UNIVERSITY of CALIFORNIA
SANTA CRUZ

DEVELOPING AN FPGA-BASED READOUT FOR THE PCT DETECTOR SYSTEM

A thesis submitted in partial satisfaction of the
requirements for the degree of

BACHELOR OF SCIENCE

in

PHYSICS

by

Jason R. Heimann

27 June 2005

The thesis of Jason R. Heimann is approved by:

Professor Hartmut F.-W. Sadrozinski
Technical Advisor

Professor Clemens Heusch
Thesis Advisor

Professor David P. Belanger
Chair, Department of Physics

Abstract

Developing an FPGA-Based readout for the pCT detector system

by

Jason R. Heimann

A novel silicon readout system is built using a Xilinx Virtex-II FPGA prototyping platform. Detector hardware from the 1997 GLAST beam test has been reworked and is read out with our new system. A charge-to-time converter is integrated into our chain to permit acquisition of calorimeter data in the same FPGA. After verifying the performance of our detectors and DAQ system, we use the equipment to perform a proton radiography experiment at Loma Linda University Medical Center.

Contents

List of Figures **v**

List of Tables **vi**

Dedication **vii**

Acknowledgements **viii**

1 Introduction **1**
1.1 Background on beam test hardware 1
1.2 Previous work on pCT experiment 3
1.3 Experimental issues 4

2 DAQ System Design **6**
2.1 Fanout board rework 7
2.2 FPGA hardware design 9
2.2.1 Initialization 10
2.2.2 Trigger response 11
2.2.3 Formatting and buffering 12
2.2.4 Handshaking and transmission 14
2.3 CTC integration 16
2.4 Software 16

3 System Verification **18**
3.1 Pulse injection 18
3.2 Noise and source measurement 18
3.3 CTC calibration 22

4 Beam Test Results **26**
4.1 Alignment 26
4.2 Tracking 26
4.3 Calorimeter 27

5 Conclusion **33**

A Dictionary of acronyms **35**

Bibliography **36**

List of Figures

1.1 Example proton paths 4

2.1 Block diagram of beam test system 7
2.2 Diagram of logic family conversions 9
2.3 Timing of signals generated during read out 12
2.4 Oscilloscope capture of signal distortion 15

3.1 Oscilloscope capture of response to calibration signals 19
3.2 Sample noise map 20
3.3 Photograph of the setup used to measure particles from a ^{90}Sr source. 21
3.4 A channel map acquired with a ^{90}Sr source 23
3.5 Physical map of front end chips on module 24
3.6 CTC calibration plot with line fit 25

4.1 A one-dimensional profile of LLUMC's proton beam 27
4.2 Correlation plot of deflection angles of protons traversing air 28
4.3 Correlation plot of deflection angles of protons traversing PMMA 29
4.4 Projections of correlation plots 30
4.5 Energy distribution from double telescope run w/ no absorber 31
4.6 Energy distribution from double telescope run w/ PMMA absorber 32

List of Tables

2.1 Event data word dictionary 14

To my late grandfather,

Andrew Richard Shore,

a true "Jack of all trades,"

as well as a dancer and a lover of life.

May his spirit live on in my words and ways.

Acknowledgements

First and foremost I must acknowledge the enthusiasm, guidance, friendship and compassion of my technical advisor, Prof. Hartmut F.-W. Sadrozinski. Over the years we've worked together I've been infected by his love for life, long lunches and science. Without him my experience as an undergraduate would not have been quite so fascinating and thoroughly enjoyable. Another character who contributed to my sustained fascination in physics is Prof. David Dorfan. Dave managed to befriend most every student he came across, and to a select few, he brought profound enlightenment with his unorthodox style of teaching. I am glad to say that I am among the select few, and will keep the lessons of "going with your gut feeling" long after I've left school. Many, many thanks are due to technicians Max Wilder and Bill Rowe for their selfless support and mentoring through the technical details of assembling and repairing our system. A tip of the hat goes to the little people who do big things in the labs: the post-docs, grads and undergrads who have shared their experiences and delights with me. This thesis would not have been possible without the loving support of my friend and partner, Sara Dale Williams, who made sure I was nourished, had clean clothes to wear, and a smile on my face throughout this torturous passage of experimental work.

1 Introduction

The Santa Cruz Institute for Particle Physics (SCIPP) has become an authority in particle detection and measurement, having developed equipment for many experiments such as GLAST, BABAR and ATLAS. Many of these developments include silicon strip detectors that provide precise (60∼25 μm) spatial resolution of energetic charged particles. SCIPP's expertise with silicon detectors lead to a collaboration with Loma Linda University Medical Center (LLUMC) in developing detector systems for several radio-biological experiments. The latest of these experiments involves a larger collaboration, called the pCT collaboration, to explore the energy loss of energetic protons in various media, for application in imaging. This experiment is part of a series of experiments investigating the feasibility of pCT using existing detector technologies.

1.1 Background on beam test hardware

A scaled-down version of hardware designed for the GLAST experiment was built and tested in a 1997 beam experiment at SLAC [1]. The hardware includes a series of tracker modules, each module containing two silicon detectors with front-end electronics. The detectors are mounted back-to-back, separated only by the thickness of the PCB on which they are mounted (1.5 mm). The detectors' strips are oriented orthogonally, allowing each module to make a two-dimensional measurement of a particle. The front-end electronics are of the second generation of GLAST ASICs fabricated for the experiment [2]. These ASICs are known as the GTFE32 chips as each chip is capable of reading out 32 channels.

Six GTFE32s are mounted on each side of the module, providing for a readout of 192 channels per detector. For each channel, the GTFE32 features a pre-amplifier, shaping amplifier and a common-reference comparator. Each comparator's output passes through a programmable digital mask into a latching shift register. Given a latch signal, the states of all 32 comparators are loaded into the shift register. The registers are daisy-chained such that all front-ends on one module operate in series; in effect, the channel registers comprise a large-scale parallel-to-serial converter. Data is shifted into the registers as well as out, allowing for the chaining of multiple modules. These registers are designed to be operated at clock speeds of up to 10MHz, allowing an entire two-detector module to be read out in 38.4 μs, or two daisy-chained modules in 76.8 μs.

To mimic the configuration planned for the GLAST experiment, the modules are mounted in an acrylic box that features a series of ten machined slots with 3 cm spacing. The box is wrapped in an aluminized Mylar foil to prevent the admission of light. With four, five or six modules in use across the ten slots, many different configurations are accommodated by the box. Detector alignment is achieved by pushing the modules (each in a different slot) against the back of the box, then clamping each one in place with a set screw. A square window is cut out of the acrylic in the front and back of the box. Similar cut-outs exist in the module circuit boards; all material overlapping the "active area" (about 4.5 cm^2) of the detectors has been removed. This minimizes energy loss and multiple scattering of the relatively slow protons encountering the materials in and of the box and modules.

A separate "fanout" board was built using perf-board with wire-wrapped interconnects, to consolidate and/or distribute signals required to collect event data from the modules. This board features seven mass-termination connectors: one to connect to the downstream DAQ (data acquisition) hardware, and one for each connected module (up to six). Additional connectors bring in power for the analog, digital and bias supply busses, as well as calibration pulses. The electronics on the board include two voltage regulators and eight PECL-TTL / TTL-PECL converter ICs. Two logic ICs, included for use with the previous DAQ system, are removed for this experiment (see

Section 2.1 for more information).

1.2 Previous work on pCT experiment

Investigations of nanodosimetry, proton radiography and pCT are ongoing at SCIPP in collaboration with LLUMC, with the aid of two x-y silicon modules purpose-built at SCIPP [3]. Previous experiments proved successful in the use of silicon detectors to locate individual particles delivered from LLUMC's proton accelerator, and images of a simple phantom were obtained [4]. The resolution-limiting effects of multiple Coloumb scattering are apparent in the results of these experiments and emphasize the need to develop a better understanding of this phenomenon.

Reconstructing an image from a collection of proton histories requires tracking individual protons, and this ability is limited by the migration of protons due to scattering. In tracking a proton, one can only measure its position at several points along its path, then extrapolate that proton's path between the measurements. Such extrapolation can be improved by measuring both the proton's position and direction [5]. Established MCS theory has been developed by collaborator David C. Williams for specific application in pCT: given the initial position and direction of a proton traversing a uniform medium, a "most likely path" (MLP) is determined when the final position and direction is known. The recently-developed theory includes calculations for a "probability envelope" that encloses the most likely path. Given the statistical nature of scattering, the MLP is a best estimate of the proton's travel, but the proton is never confined to this path. The probability envelope gives one the likelihood of finding the proton in relation to the MLP, should we measure the proton within the medium being traversed. See Figure 1.1 for example MLPs with probability envelopes, generated in a Monte Carlo simulation [6].

Several experiments are planned by the pCT collaboration to test the MLP theory, as well as measure the probability envelope. The medium selected for this experiment is PMMA, more commonly known as Plexiglas or Lucite. This medium is selected for its radiation length, similar to that of human tissue, the medium we ultimately wish to image. The experiment, in concept,

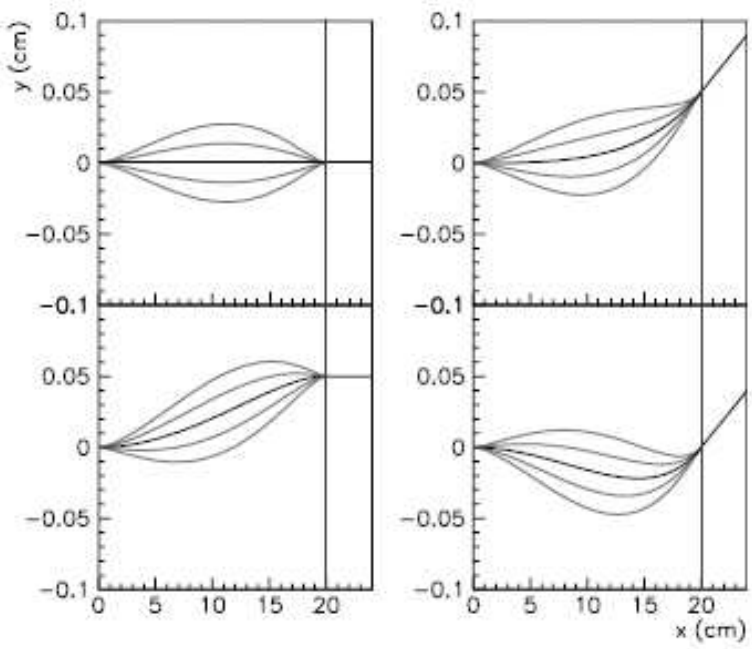

Figure 1.1: Example proton paths. The central, darker line is the MLP (given a constrained entry and exit position and angle), one- and two-sigma probability envelopes appear to the sides of the MLP [6].

is quite simple: measure the position and direction of protons before and after they traverse a known quantity of PMMA. A "roving" detector is placed at some interval within the PMMA to measure each proton's position as it is traversing the medium. The probability envelope can then be extracted from the roving detector's measurements, when viewing a subset of proton events (based on predetermined entry and exit positions and angles).

