Chapter 1

Introduction

1.1 Purpose and Scope

This thesis describes a system capable of digitally recording the analog charge deposited on the Region 1 tracking system detectors used by the Qweak experiment at Jefferson Lab. Charged particle position information is made available by using a process known as Gas Electron Multiplication (GEM) whereby the original particle generates an avalanche of charged particles large enough to be detected by large-scale electronics. The final stage of this GEM detector contains analog voltage signals on individual copper strips within the charge collector. The analog signals on these individual strips are processed by a custom-made integrated circuit (IC), called VFAT2 (or VFAT), containing preamplifiers and pulse-shaping networks. [1][2][3]

For the purpose of recording the content of multiple VFAT cards simultaneously, it was necessary to use the V1495 Field Programmable Gate Array module built by CAEN which operates on a VME backplane.[4] The block diagram for the Region 1 detector is shown in Figure 1-1 on page 2.

Where necessary this thesis will reference and attempt to explain other portions from various manuals; all of these pertinent manuals are included on a CD with the original copy of this thesis. For the user firmware and the ROC (Read Out Controller) portions of the detector, UML diagrams have been included to simplify the process of familiarizing the reader with this detector and its respective algorithms for data acquisition from the VFATs. Unfortunately, due

to the finite nature of this thesis, there are many assumptions that are made about the level of technical proficiency of the reader; hopefully, these are minimal, but certainly they are unavoidable.

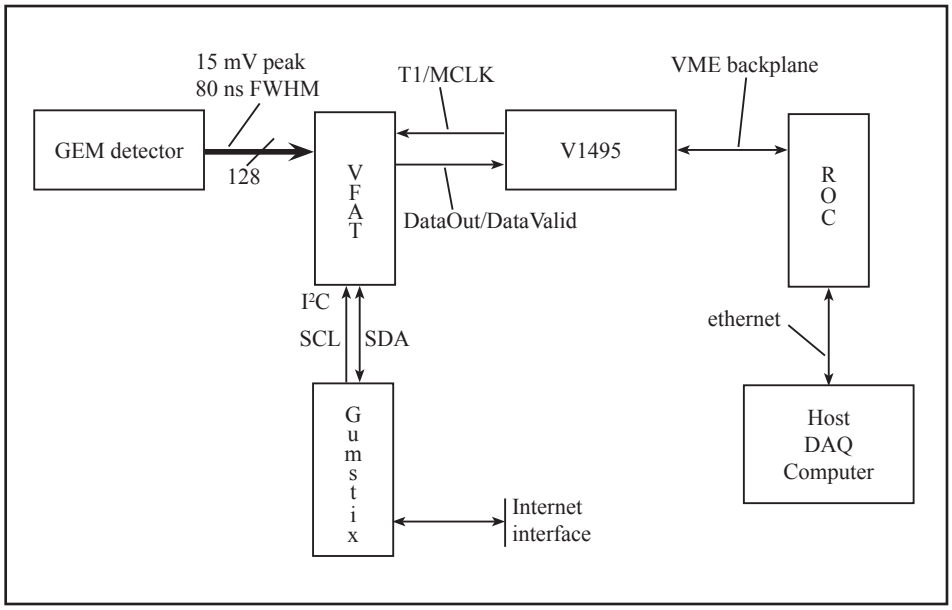

Figure 1-1: Block diagram for Region 1 detector electronics

The main portion of this thesis begins where the analog voltage signals exit the GEM detectors and ends where the data is transferred to the Linux DAQ computer. The most important topics, for which discussion is necessary to understand the full scope of this project, will be briefly touched upon in this introduction.

1.2 Region 1 Detector Description

The Region 1 Detector is a particle position detector for the Qweak experiment at Jefferson Lab. A graphical depiction of the Qweak experiment, along with the relative location of the Region 1 detector within this experiment, can be seen in Figure 1-2 on page 3. Very precise cylindrical position can be obtained from a particle by a process known as Gas Electron Multiplication (GEM). This process is a well-established technology that uses two different processes to obtain gains of ~10^6 charged particles.

