes encoded LL1 reduced-bit data to issue up to 32 triggers. In this paper we discuss a new set of LL1 trigger systems that will be deployed for the third RHIC run. These new trigger systems utilize the same basic hardware designed around commercial FPGA logic in order to provide a common platform for a wide variety of physics triggers.

I. INTRODUCTION

The Relativistic Heavy Ion Collider at Brookhaven National Laboratory has embarked on an ambitious program to study the nature of matter under extreme conditions through the collision of heavy ions at ultrarelativistic velocities. It is predicted by Lattice QCD simulations that the large energy density created in these collisions will be sufficient to study a phase transition between ordinary matter (in which quarks are confined to hadrons and mesons) to a new state of matter known as the Quark Gluon Plasma (QGP) where the quarks and gluons are free to roam over large distance scales [1]. The energy density required for this phase transition is similar to that when the universe was a few μ seconds old.

The PHENIX experiment [2] at RHIC is a large, multi-armed spectrometer optimized for the detection of rare probes of the earliest stages of the collision (photons, jets and J/ψ production) as well as hadronic observables that observe the

created system at freezeout. Already from the first two runs of RHIC intriguing measurements have been made that suggest that these collisions are creating matter in a state that has not been previously observed [3]. In addition to the heavy-ion physics program, PHENIX is also pursuing a physics program studying polarized proton collisions at 200 GeV/c in an effort to shed light on the origin of the nucleon spin [4].

In order to pursue a rare-event physics program the PHENIX experiment requires high luminosity from the RHIC accelerator and highly selective Level-1 and Level-2 trigger systems. In particular, the Level-1 trigger system is limited to a maximum rate of 12.5kHz by the readout rate of the detector front-end electronics. This will require an event rejection of up to ~ 1000 when RHIC reaches its goal of ten times design luminosity for protons in future runs. The PHENIX Level-1 trigger is fully pipelined and provides an event decision in less than 4.2μ s (40 RHIC clock ticks).

In order to meet this challenge we have designed and tested a new set of PHENIX Level-1 trigger electronics. In contrast to previous PHENIX Level-1 designs, which were hardware-specific to a given detector, we have designed electronics that are entirely reprogrammable and able to operate with a variety of PHENIX detector systems. The trigger boards receive data from a set of up to 20 optical fiber links at a maximum rate of 19.2 Gbit/s per board and process the data in up to five Xilinx XCV2000E programmable gate arrays.

This new design is being implemented to handle trigger information from the Zero Degree Calorimeters (ZDC, which monitor participant nucleons in nucleus-nucleus collisions) and Normalization Trigger Counters (NTC), a combined trigger using the PHENIX electromagnetic calorimeter (EMCal) and ring-imaging Cerenkov counter (RICH), and from the Muon Identifier (MuID) system. For the MuID the full trigger will process information from a set of Iarocci tube counters between walls of steel absorber. The trigger algorithm will allow for dynamically tracking particles in both the horizontal and vertical projections of the MuID, allowing for skipped hits and multiple scattering of the muons in the absorber.

II. THE PHENIX LEVEL-1 TRIGGER SYSTEM

The PHENIX Level-1 Trigger is a parallel, pipelined trigger system designed to provide a highly configurable trigger capable of meeting the demands of the PHENIX physics program. The Level-1 Trigger consists of two separate subsystems. The Local Level-1 systems

Manuscript received November 14, 2002. This work was supported by a grant from the U.S. Department of Energy, Nuclear Physics Division

J. Lajoie is with the Department of Physics and Astronomy, Iowa State University, Ames, IA 50011 (telephone: 515-294-6952, e-mail: lajoie@iastate.edu).

J. Hill is with the Department of Physics and Astronomy, Iowa State University, Ames, IA 50011 (e-mail: jhill@iastate.edu).

S. Belikov is with the Department of Physics and Astronomy, Iowa State University, Ames, IA 50011 (e-mail: belikov@bnl.gov).

H. Skank is with the Department of Physics and Astronomy, Iowa State University, Ames, IA 50011 (e-mail: skank@ecad3.ameslab.gov).

G. Slegee is with the Department of Physics and Astronomy, Iowa State University, Ames, IA 50011 (e-mail: slegee@ameslab.gov).

S. Skutnik is with the Department of Physics and Astronomy, Iowa State University, Ames, IA 50011 (e-mail: skutnik@iastate.edu).

communicate directly with the participating detector systems. The input data from these detector systems is processed by the LL1 algorithms to produce a set of reduced-bit input data for each RHIC beam crossing. The Global Level-1 system receives and combines this data to provide a trigger decision. In addition, buses (both global and trigger) are managed by GL1.