1.3 Experimental issues

To accommodate the "most likely path" experiment, a flexible, responsive and reliable tracking and calorimeter system is required. Readout electronics are needed that can easily handle any configuration of modules. Given the nonuniform spill structure of LLUMC's proton beam, these electronics need to process each particle event as quickly as possible (within 100 μs) to minimize dead time, while transmitting the event data efficiently to minimize data pile-up. Limited beam

availability demands a system which can operate with little or no down time, one that is tolerant to erroneous conditions and can recover without user intervention (avoiding the need for an operator in close proximity to the beam).

Since the hardware used in the experiment had been built for use with a DAQ system that is inaccessible to our collaboration, significant rework is required. This effort is compounded by the fact that some of the hardware was damaged in transit from its previous location. Before the MLP experiment can be performed, the rework and repairs need validation. The rework and development is detailed in Section 2, while the verification is described in Section 3.

2 DAQ System Design

From its successful implementation in a previous SCIPP experiment, a high-speed digital platform is selected to compose a data acquisition system for the beam test hardware. A Xilinx "Proto Board" development kit featuring Xilinx's Virtex-II FPGA was developed to read out events from a newer silicon detector front end, the PMFE (Particle Microscope Front End).[7] To facilitate read-out, the Proto Board is configured to connect between the DAQ PC (through a signal translation board) and a test board, a small PCB containing the PMFE and silicon detector. The FPGA is configured to generate clock signals, send them to the PMFE and read in event data. Also within the FPGA configuration are several layers of buffering and error-checking, to ensure the quality and continuity of the data stream. The previous successful deployment of the FPGA established its practicability and economy as a readout controller, in comparison to a dedicated CAMAC or VME system.

The GTFE and PMFE are similar in their signal input chain: they both contain a pre-amplifier, shaping amplifier and comparator. However, their data output chains are remarkably different. While the GTFE is designed to perform its read out once per triggered event, the PMFE produces a continuous stream of read out data. The PMFE employs several chained groups of registers to output 64 synchronously-latched channels of data over 8 output lines, in a dual-data rate format (utilizing both rising and falling edges of the clock to transfer data). The PMFE has no self-triggering ability; it is operated in a continuous fashion, whereby the chip's 64 channels are re-latched immediately after the data from the previous latch is output. In contrast, the GTFE

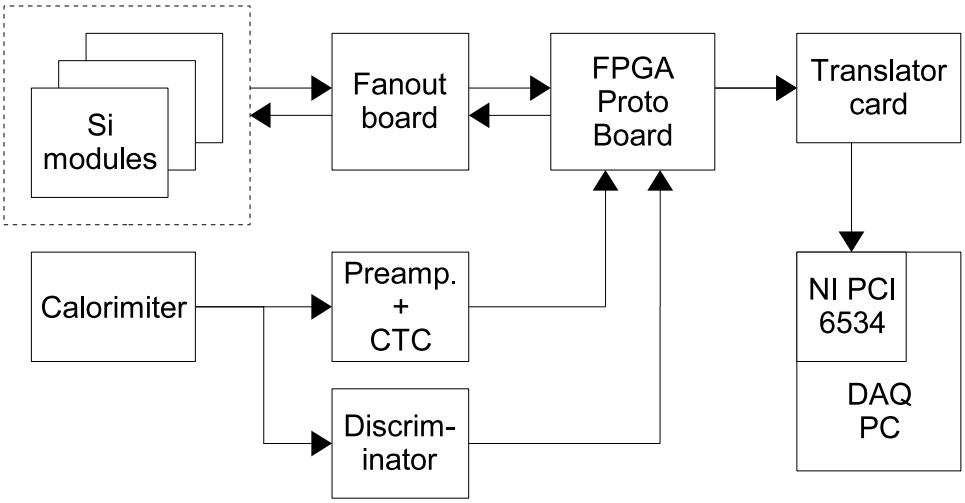

Figure 2.1: A high-level block diagram of the new beam test system.

provides a self-triggering capability, and its signals are propagated only on the rising edge of the input clock. It outputs its channel data serially through a single shift register, hence, the GTFE's output is comprised of a single line.

With the numerous differences between the GTFE and PMFE, and the versatility of the FPGA Proto Board, the decision was made to create an entirely new readout controller. Aside from the fabrication of a few twist-flat cables, no new physical hardware is needed. The new readout controller would exist entirely in the FPGA's internal configuration. To minimize the total amount of work needed to complete our DAQ system, the same "back end" as the PMFE system is re-used. This back end is comprised of a LVDS-CMOS signal translation board, and a PC with National Instruments' PCI-6534 high-speed digital I/O card [8]. See Figure 2.1 for a block diagram of the new beam test system.

2.1 Fanout board rework

The original beam test DAQ system utilizes an older differential signaling standard known as PECL (Positive Emitter Coupled Logic), a standard that is not directly supported by the FPGA.

Since the fanout board already contains PECL-TTL converter ICs, the decision was made to rewire and augment this board to achieve compatibility with the FPGA. To minimize the amount of hardware needed for this rework, two signal families are used in the FPGA-fanout interface: LVDS and LVTTL. Low Voltage Differential Signaling (LVDS) is a high-bandwidth, low power logic family that exhibits good noise and power supply immunity. Low Voltage Transistor-Transistor Logic (LVTTL) is a single-ended signal family that is level-compatible with standard TTL [9]. The compatibility of LVTTL drivers and TTL receivers is a stroke of good luck, as the FPGA is only capable of driving low voltage signals.

The TTL families in our system are referenced to separate ground planes (one on the fanout board and one on the Proto Board), inevitably distorting the rising edge of transitioning signals. The ground plane implemented on the fanout board consists only of narrow strips of copper tape, making the entire board especially sensitive to large ground currents. Unfortunately, PECL drivers consume large amounts of current ($\sim$16 mA per driver). This leads to fluctuations of several hundred millivolts in the ground potential over periods of tens of nanoseconds, when many signals transition simultaneously. The LVTTL standard allows a noise margin of 400 mV, and with a comparable "ground bounce" present, one can not rely on the precise timing of single-ended signals. Hence, the TTL family is only used to transmit signals that are tolerant of settling times on the scale of microseconds.

Existing driver ICs on the fanout board accept PECL or TTL input signals, and drive those signals out to each detector module at PECL levels. Data signals returning from the modules are converted through separate ICs from PECL to TTL, sent through multiplexing logic, then are converted back to PECL for output to the original DAQ system. For the new system, the logic ICs are replaced with TTL-LVDS and LVDS-TTL converters. Now the TTL data signals are sent out to the FPGA through the TTL-LVDS converters. Meanwhile the TTL-PECL converters are reused to fan out clocking signals from the FPGA, through the newly mounted LVDS-TTL chips, to each of the modules. See Figure 2.2 for a graphical "road map" of the signals as they undergo

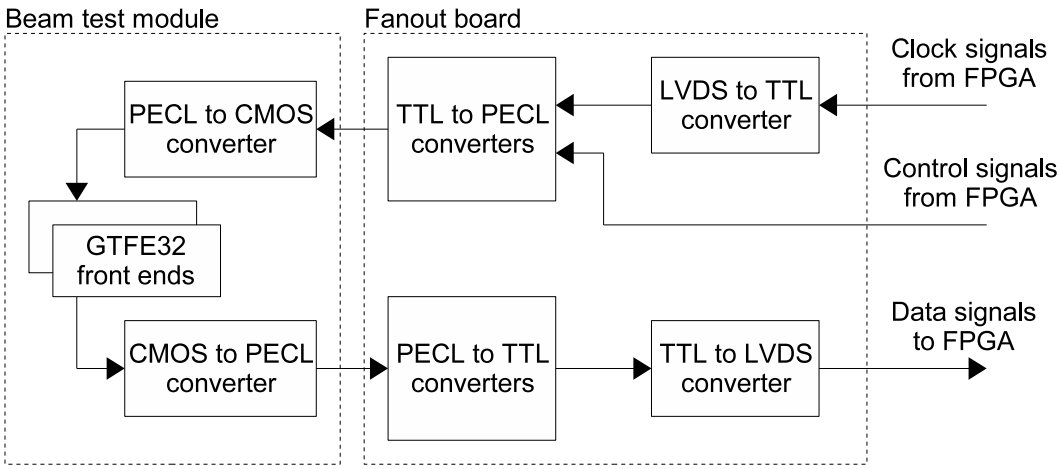

Figure 2.2: A diagram of the logic family conversions that occur between the detector modules and the FPGA Proto Board.

the above-mentioned conversions.

The use of the LVDS and ECL families are crucial in routing high-speed signals. The large ground fluctuations and subsequent signal distortion are overcome by using differential signals (two lines referenced to each other, not to a common ground), since the 10 MHz serial clock transitions every 50 ns. Distortion and pick-up are observable on the differential lines, but since these artifacts are common to both lines, the differential signal receivers eliminate such effects (by at least 60 dB, well within the allowable noise margins). Fluctuations are mitigated further by staggering the timing with which signals were switched; see Section 2.2 for more details.