The first process the GEM detector uses is ionization of an otherwise inert gas via inelastic scattering. These liberated charged particles are accelerated by a static electrical field of ~3 kV/ cm. Because either process can be responsible for

starting the ionization process, either high-energy charged particles or photons can be detected.

After the initial ionization process, ionized particles are accelerated towards insulating Kapton layers sandwiched between layers of copper. This copper-Kapton-copper assembly is then placed within a series of other similar assemblies and electronically biased at a very high static voltage. As electrons come into close proximity with the very high electric fields of the perforated holes in the Kapton assembly, even more electrons are released due to the acceleration the particles undergo thereby causing even more ionization events. Figure 1-3 on page 4 depicts these perforated holes along with the field lines that they generate. Because a vast majority of these ionization events occur at the high-field regions near the holes of the copper-Kapton-copper layers, detectors of this type are useful for having both large preamplification and higher spatial resolution than many other detectors.

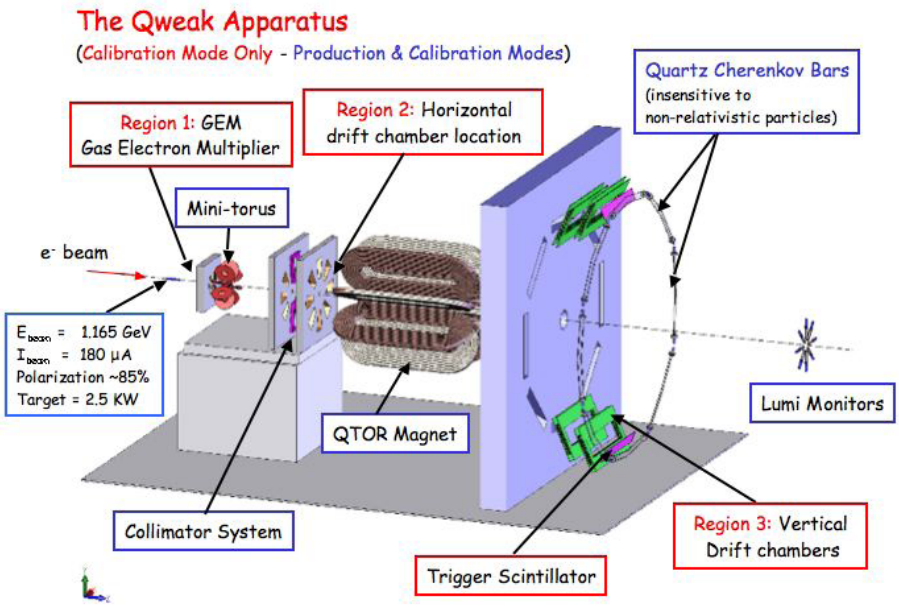

Figure 1-2: Graphical Depiction of Qweak experiment including the Region 1: GEM detector

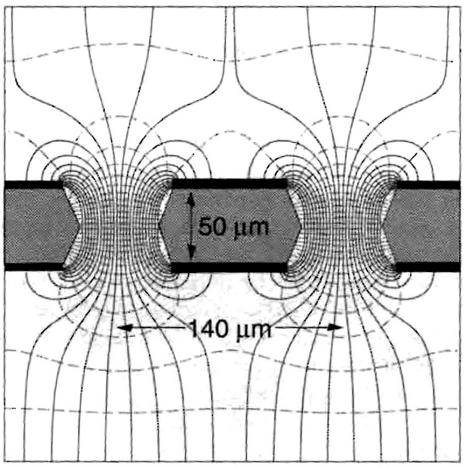

Figure 1-3: Perforated hole structure of copper-Kapton-copper layer[12]

1.3 VFAT2 Readout Card

A digital hit-or-miss comparison is based on a programmable threshold voltage after having passed through the preamplifiers and pulse-shaping networks. The VFAT cards are fast-acting integrated circuits (ICs) developed by CERN. Each one possesses the ability to process 128 channels simultaneously. The VFATs' discriminators used to identify charge levels are programmable via the VFATs' registers. All of the programmable registers of the VFAT cards are accessed using its I^2C interface.