The PHENIX detector system readout is divided into two sets of elements: *granules* and *partitions*. A granule is the smallest detector element that communicates with the PHENIX timing and control system via a Granule Timing Module (GTM). The GTMs distribute the local 9.4Mhz RHIC beam clock as well as control bits and event accepts to the granule. A partition is an administrative configuration of granules that share both buses and Level-1 triggers.

Both the LL1 and GL1 systems are implemented in nine or eleven layer 9U VME-P format boards using Actel [5] FPGA logic. In order to minimize design time and effort, all Level-1 boards share a great deal of common infrastructure logic (for example, VME communication). All boards incorporate the ability to inject test patterns at their inputs under both VME and timing control. In addition, the data flow on the board can be monitored at the beam clock level via a set of synchronous fifos that store 1024 beam crossings worth of data from all stages in the board algorithm. The combination of test patterns and monitor fifos allows real-time monitoring of the Level-1 trigger and allows users to quickly pinpoint hardware failures or misconfigurations. A separate set of FIFOs store additional information for inclusion in the data stream for accepted events.

III. THE GENERIC LOCAL LEVEL-1 HARDWARE

The Generic Local Level-1 Hardware (GenLL1) was designed to address a number of difficulties faced in extending the original design concepts for LL1 systems to address the full range of PHENIX trigger needs. First, the hardware had to be reprogrammable to allow not only quick corrections to the trigger logic, but modifications of the trigger system as the PHENIX physics program evolved. This also had the added advantage of lowering development costs and allowing quicker prototyping of new trigger systems. Second, the new hardware had to be able to manage a data throughput of $\sim 20\text{Gb/s}$ in order to be able to handle large data volume detectors (such as the Muon Identifier) without prohibitive cross-stitching between multiple boards (and crates of boards). Finally, in order to keep power consumption and heat load on the board within reasonable limits the GenLL1 design was required to use the next generation of HP GLINK receiver/transmitter logic (the HDMP 1032/1034), which runs at 3.3V. Smaller format transceivers (Agilent HFBR 5912) were also required in order to be able to handle up to twenty fiber inputs on a single 9U VME board.




Fig. 1. Block diagram of the GenLL1 trigger hardware, showing the fiber inputs (20), GLINK logic, FPGA and termination card locations, and VME interface.

The basic design of the GenLL1 board is shown in block-diagram form in Figure 1. The board is designed so that data from the detector systems flows into the front of the board via the optical transceivers, is processed in the FPGA's located near the center of the board, and the reduced trigger information flows out the back connectors (via a VME transition card).

The choice of Xilinx FPGA's was determined by several factors. The availability of large pin-count packages, such as the FG680, was essential in order to design the board so that each Xilinx chip could receive data from all twenty GLINK receivers. In addition, large logic densities (~ 2 million gates per FPGA) were required. In the end we chose the Xilinx XCV2000E FPGA as it fulfilled the above requirements and there was already substantial Xilinx experience within the PHENIX collaboration.

The GenLL1 boards were manufactured as 11-layer boards [6] using an innovative via/pad structure for the FPGA footprints. Because the design called for using essentially all of the 680 pins available to us on each FPGA, the trace routing density was extremely high, especially when pin escape vias were taken into account. Instead of a traditional footprint, the pad for each FPGA pin was drilled out after the board was laminated and filled with a conductive epoxy. After curing, the top and bottom of the pad was plated, essentially forming a "through-hole" footprint for the FPGA. Using this manufacturing method allowed us to route the connections to the FPGAs without pin escapes, dramatically simplifying the routing and manufacturing process. In addition, the pin connections to the power and ground planes had a reduced inductance due to the lack of short traces to pin escape vias, which helped limit noise on the FPGA power. Each GenLL1 board has ten locations that can accept an XCV2000E FG680 FPGA.

In order to make maximum use of the available board area for FPGA logic, the data busses from the GLINK receivers were routed directly to the FPGAs and source-terminated. Full termination was not practical because the GLINK pin drivers are unable to drive a full termination, and the

additional power consumption of a full termination was unacceptable. Upon testing the first GenLL1 boards produced it was found that source termination was adequate for small systems operating at $4 \times \text{BCLK}$ (such as the NTC/ZDC trigger logic), where the fibers used could be chosen so that their data busses terminated on the FPGA chip used, limiting the influence of reflections at the FPGA pins. However for systems with a large number of fibers, or systems operating at $6 \times \text{BCLK}$ (such as the EMCAL/RICH and MuID systems), reflections between the end of the bus and the source termination were significant. This was solved by designing a set of “termination daughtercards”, or small PC boards with a FG680 format, containing the termination resistors for a specific FPGA location on the board (four fiber busses terminate at each of five FPGA footprints). These termination daughtercards had small solders “balls” installed on the FG680 footprint on the bottom of the cards, and were then installed on the main board in much the same way as the Xilinx FPGAs [7].