2.2 FPGA hardware design

The Xilinx Proto Board provides a convenient interface from the FPGA chip to the outside world, via several arrays of .100" header pins. These pins are connected to the internal logic of the FPGA through highly configurable IOBs (Input-Output Blocks) that allow the designer to select the behavior of each pin: if the pin performs input or output, and which signal family to use. The Proto Board also features a service FPGA that (among other tasks) creates a system clock signal whose frequency is adjustable in 10 MHz steps up to 90 MHz. For reasons that will soon become

clear, a 60 MHz system clock is selected. Following the general recipe used to read out the PMFE, the hardware design is straightforward [11]. The FPGA is configured to perform the following tasks:

1. Initialize the GTFE32s in all modules
2. Respond to triggers (read out event data)
3. Format event data for output
4. Load data into a first-in, first-out (FIFO) buffer
5. Handshake with the PC I/O card and transmit data

The hardware design contains several discrete modules to allow the parallel processing of these tasks. To maintain synchronization between the modules, a common clock signal (the 60 MHz system clock) is distributed throughout the FPGA. In addition, a common reset signal is implemented. Each module contains a FSM (Finite State Machine) that synchronizes its operations with the system clock. The system reset button (which generates the reset signal), mounted on the Proto Board, initializes all modules and state machines into a known state. Once initialized, the system's behavior is entirely determined by external stimuli (detector and/or calibration signals). The hardware implemented to perform each of the above tasks is detailed in the following sections.

2.2.1 Initialization

In its power-on state, the GTFE32's mask register is pre-loaded such that all channels are disabled (masked). The mask register is programmed much like the channel registers are read out; i.e. each GTFE32's mask register is chained to the previous one so all front ends on one (or two) modules are programmed with a pair of lines. The "CLKD" line is used to clock data into the register chain, while the mask data appears on the "DDIS" line. Since a maximum of six modules must be supported, and modules are chained in pairs, three DDIS lines are implemented at the fanout board: DDIS0, DDIS1 and DDIS2.

Upon system reset, a small FSM in the top level of FPGA hardware inhibits trigger responses while clocking data into the modules' mask registers. A mask ROM is implemented in the FPGA, outputting its contents on the DDIS lines, to allow selective channel masking. Once the mask registers have been loaded, the system enables the trigger input and waits for trigger events. The mask lines (CLKD, DDIS0, DDIS1, and DDIS2) remain inactive while the trigger input is enabled.

Masking channels in hardware is necessary when there are many extremely noisy channels (approaching 100% occupancy); these "bad channels" would consume bandwidth in the data-driven FPGA-PC connection. When bandwidth becomes a premium, then hardware masking is a solution. Currently, the system provides enough bandwidth to read out several hundred channels per event without filling the FIFO buffer, so the decision was made not to implement any hardware masking. This approach has the advantage that the "health" of the system can be monitored in the computer (i.e. channel maps) without bias.

2.2.2 Trigger response

Because the GTFE32 does not have any internal event buffering, it is essential to read events out completely and quickly upon the arrival of a trigger signal. The self-triggering capability of the GTFEs is not used in this experiment because it is not reliably working on all modules. Instead, an input pin is implemented on the FPGA to accept any rising-edge signal as the system's trigger. Upon receipt of a trigger signal, the event handler FSM comes out of its wait state, pauses for 1.5 μs, then sends the signal required to latch the channel registers in all GTFEs. The 1.5 μs pause is inserted to compensate for the time constants of the front end amplifiers; this ensures maximum detector efficiency. To latch the channel registers, the serial register clock lines (CLKS0, CLKS1, CLKS2) are pulsed once, with the "SEL" (clock select) line held low. The SEL line simply steers the clock signal to either the latch (when low) or the shift register (when high).

Once the latch pulse is sent out, the FSM pauses for 950 ns before raising the SEL line. Another 50 ns later, the first serial clock is started. To facilitate the readout of three serial data lines

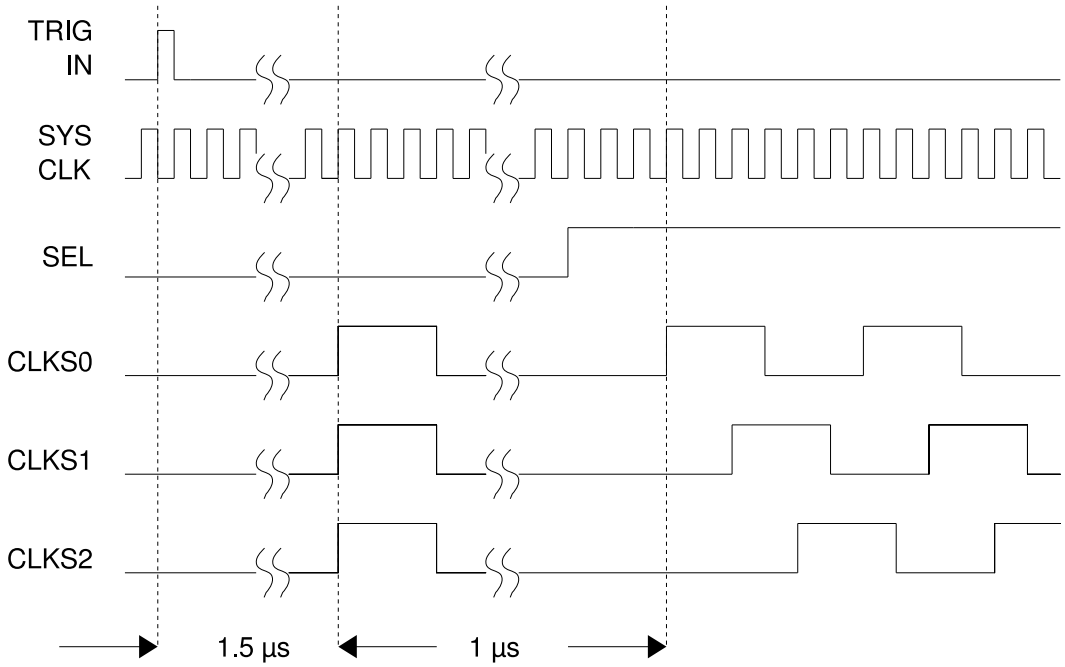

Figure 2.3: Timing of signals generated during read out, beginning with the trigger signal. The system clock (SYS CLK) frequency is 60 MHz, the serial clocks (CLKS0, CLKS1, CLKS2) are 10 MHz.

(DATAA, DATAB, DATAC) into a single buffer, the serial registers are clocked 120° out of phase in respect to each other, while the data lines are demultiplexed into a single signal. See Figure 2.3 for a timing diagram describing the structure of these signals. With this staggered timing, serial data is formatted on the fly and loaded into the FIFO buffer at rates up to 30 MHz (see the following sections for more details). Once 768 clocks have been delivered to the serial registers, the event handler FSM enforces a 17 μs dead time before resetting itself into the wait state. This dead time is important to allow any ground fluctuations induced during read out to settle. Without dead time the sensitive front end pre-amplifiers could interpret these fluctuations as signal, reporting false "hits" in the channel registers should a latch occur immediately after readout.

2.2.3 Formatting and buffering

The serial data appearing on the DATA lines are strings of ones and zeros. To translate these ordered bits into information that can be transmitted asynchronously, a channel number is

assigned to each bit. A counter is incremented, starting from zero, as the serial clock (specifically, CLKS2) pulses to read out event data. When a one appears on the data line, the value of the counter indicates the channel in the shift register that was hit. With up to 768 channels in the register, 10 bits are required to represent the channel number. Information is appended to this value to indicate the DATA line (i.e. module group) from which the channel hit is read; two bits are used to represent this value.

To pad the event data out to 16 bits, the input width of the I/O card, the 12 bits representing channel and module group are "packed" in the lower nibbles of the word, while the top (most significant) four bits are set to zero. See Table 2.1 for an explanation of the output data format. This data word is presented to the FIFO buffer for every channel during read out, but the buffer does not clock in the word unless the serial data line is high; the demultiplexed serial data line is tied to the write enable input of the FIFO buffer, while the buffer is clocked with the 60 MHz system clock. In alternate cycles of the system clock, the write enable input is forced low, to prevent duplicate writes.

A careful study of propagation times was performed to determine when a data bit would arrive, in relation to the transmission of a clock pulse to the corresponding serial register. Propagation delays of approximately 20 ns are incurred each time a signal (clock or data) passes through a translator chip. With two translator chips and 1.2m of cabling encountered by the clock, and the same for the returning data signal, approximately 100 ns (+/- 10 ns) of delay occurs between the rising edge of the clock and the rising (or falling) edge of the arriving data signal. This delay is equivalent to a single 10 MHz clock pulse; this delay is accounted for by waiting one serial clock pulse before starting the channel counter. The data line is not read on its edge, instead, it is read on the falling edge of the serial clock, so the signal has several tens of nanoseconds to stabilize.

To separate one event's data from the next, start and stop code words (see Table 2.1) are inserted into the FIFO buffer at the beginning and end of each event. To prevent confusion between channel hit data and code words, at least one of the top four bits in a code word is made nonzero.

Word	Meaning
0x0400 - 0x06FF	Channels 0 - 767, Module group 0
0x0800 - 0x0AFF	Channels 0 - 767, Module group 1
0x0C00 - 0x0EFF	Channels 0 - 767, Module group 2
0xAAAA	Event start code
0xCCCC	Filler code (ignored)
0xFFFF	Event stop code
0xDFFF	CTC readout failure
0xDnnn	CTC readout: result = nnn

Table 2.1: Event data words output from the FPGA to the DAQ PC

The start code is inserted immediately before the first channel is read out, and the stop code is inserted immediately after the last channel is read out. Lastly, "filler" code words are employed to overcome the discrete data transfer mechanism that exists in the PC I/O card. This card will only transfer data into the PC's memory in four word blocks, so if a number of words are sent to the card that is not a multiple of four, several words will not appear in the DAQ software. To prevent any such occurrence, sixteen filler words are written to the FIFO buffer after each event.