1.3.1 Preamp/Filter/Comparator

The front-end of the VFAT consists of a transimpedance amplifier, an integrator, an adjustable differential threshold detector, a differential-mode comparator, and ends with a common-mode comparator for digital output. These different sections are shown in Figure 1-4 for demonstration purposes. Unfortunately, as will be discussed later in this section, this schematic is only for demonstration purposes and cannot possibly strictly adhere to the actual circuitry in the VFAT front-end because there is little correlation to the adjustable I^2C registers.

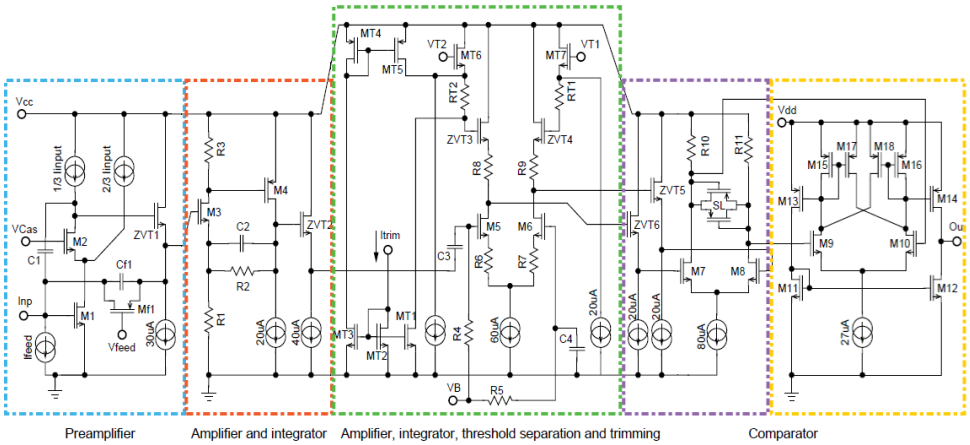

Figure 1-4: The VFAT2 front-end with different sections outlined for clarity

The transimpedance amplifier (shown in Figure 1-4 as surrounded by blue) converts a current (or charge) signal into a voltage signal. Adhering to the naming convention of these various amplifiers, one takes the amplifier gain as the output signal over the input signal. Therefore, taking the output signal (voltage) and dividing it by the input signal (current) provide us with the units of impedance and hence the name of the amplifier.

Per the literature on the VFAT, we know that the pre-amp is comprised of a cascode-configuration amplifier with an NMOS input transistor operating close to weak inversion. It can be surmised that operating this transistor near the weak inversion layer is necessary because this is precisely where the equivalent gate capacitance of the input transistor is at (or nearly at) a minimum – particularly at higher frequencies, i.e. 100Mhz. This is because the NMOS transistor equivalent gate capacitance is dominated by the series capacitance of the depletion layer at this point.[13][14]

Operating a transimpedance amplifier at a lower capacitance means that a greater voltage output gain can be achieved given the same amount of input charge. This follows if the input capacitance is the main (or one of the main) transimpedance gain factor. Also, the cascode configuration gives the amplifier very high gain with the only penalty being that the output voltage range decreases by the overhead of one overdrive voltage, e.g. $V_{gs} - V_{th}$.

This preamplifier stage also contains a tunable active-feedback capacitance. This provides a 13ns time constant for a 0.8 µA feedback current. Using this active-feedback one can theoretically set the lower level of accepted frequencies into the pulse-shaping network.

The second distinct stage of the analog front-end is another amplifier with an integrator (shown in Figure 1-4 as red). Using this stage, the upper end of the accepted frequencies for the overall amplifier can theoretically be set.

After the integrator, the signal passes through the adjustable threshold discriminator. DC signals are blocked coming into this stage by the input capacitors. This stage appears as a differential amplifier with adjustable loads placed in the differential current paths. This clever configuration allows for adjustment of the input voltage necessary to flip the respective voltage outputs taken from the drains of the input NMOS transistors shown in Figure 1-4 as M5 and M6. Setting the proper current for this differential amplifier and then shifting the input voltage will steer the current from one side of the differential amplifier to the other thereby switching the voltage at the outputs drastically. Changing V_{th2} and V_{th1} causes there to be more or less voltage necessary to shift the voltage outputs. During this project we have operated under the instructions that $V_{th2} - V_{th1} > 0$ is the proper setting for detecting positive input pulses and $V_{th2} - V_{th1} < 0$ is the proper setting for detecting negative pulses. Since we are detecting pulses of accumulated electrons, we have our threshold settings set to the latter.