Our experience with the FG680 BGA format, as well as the conductive epoxy footprint and termination daughtercards has been excellent. Board stability was improved by the addition of three aluminum stiffening rails surrounding the FPGA locations, forcing the board to remain flat during the FPGA mounting process. Some warping of the termination daughtercards was a problem, however, requiring careful adjustment of the oven used to heat the assembly and melt the solder. At the present time we have not found any broken connections or open solder joints on the assembled GenLL1 boards.

The completed GenLL1 assembly (for a MuID trigger configuration) is shown in Figure 2. The fiber termination cards can be clearly seen in the photograph.




Fig. 2. A completed GenLL1 assembly (in a MuID trigger configuration, showing the fiber receivers and termination daughtercards.

IV. FPGA PROGRAMMING

Due to the fact that a great deal of “infrastructure” logic had already been developed for LL1 systems in a schematic capture system, it was decided to break the programming into two pieces. The foundation logic in the chip, such as the GLINK demultiplexing circuit, VME interface, accepted event data and monitor fifos, and data bit masks were

implemented in schematic capture. This allowed us to make use of existing designs used for earlier systems. However, the channel mapping and physics algorithms for the chip are implemented in VHDL. The FPGA logic is designed in a modular fashion, so that the same infrastructure can accommodate multiple physics algorithms. We are currently using the Xilinx ISE 4.2i [8] toolset, running under both the HP-UX and Windows platforms.

The design of the GLINK demultiplexing circuits within Xilinx logic presented a serious challenge. The demultiplexing process requires that we take four or six frames of data from the fiber data bus (at four or six times the RHIC clock frequency of 9.4Mhz, or 38.4 and 57.6Mhz, respectively) and reconstructed them into a single beam crossings worth of data at the RHIC clock frequency. Particular attention must be paid to the phase detection and synchronization of the data to avoid combining data from adjacent beam crossings. Because the XCV2000E only offers four global clock buffers, there were insufficient global clock resources available to allow both the RHIC clock and the incoming fiber data clocks to both use the global clock buffer resources.

We decided to place the RHIC clock on a global clock buffer and use the “low skew lines” resources (22 of which are available in the XCV2000E) for the incoming GLINK clocks. We found that in addition to the timing constraints implemented in the design, we were required to use specific RLOC constraints on the elements of the GLINK demultiplex circuits in order to minimize the skew between the various elements in the design. With this approach we were able to successfully design GLINK demultiplex circuits operating at four times and six times the RHIC, combining as many as twelve of these circuits into one FPGA.

For the Muon Identifier trigger, the complicated cable mapping from fiber bit to detector channel was done using autogenerated VHDL code. The MuID cable mapping is stored in an Objectivity database and maintained by the detector group. A C++ program was written to extract this channel mapping and write VHDL code for the algorithm logic that combines the input fiber bits to a set of “logical tubes” used in the algorithm processing. This VHDL code is translated into an EDIF netlist using the Xilinx XST synthesis tool.

The MuID trigger logic that processes the logical tubes and generates trigger primitives (see the algorithm discussion that follows) is also written in VHDL, allowing the functional aspects of the physics trigger to be tested and modified quickly.

V. THE GENLL1 TRIGGER ALGORITHMS

The trigger systems outlined below demonstrate the flexibility of the GenLL1 trigger boards, from systems with multiple independent trigger entities per board, to systems that combine detectors to form trigger decisions, and finally a system that uses multiple boards to process a single detector.

A. The NTC/ZDC LL1 Trigger

The PHENIX detector uses a variety of counter systems for interaction determination and characterization of the particle multiplicity in a collision. The Normalization and Trigger Counter (NTC) is used to extend the phase space coverage for minimum bias triggering in proton-proton collisions. It is essentially four quadrants of scintillator on each side of the interaction region, allowing a determination of the interaction timing a z-vertex location within the PHENIX detector. The Zero Degree Calorimeter (ZDC) is a neutron calorimeter located at either DX magnet that bends the beams into (and out of) the PHENIX interaction region. It is used primarily in heavy-ion collisions to characterize the impact parameter of nuclear collisions, and can provide an additional measure of the z-vertex of the interaction using the timing between the calorimeters at either end of the interaction hall.

The NTC/ZDC LL1 trigger system is, in fact, two separate trigger systems running concurrently in a single GenLL1 board. Because only five fibers at four times the RHIC clock are required (four for the NTC and one for the ZDC), the GenLL1 board is only partially populated and utilizes only one FPGA for algorithm processing. The algorithm consists of forming a mean time for each side of the detector systems using TDC inputs from the fiber links, followed by a subtraction to produce a measure of the beam vertex. Bounds can be placed on valid TDC values as well as the interaction vertex, and this information can be passed on to GL1 for triggering of luminosity monitoring purposes.