2.2.4 Handshaking and transmission

Once the FIFO buffer begins receiving data, the output controller FSM springs into action. This FSM has the task of communicating the data to the PC's I/O card. The I/O card is configured to receive data using a two-way handshaking protocol. Two dedicated lines are implemented to perform this handshaking: REQuest and ACKnowledge. When the output FSM wants to write data to the PC, it raises the REQ line and waits for the ACK line to go high before sending out any data. Once this condition is met, 16-bit words are sent directly from the FIFO buffer to the PC DAQ card at 10 MHz, one word per rising edge of the DAQ clock. Should the ACK line ever fall, the last word sent is retained in the FIFO buffer and the output FSM goes into an idle state. Once the handshaking is re-established, the data transmission resumes.

An additional challenge was imposed on us in the development of the output chain. For safety reasons we are required to locate our DAQ PC approximately 20m away from our detector,

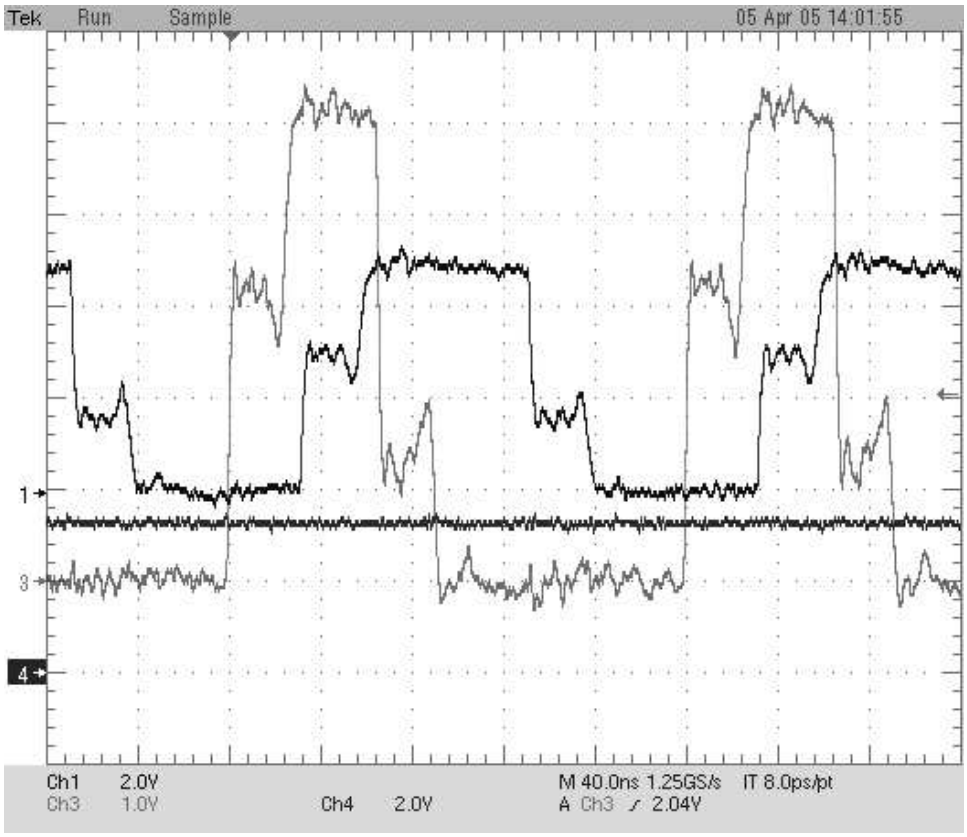

Figure 2.4: An oscilloscope capture of the signals at the PC end of the output chain, after 20m of cabling. The taller pulse is the REQ signal, while the shorter pulse is the clock signal. Due to distortion, the REQ signal appears to fall in sync with the rising edge of the clock. This scheme of handshaking on every clock pulse was abandoned because the excessive distortion lead to read-out malfunction.

requiring a lengthy cable between the FPGA and signal translation board. At first the handshaking signal structure from the PMFE experiment was implemented, but the inductive and capacitative effects of the long cable distorted the handshaking signals into obscurity. This distortion lead to read-out malfunction, and thus this signal structure was abandoned. See Figure 2.4 for an oscilloscope capture of the signals. The problem arose due to the handshake (exchange of ACK and REQ signals) occurring with every data word transmitted. The problem was solved by performing the handshake once, at the beginning of data transfer, and holding the handshake signal constant until the transfer is either interrupted or completed.

2.3 CTC integration

After completing the silicon detector read out system, the next task was to incorporate an energy measurement into the event data. The calorimeter is used to trigger our read out controller during data taking. Since the energy information is digitized within the calorimeter read out electronics, little additional work was required to make the incorporation. The calorimeter read out terminates in a rack-mounted device known as a charge-to-time converter (CTC), which converts the analog calorimeter signal into a digital pulse train. Encoded in the time between digital pulses is the desired energy information.

An additional input is added to the FPGA design that accepts the TTL signal from the CTC. A simple counter is incremented by the 60 MHz system clock, and gated by the pulse train coming from the CTC. The train's first pulse starts the counter from zero, while the third pulse (of four total per event) stops the counter. At the end of the GTFE serial register read-out, the CTC counter value is inserted into the FIFO buffer (formatted as in Table 2.1), just before the insertion of the stop code. If the CTC signal is not properly interpreted by the end of the serial register read-out, a failure code is inserted into the buffer in place of the counter value, so every event is guaranteed to have information from the CTC read-out.

2.4 Software

The software required to operate the PC I/O card was made fairly simple in the software provided by National Instruments [8]. As little as twenty lines of code are necessary to perform a read operation and save the data to disk, and this code is provided by NI with the purchase of the hardware. The code is augmented to do a very simple interpretation of the output data: decode the channel and module number, as well as CTC measurement, before writing the data to disk. The software also tags each event with a unique number (known as the event ID) to aid later analysis. Lastly, the necessary libraries are added to write our data in the ROOT file format [10]. The ROOT

libraries use a compression algorithm when storing data on disk, such that tens of thousands of events are stored in a file of just a few megabytes. With the additional programming the read-out software consisted of about 200 lines of C code; this program runs very efficiently and reliably on the DAQ PC.

3 System Verification

Before taking the system to LLUMC for the experiment, it is necessary to validate the read out controller design, as well as verify its operability with the detector system. Several systematic tests, detailed below, are performed to ensure each part of the read out chain works as expected.

3.1 Pulse injection

The GTFE32 includes four calibration busses that allow the user to inject charge into selected pre-amplifier inputs. Such injections can mimic the signals produced by an attached silicon detector, providing an ideal stimulus with which to study the front-end's response, as well as that of the read out chain. A bench-top pulse generator is used to send a 1 kHz train of pulses to a calibration bus while sending simultaneous pulses to trigger the read-out. The response from the GTFEs is monitored on an oscilloscope (see capture in Figure 3.1), while pulses are counted "by hand" and compared to the results reported by the DAQ PC. After several flawless trials, it was decided the read out controller was accurate.

3.2 Noise and source measurement

The following tests confirm the physical behavior of the detectors and front-ends. "Noise maps" of each module are obtained by setting up the modules for normal operation and reading them out with no stimuli. Again, the pulse generator is used to send trigger signals at 1 kHz, but

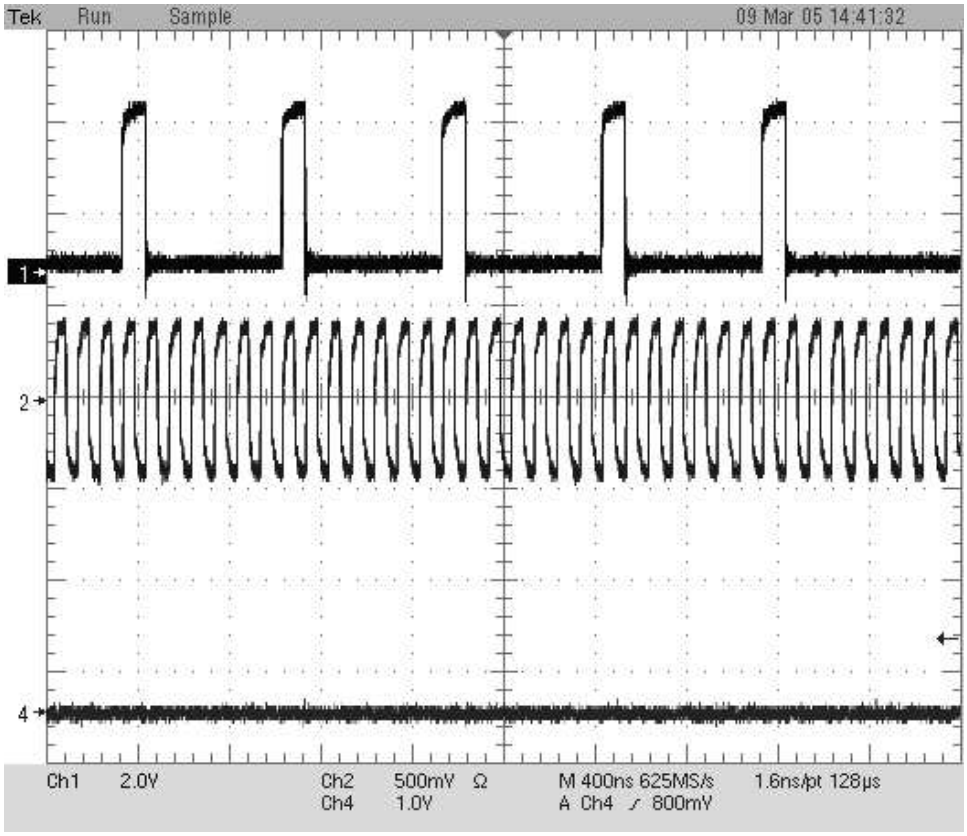

Figure 3.1: Oscilloscope capture of response to injected signals. The topmost signal is the serial data output from a module; the central signal is the serial clock output from the FPGA. Both signals are measured at the fanout board.