The final two sections of the front-end detectors provide a differential-mode comparison (shown in purple) of the output voltages from the differential amplifier followed by a single-ended comparison (shown in orange).

One difficult aspect of this project is that the schematic in Figure 1-4 can only be used to indirectly determine the values of the adjustable I^2C registers, found as primary registers 0x02 - 0x07. For instance, it would appear that the integrator portion of the schematic is adjustable via the IShaper or IShaperFeed registers, but one cannot made a direct correlation between the two. Not knowing exactly how to adjust these registers to tune in or tune out particular aspects of our pulses forced us to treat this aspect of the Region 1 Detector as a black box. Fortunately, we did receive special instruction on how to set these values for our experiment from Paul Aspell, one of the designers for the VFAT. Note that these actual values varied from the default settings found in the VFAT2 Operators Manual. These values are listed in Table 1-1 below.

Register	Value
IPreampIn	142d = 0x8E
IPreampFeed	70 d = 0x46
IPreampOut	130 d = 0x82
IShaper	127 d = 0x7F
IShaperFeed	85 d = 0x55
IComp	100 d = 0x64

Table 1-1: Recommended I^2C Preamp Settings

1.3.2 T1 Command Interpreter

There are four basic commands, called T1 commands. Each one of these is comprised of an LVDS, 3-bit pattern sent in sync with the IC's external clock pulse (MCLK), used to direct the action of the VFAT.[1] It must be noted that no two T1 commands can arrive closer than 3 clock periods due to their encoded nature.

Pattern	Command Name	Function
0b100	LV1A (Level 1 Accept)	Trigger
0b111	CalPulse	Timing of calibration pulse
0b110	ReSync	Resynchronisation of all state machines
0b101	BC0	Bunch crossing zero identifier

Table 1-2: T1 Commands

The LV1A signal is the most commonly used signal; this signal tells the VFAT to search into its recorded memory for hit data and package this data into a serial packet to be read out on the LVDS DataOut line synchronously with the MCLK signal. The VFAT has two SRAM registers for "hit" information. The write address to the first, SRAM1, is incremented every clock cycle of MCLK. When a LV1A signal is sent to the memory module, the hit record contained in the current register minus the latency value, I^2C extended register 0x00, is transferred from SRAM1 to SRAM2. This hit record is then serialized, packaged, and sent out via the DataOut line.

For reference, the structure of the serial data reported per each LV1A signal is shown below.

1010	BC<11:0>	
1100	EC<7:0>	Flags<3:0>
1110	ChipID<11:0>	
Channel Data <127:0>		
CRC 16 checksum <15:0>		

Table 1-3: Structure of DataOut Serial Packet

The CalPulse signal is used to send a calibration pulse onto the signal lines of the VFAT. We were successfully able to verify that a pulse could be detected on each individual signal line of the VFAT ICs via this method. One can see the results in section "4.6 CalPulse" on page 65. Beyond this, the operation details are clearly described in the VFAT Operation Manual.[1]

At times it may be necessary to reset the values of the state machines. This happens automatically at the beginning every time the power is turned on to the VFAT. It is also possible to do this during normal operation by sending a ReSync pulse on the T1 line to the VFATs. This can be particularly useful between the time when the VFATs are first powered on and when actual hits are to be recorded. As described in “Readout Controller Library“, this signal is automatically sent to the VFATs when the function “v1495Reset()” is called on the Readout Controller.

Finally, rather than completely resetting all of the state machines contained in the VFAT, the BC0 signal can be sent to only reset the Bunch Crossing Number (BCN). The BCN is a 12-bit counter that increments once for every clock period. It might be useful, for example, to use this number to determine if there are missing LV1A requests in the data.