B. The EMCal/RICH LL1 Trigger System

The EMCal/RICH LL1 trigger system is a highly flexible trigger system for the PHENIX central spectrometer arms, combining information from both the Electromagnetic Calorimeter (EMCAL) and Ring Imaging Cerenkov detector (RICH). The electromagnetic calorimeter measures the energy of incident electrons and photons, while giving a minimal response for hadrons. The RICH detector will give a signal for electrons or high momentum (>5 GeV/c) pions. The EMCal/RICH LL1 system allows coincidences to be made between overlapping regions of the EMCal and RICH in order to generate a wide variety of possible triggers.

Because the EMCal/RICH system serves two separate spectrometer arms of the detector, it is implemented in two GenLL1 boards, each with one FPGA and eight fiber links. The FPGA programming for the EMCal/RICH LL1 system is being done by Brookhaven National Laboratory.

C. The MuID LL1 Trigger System

The PHENIX Muon Identifier (MuID) system is two separate detectors at the ends of the two muon spectrometer arms. Layers of steel are interspersed with five layers of Iarocci tube charged particle detectors, and muons are identified by their ability to penetrate the steel absorber (hadrons will shower within the absorber material). A muon with momentum greater than ~ 2.5 GeV/c will penetrate through all layers of the detector.




Fig. 3. MuID tubes are combined between panels in the same gap to form *logical tubes* that span the entire layer. Logical tubes between layers are combined into *symsets* that form a consistent road pointing back to the interaction point.

The MuID LL1 system will use a set of four fully populated GenLL1 boards (each with twenty fibers and five FPGA's). Each board receives the complete data from one projection (horizontal or vertical) of one of the MuID detector arms at a data rate of 18.4Gbit/s. The first step in the LL1 algorithm is to combine all of the hit tubes from subpanels within a given layer of Iarocci tubes, forming a "logical" tube that spans the entire acceptance of the detector at a given layer (see Figure 3). These combinations of logical tubes are organized into "symsets", indexed by the tubes number in the first gap. The "symsets" are collections of tubes that form a consistent path for a track from the origin. The range of tubes included in a symset at each layer is broadened to allow for the multiple scattering of the muons through the absorber material. Note that a given symset may share a great deal of tubes with neighboring symsets.

A symset is said to be "hit" if the pattern of hits logical tubes in the symset satisfies criteria consistent with a penetrating charged track. These criteria can allow for skipped gaps (gaps without hits) to increase the efficiency of the trigger. In addition, it is possible to process in parallel symsets that require different maximum depths in the absorber to allow triggering on "shallow" muons, which come from different physics processes.

The MuID LL1 trigger implementation is the most demanding both in terms of hardware and the FPGA logic requirements. Typical logic usage in the MuID LL1 FPGA programs ranges from 45-65% of the total available logic elements for a simple symset implementation. This allows adequate spare logic for future expansion and development.

VI. SUMMARY

We have described a new set of generic trigger hardware for the PHENIX experiment that allows a wide range of flexibility in developing Level-1 triggers. The GenLL1 boards can be configured for a variety of hardware needs, from small systems with a few input fiber connections (such as the NTC/ZDC LL1), to large systems that process data

very high rate (such as the MuID LL1). Many of the system design, hardware and programming techniques are directly applicable to other particle physics experiments.

VII. REFERENCES

- [1] J. Harris and B. Muller, "The Search for the Quark Gluon Plasma", *Ann. Rev. Nucl. Part. Sci.* (1996)
- [2] D. P. Morrison et al., "The PHENIX Experiment at RHIC," *Nucl. Phys. A638,565c* (1998)
- [3] K. Adcox et al., "Suppression of Hadrons with Large Transverse Momentum in Central Au+Au collisions at $\sqrt{s}=130$ GeV," *Phys Rev. Lett.* 88 022301 (2002)
- [4] N. Saito et al., "Spin Physics with the PHENIX Detector System," *Nucl. Phys. A638, 81c* (1998)
- [5] Integrator Series FPGAs –40MX and 42MX Families Preliminary Data Sheet, Actel Corp., September 1997
- [6] The GenLL1 boards were manufactured by Accurate Circuit Engineering, 30195 Kilson Dr., Santa Ana, CA 92707.
- [7] FPGA mounting and assembly of the termination daughtercards was done by Electromax, Inc., 1960 Concourse Drive, San Jose, CA 95131
- [8] For more information on the Xilinx development environments, see their web site at <http://www.xilinx.com>