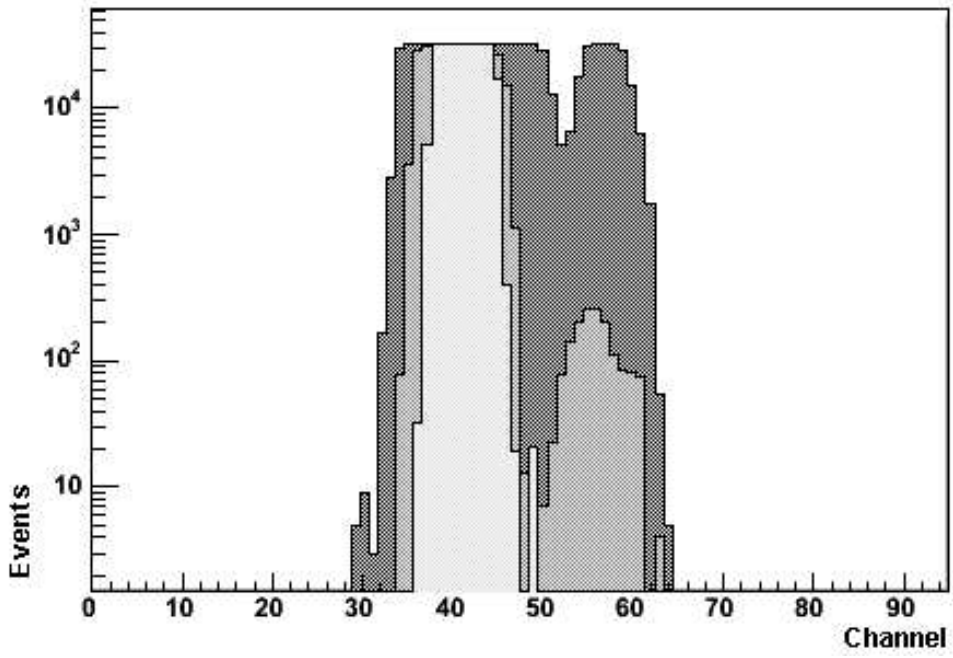

Figure 3.2: A noise map of the first 96 channels of a module. Three runs of data were collected at increasing thresholds, and the results are superposed here, highest threshold in front. Channels 32 - 63 show a somewhat noisy GTFE32. All other channels in this figure are very quiet at the selected thresholds.

now this signal is not applied to the calibration busses. Several runs of data are collected, each with the same number of triggers, but at different threshold voltages. See Figure 3.2 for a sample noise map with the results from three thresholds superposed.

In this plot one can see a decreasing response from channels 29 - 64 as the threshold is increased, just as expected. To ensure the front-end response is not a systematic effect, the module connections are reordered and the noise maps are reproduced. As expected, the only change appearing in the data is the order of the modules. This proves the utility of noise maps in identifying the connected modules; each module has its own noise signature that is readily distinguishable from the other modules.

A final test of the silicon read-out involves actual particle measurement. Particles are measured from a sealed, coarsely collimated ^{90}Sr source, placed at the approximate center of the

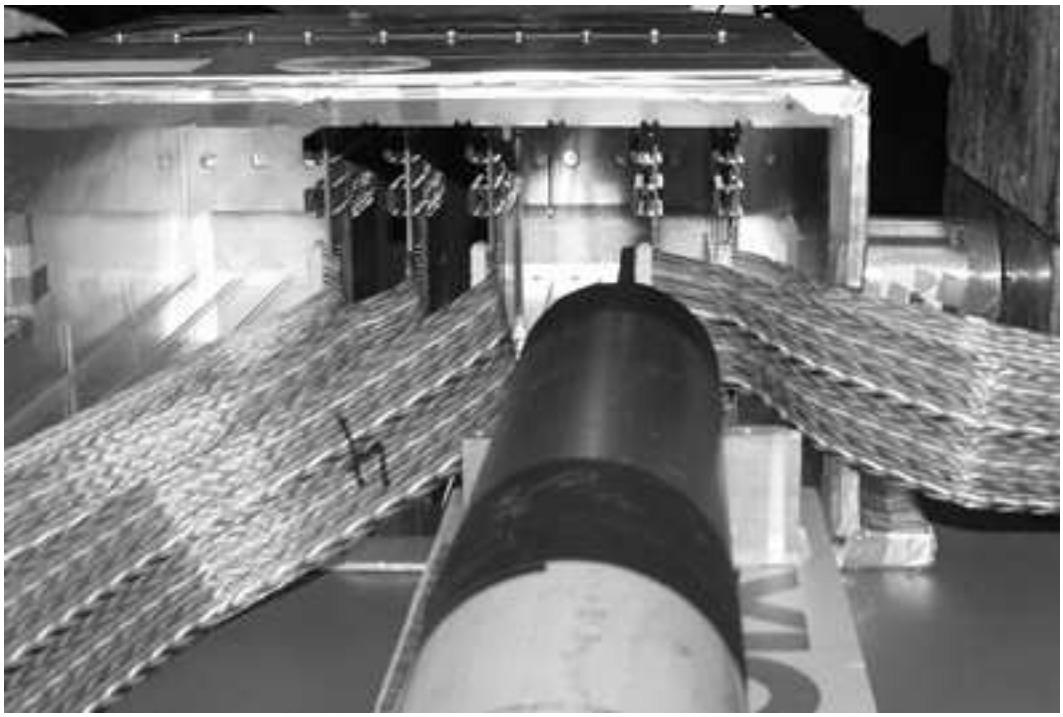

Figure 3.3: A photograph of the setup used to measure particles from a ^{90}Sr source. The source (held in place by lead bricks) is partially visible at the right side of the picture, against the acrylic box with five modules installed. A scintillator with PMT is placed two modules behind the source.

entry window, against the acrylic housing. A plastic scintillator with PMT is employed to generate a trigger signal. Beta particles emitted from the ^{90}Sr source have low energy (2.284 MeV) when compared to the 200 MeV beam at LLUMC. Through an informal experiment it was found these particles can traverse no more than two modules and still trigger the scintillator at an acceptable rate (at least 5x the dark count from the PMT, which was $\sim$10 Hz), hence, the modules are tested two at a time. See Figure 3.3 for a picture of our test configuration.

As with the noise maps, data is collected with equal amounts of triggers at several thresholds. The superposed results are collected into a "channel map." See Figure 3.4 for one such map. The coarse collimation of the source is visible in the large Gaussian-like distribution of channel hits, roughly centered in the detector. Aside from noise "spikes" appearing in some distributions, one can see very little difference between the various thresholds. This is a good indicator that the detectors are operating with a high level of efficiency. As an exercise in orientation, the source is moved to

one corner of the window and a small sample of data is collected. From the data, one is able to produce a physical map of the modules, to correlate each channel number with an actual location on each detector. See Figure 3.5 for our physical map.

3.3 CTC calibration

The CTC read-out is verified, then calibrated, in two separate exercises. The first exercise is performed at SCIPP, where a well-defined pulse train is created with a precision pulse generator to simulate the CTC output. This pulse train is fed into the FPGA while a noise map is collected, and the CTC measurement results are tabulated for comparison with the pulse generator settings. Several different pulse trains are generated to verify the linearity of the FPGA's response. In each case, the measured CTC times are compared to the timing of the generated pulse train as follows:

$$\frac{n_{counter}}{60\text{MHz}} = t_{CTC,measured} \tag{3.1}$$

We find the FPGA reports the CTC times quite accurately, plus or minus one count, giving an uncertainty of 1 / 60 MHz = 16.7 ns.

The second exercise is performed at LLUMC, where the scintillating crystal is connected to its associated read-out electronics (including the CTC) and placed in the beam line. Several beams are delivered, each with a well-known energy, and the CTC response is measured over several million events. A linear fit is performed as the CTC responses are plotted against the beam energies, and a "calibration equation" is produced from this fit (Equation 3.2; see Figure 3.6 for the latest calibration plot). With this equation one can determine the calorimeter energy in a particular event, given the CTC time reported by the FPGA. The results of the calibration exercise are detailed in a paper from our collaborators at LLUMC.

$$E\ (\text{MeV}) = t_{CTC} * 6.115\frac{\text{MeV}}{\mu\text{s}} - 36.23\text{MeV} \tag{3.2}$$

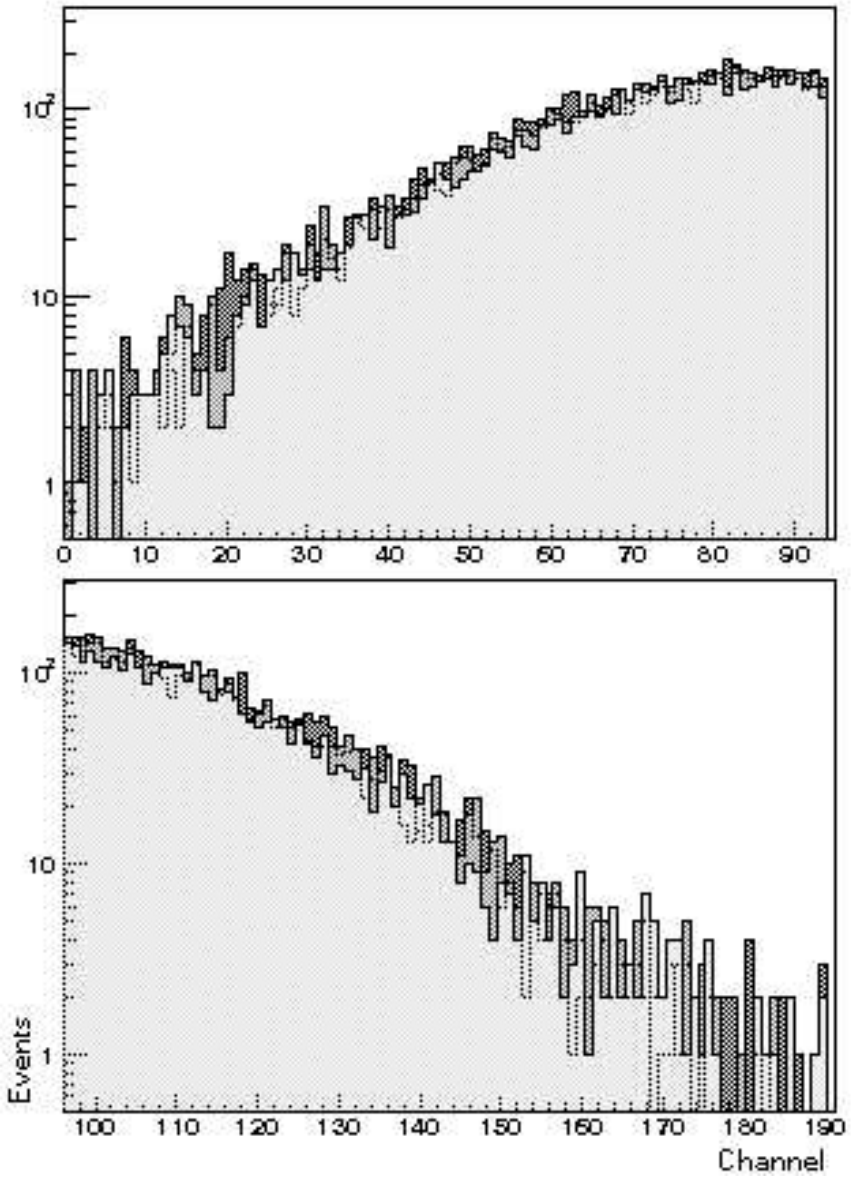

Figure 3.4: A channel map acquired with a ^{90}Sr source. Measurements are taken at three thresholds and superposed. The lack of a significant different between the superposed results indicates efficient detector operation.

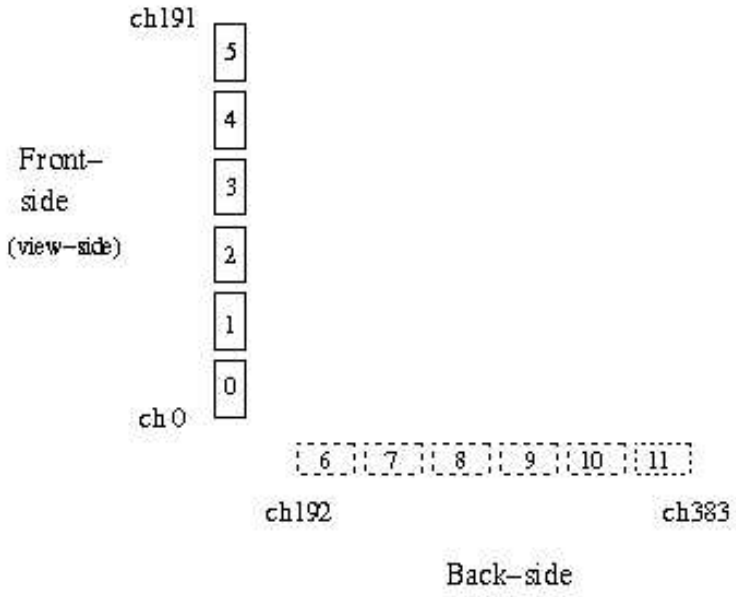

Figure 3.5: A physical map of the front end chips mounted on each silicon detector module.

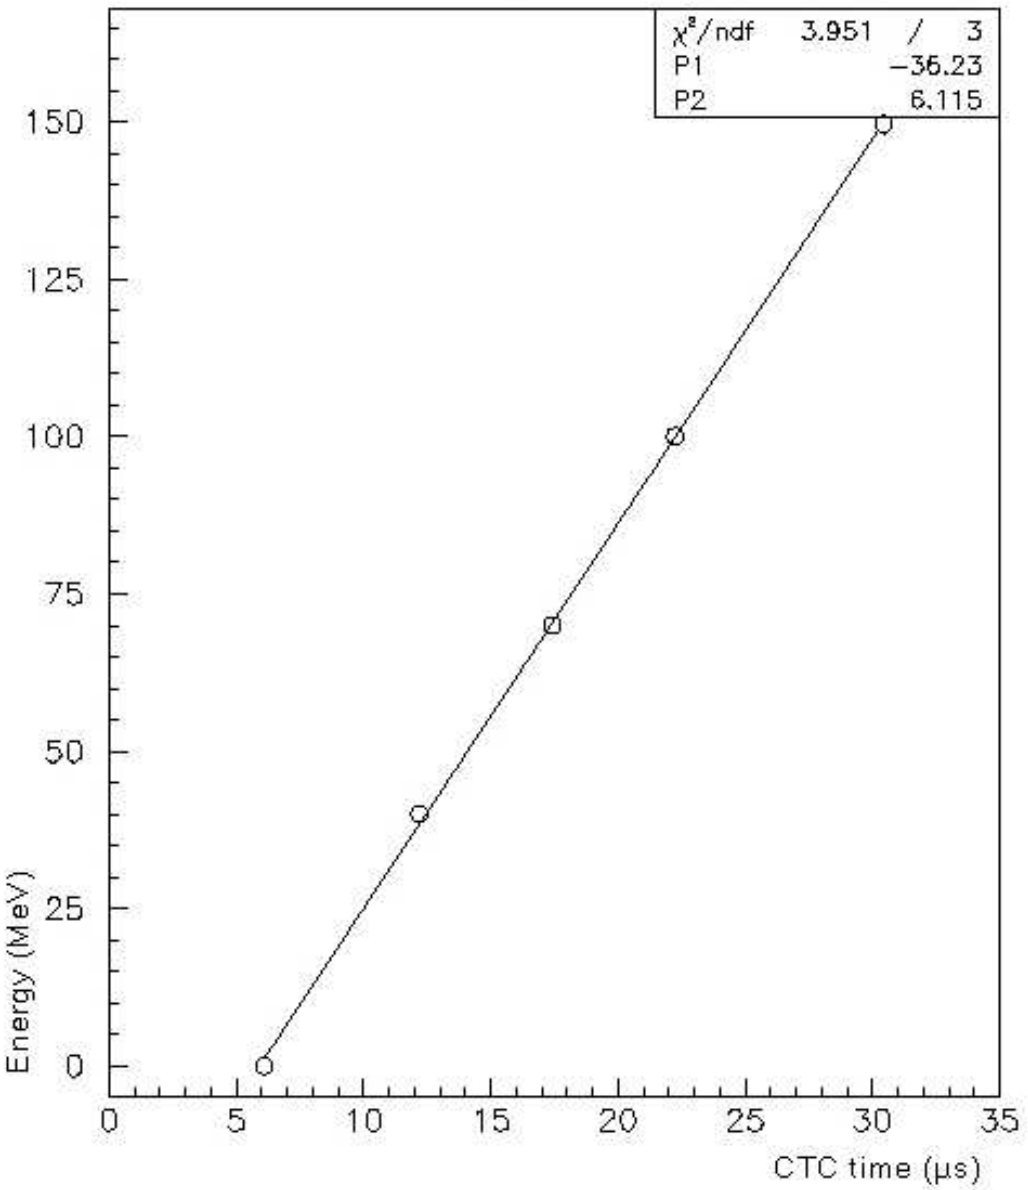

Figure 3.6: A plot of the most recent CTC calibration measurements, with line fit. The parameters obtained from this particular line fit are used to determine energy measurements from the recorded CTC times.

4 Beam Test Results

In this section selected results are presented that were acquired by the beam test system at LLUMC in May, 2005.

4.1 Alignment

The beam test box is attached to an optical bench and aligned with the beam through a series of measurement and adjustment cycles; a short run of data is collected, and the results are analyzed to determine how to better position the box, with respect to the beam profile and direction. The results from these runs are similar to the ^{90}Sr source measurements, in the appearance of a roughly Gaussian distribution across the channel map. See Figure 4.1 for a one-dimensional profile of the beam: the result of an alignment run.

4.2 Tracking

Once the box is in approximate alignment, the modules are placed in various arrangements and further beam runs are made, both with and without the presence of the PMMA material. The modules are arranged in a "double telescope" setup such that two adjacent modules are positioned at the front and two at the back of the beam test box. In this arrangement one can measure the direction of particles before and after they encounter the medium in the middle of the box, be it air or PMMA. Figure 4.2 shows the correlation between the entry and exit deflection angles (as projected

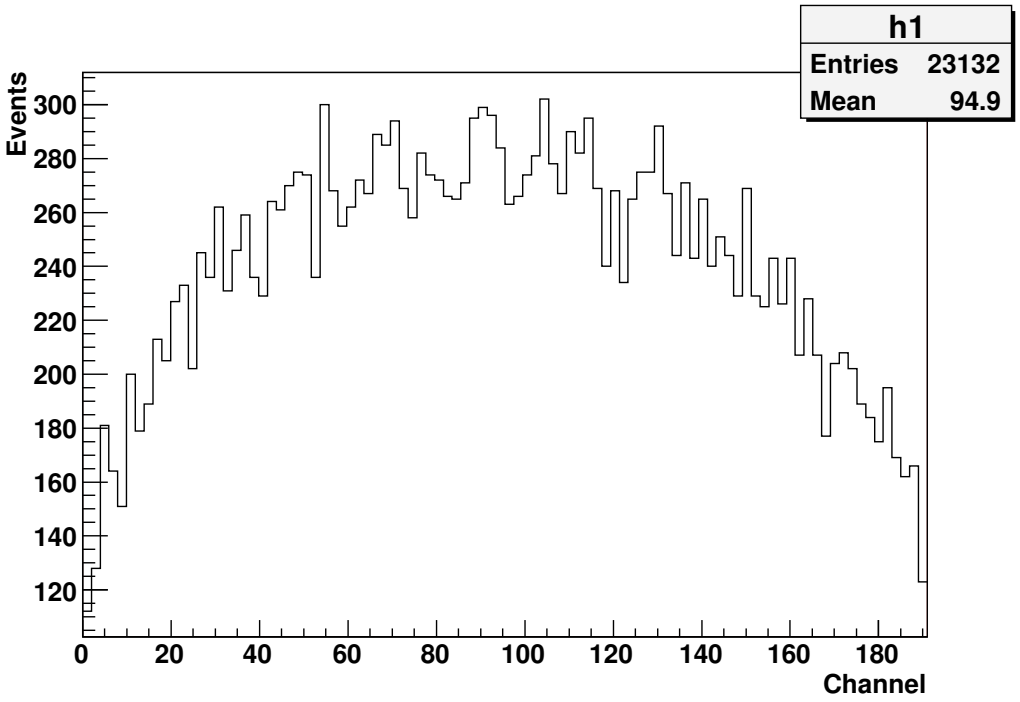

Figure 4.1: A one-dimensional profile of LLUMC's proton beam, as "viewed" from the first silicon plane in our detector array.

on the x-axis) of protons traversing air, while Figure 4.3 shows similar results with PMMA. These results are quite exciting in that they give direct evidence of the relation between material density and angular deflection incurred through multiple scattering (see a comparison of the previous plots in Figure 4.4).

4.3 Calorimeter

As discussed earlier, energy measurements are recorded with each particle event. Figures 4.5 and 4.6 indicate the energy distributions obtained from the double telescope runs. These results indicate the energy loss incurred through multiple scattering, and agree quite well with theoretical predictions: for the given length of PMMA used (approx. 12 cm) and the 200 MeV energy of the incident protons, 98 MeV is the predicted exit energy. Additional spikes appear in the energy distribution from the PMMA run, whose origin is still under study at SCIPP. Events with protons

scattering "around" the absorbing material are possible, and such events may produce the extraneous spikes.

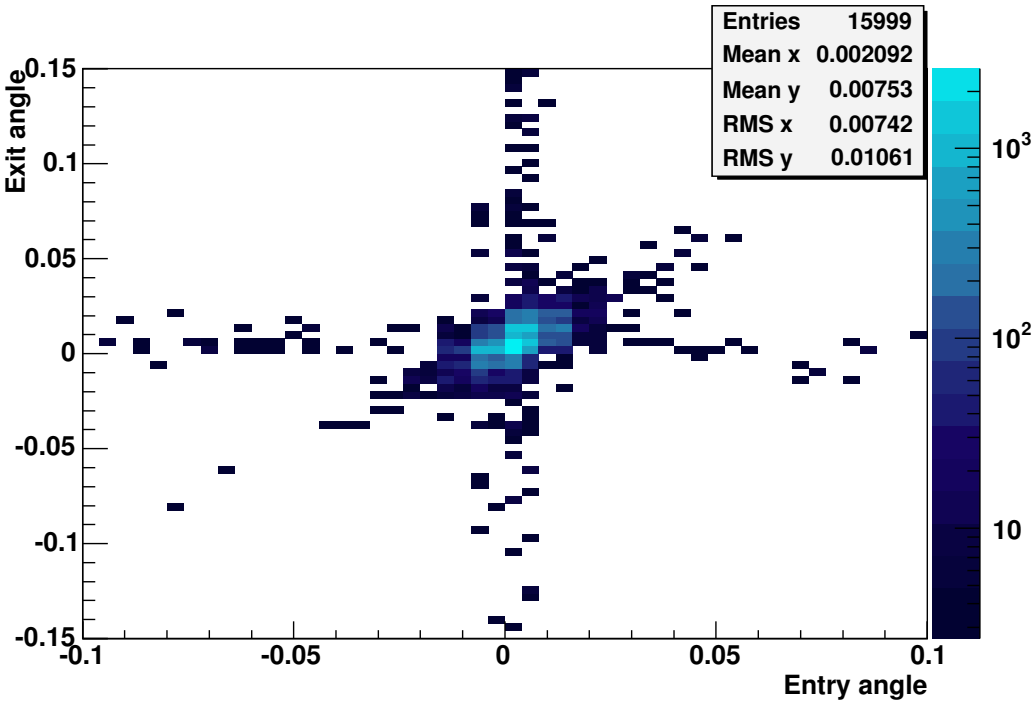

Figure 4.2: A correlation plot of the deflection angles of protons traversing air.

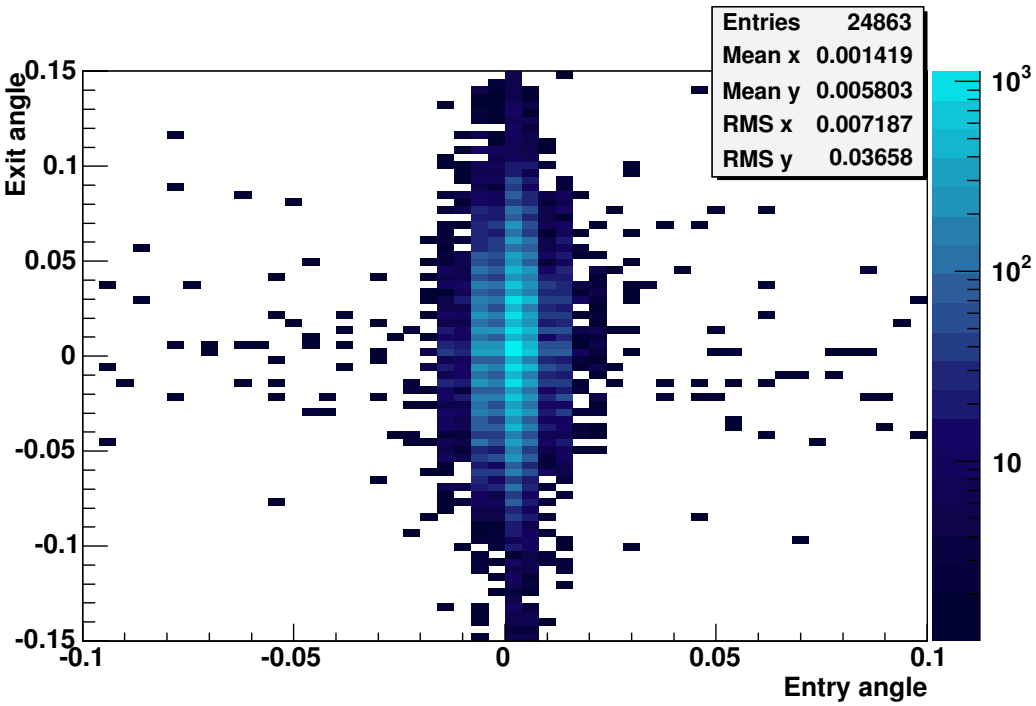

Figure 4.3: A correlation plot of the deflection angles of protons traversing PMMA.

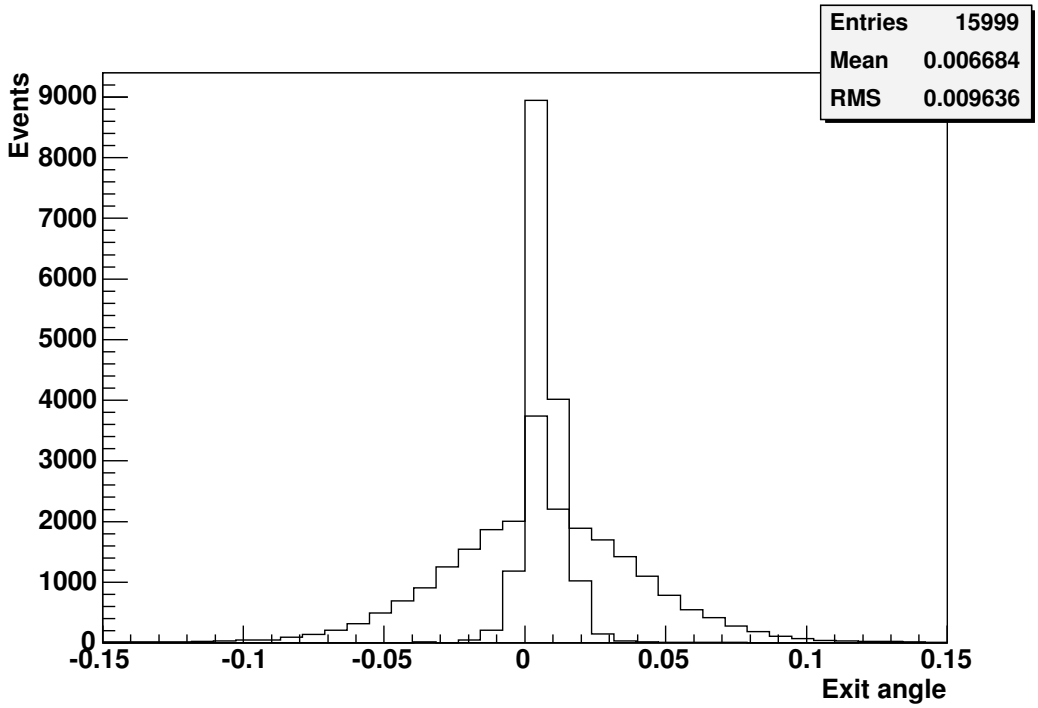

Figure 4.4: The previous two plots are projected onto the Y-axis and superimposed in this plot. The tall, sharp peak is the projection of the plot for protons traversing air; the low, broad peak is from protons traversing PMMA. MS theory predicts the most likely scattering angle as a function of the traversed material's density; this plot is direct evidence of this phenomenon.

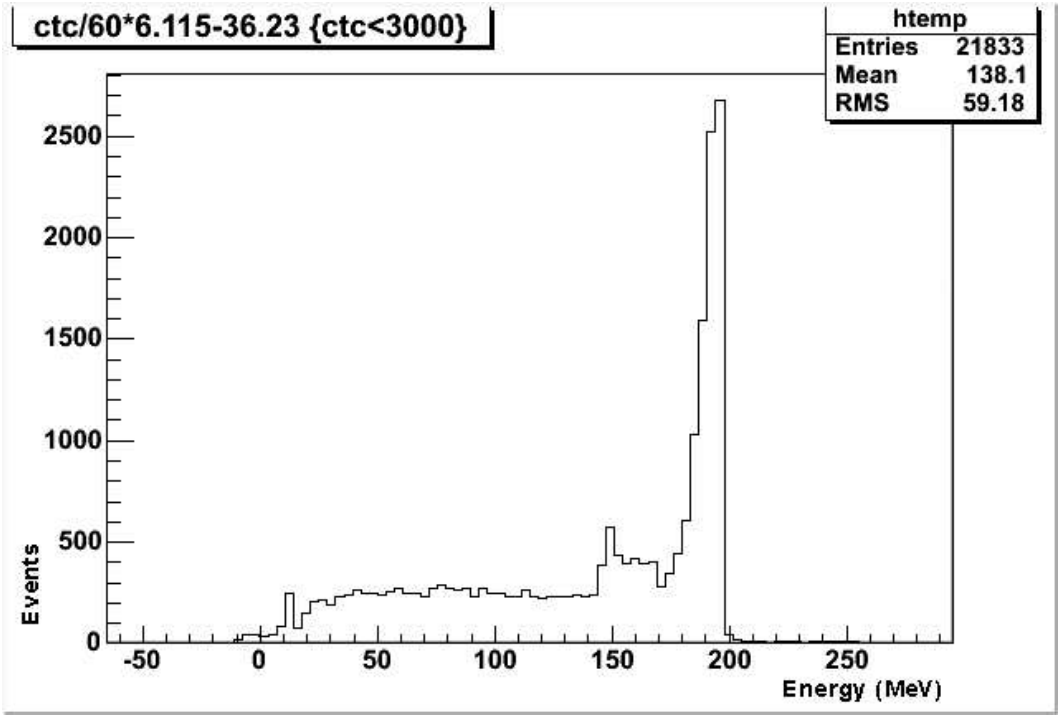

Figure 4.5: A plot of the energies measured during a double telescope run with no intermediate absorber. The 200 MeV beam energy for this run is manifested in the sharp peak at the "end" of the plot. The low energy "tail" is under investigation; this could be caused by protons leaving the calorimiter before stopping.

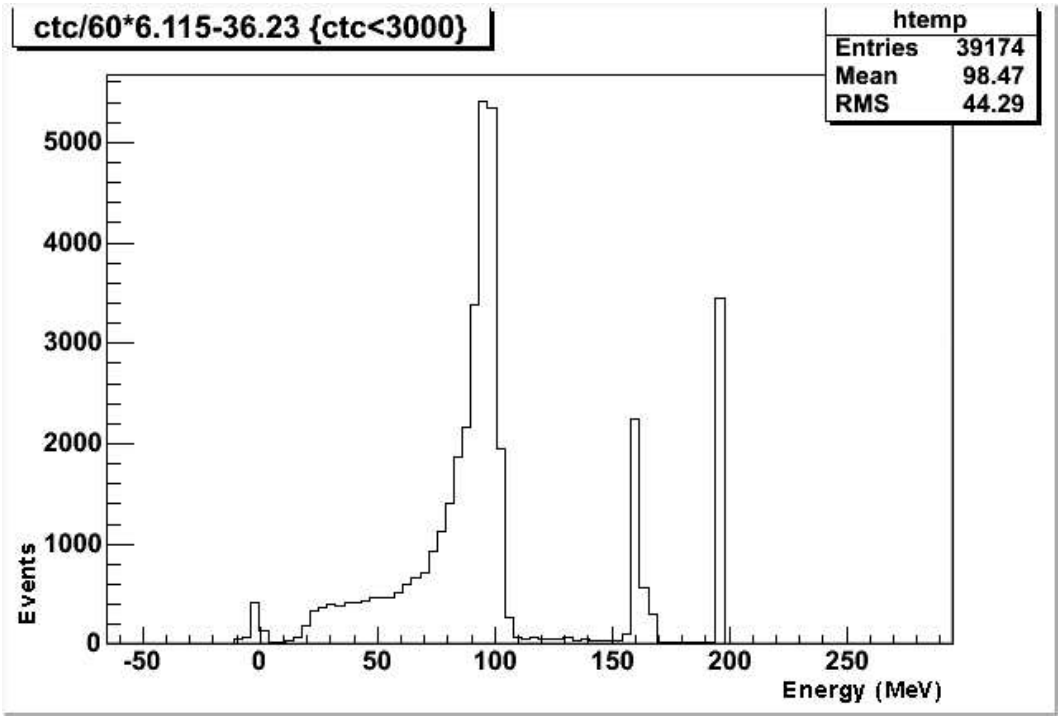

Figure 4.6: A plot of the energies measured during a double telescope run with 12 cm of PMMA absorber. The predicted exit energy for a straight-through proton is 98 MeV; a sharp peak is visible at this value.

5 Conclusion

Following the success of the PMFE test bed, a single-chip solution was implemented to solve the complex problem of heterogeneous, high-speed data sources. Where a rack-mounted solution was previously required, the compact and low-power FPGA Proto Board offered the same modularity and expansibility. These capabilities are manifested in the design of the latest pCT experiment, accommodating the simultaneous read-out of up to 6 x-y silicon detector modules with 72 front end chips, and a calorimiter, all into a single data stream. The ability to combine disparate data sources at a single point demonstrates the power offered by an FPGA-based system. This indicates the FPGA is a good solution for small-scale applications; the FPGA is an outstanding candidate in selecting a technology with which to implement the final pCT configuration, an even larger detector system that includes much larger silicon planes and 36 calorimeters.

The results obtained during the verification and calibration tests at SCIPP are convincing evidence of the effectiveness of an FPGA-based detector system, as every predicted or expected value was found in the measurements obtained with the beam test hardware and read-out electronics. Injecting calibration pulses demonstrated the correct operation of the GTFEs on each module, as well as the correct interpretation of the GTFE data signals by the FPGA. A relatively low background was observed after taking various noise maps. Measurements of a ^{90}Sr source yielded results that included very few "dead channels." The source measurements also indicated the detectors are efficient in their ability to track individual particles. After incorporating the necessary hardware to read out the CTC signals, initial calibrations showed the FPGA's time measurement abilities are

precise and accurate.

The verification results were confirmed with the beam test at LLUMC, as excellent results were obtained in the tracking and energy measurement of individual protons. Several profiles of the beam were obtained to get a picture of its geometry and structure, then many runs of data were collected with varied module and absorber configurations. These data runs contain hundreds of thousands of proton events that can help the collaboration develop a better understanding of the multiple scattering and energy loss of protons in a low-density medium. Ongoing analysis is being performed at SCIPP to extract the physical information sought in this experiment.

Appendix A

Dictionary of acronyms

ACK	acknowledge (handshaking signal)
ASIC	application-specific integrated circuit
CAMAC	computer automated measurement and control
CMOS	complementary metal-oxide semiconductor (logic family)
CTC	charge to time converter
DAQ	data acquisition
FIFO	first-in first-out
FPGA	field-programmable gate array
FSM	finite state machine
GLAST	gamma-ray large area space telescope
GTFE	GLAST tracker front end (front end ASIC designed for GLAST)
IC	integrated circuit
I/O	input / output
IOB	input / output block (usu. associated with FPGA)
LLUMC	Loma Linda University Medical Center
(LV)TTL	(low-voltage) transistor-transistor logic (logic family)
MCS	multiple Coloumb scattering
MLP	most likely path
NI	National Instruments (vendor of DAQ hardware)
PC	personal computer (a.k.a. desktop computer)
PCB	printed circuit board
pCT	proton computed tomography
(P)ECL	(positive) emitter-coupled logic (logic family)
PMFE	particle microscope front end (front end ASIC designed at SCIPP)
PMMA	polymethyl-methacrylate (trade names: Plexiglas or Lucite)
PTSM	particle tracking silicon microscope (built using the PMFE)
REQ	request (handshaking signal)
SCIPP	Santa Cruz Institute for Particle Physics
VME	versa module Europa (backplane system)

Bibliography

[1] W. B. Atwood, S. Ritz, P. Anthony, R. P. Johnson, W. Kroeger, H. F.-W. Sadrozinski *et al*, "Beam Test of Gamma-ray Large Area Space Telescope Components," *Nucl. Instrum. Meth.* A **446** (2000) 444-460.

[2] R. P. Johnson, P. Poplevin, H. F.-W. Sadrozinski and E. Spencer, "An amplifier-discriminator chip for the GLAST silicon-strip tracker," *IEEE Trans. Nucl. Sci.*, **45** (1998) 927-932.

[3] B. Keeney *et al.*, "A Silicon Telescope for Applications in Nanodosimetry," *IEEE Trans. Nuc. Sci.* **49** (2003) 1724-1727.

[4] L. R. Johnson *et al.*, "Initial studies from on proton CT using a silicon strip detector telescope," *Nucl. Instrum. Meth.* A **514** (2003) 215-223.

[5] H. F.-W. Sadrozinski *et al.*, "Toward Proton Computed Tomography," *IEEE Trans. Nucl. Sci.* **51**, No. 1, (2004) 3-9.

[6] D. C. Williams, "The most likely path of an energetic charged particle through a uniform medium," *Phys. Med. Biol.* **49** (2004) 2899-2911.

[7] H. F.-W. Sadrozinski *et al*, "The Particle Tracking Silicon Microscope PTSM," *IEEE Trans. Nucl. Sci.* **51**, No. 5, (2004) 2032-2036.

[8] Browse to `http://www.ni.com` for more information.

[9] Browse to http://www.interfacebus.com/voltage_LV_threshold.html for a comparison of logic families. See P. Horowitz and W. Hill, "The Art of Electronics," New York: Cambridge University Press (1980), pp. 326-9 for more information.

[10] Browse to http://root.cern.ch for more information.

[11] Brian Keeney, "The Design, Implementation, and Characterization of a Prototype Readout System for the Particle Tracking Silicon Microscope," Master's Thesis, Univ. of California at Santa Cruz, 2004