1.3.3 I^2C Communication

The physical mode of communicating with the configuration registers of the VFAT chip is performed via an I^2C protocol.

I^2C is a synchronous, full-duplex communication scheme albeit not a simultaneous full-duplex signal. The clock signal is always sent from the master microcontroller via the SCL line and the data is sent either by the master (in the event of a write or request) or by the slave (in the case of a read) via the SDA line. Both lines are sent via an open-drain configuration. This means that while the I^2C is being used by one device, another device can look at the I^2C lines to see if they are pulled low before proceeding to send data.

Each VFAT has a unique 7-bit I^2C address for which the first three MSBs indicate the specific VFAT and the four LSBs indicate which primary address is to either be written to or read. A description of how to set the I^2C address on the VFAT Breakout Board can be found in section “4.2.2 I^2C Address and Scan Enable Jumpers” on page 42. If one of the extended registers is to be written to or read from, one must first fill the “ExtRegPointer” primary register 0x0D with the address of the respective extended register. The data can then either be written to or read from the “ExtRegData” primary register 0x0E This allows the VFAT to contain an extra 135 registers.

There are many references on utilizing the I^2C interface. For more information on this common digital communication scheme, the reader is encourage to review the content of other references on the subject.

1.3.4 Gumstix Microcontroller

In order to communicate to the VFAT registers via I^2C and to know their respective contents for our experiments, it was determined that this could most efficiently be implemented by using a Gumstix Microcontroller. The name for Gumstix aptly describes these powerful microcontrollers as their size roughly resembles that of a stick of gum. These Computer-on-Module (COM) microcontrollers include many extensible modules including an I^2C communication board and an ethernet board.

The actual setup and implementation of these microcontrollers was performed by Brian Oborn. The Gumstix microcontroller actually uses a form of C-Shell from Linux and is supported by the OpenEmbedded Project. Thus, we were able to use the I^2C and CGI capabilities of this distribution of Linux to not only read and write to the I^2C registers on the VFATs but also to display the contents of these registers via a website hosted on the Gumstix controller.

A picture of one of our Gumstix computers is shown below in Figure 1-5.

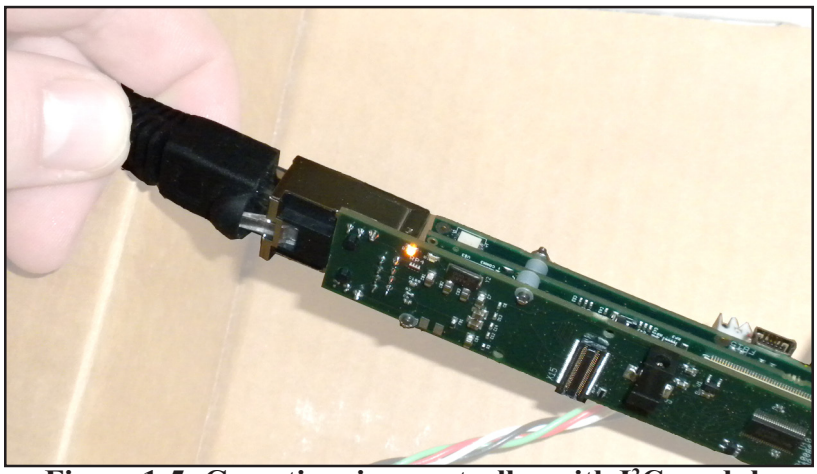

Figure 1-5: Gumstix microcontroller with I^2C module and ethernet module attached

1.4 V1495

The V1495 module contains two Field-Programmable Gate Arrays (FPGAs). The first FPGA is programmable by the user and is used as the main interface to the external I/Os on the faceplate of the V1495. The second FPGA controls the interfacing of the User FPGA and the VME backplane. Both FPGAs are members

of the Cyclone device family.[5] The complete operation and underpinnings of the V1495 by itself are worthy of a separate thesis; the Cyclone Device Handbook, Volume 1, contains 385 pages. Of this, an understanding of perhaps 10 percent is sufficient to comprehend, operate, and effectively alter the properties of the Cyclone FPGA for the Region 1 detector.

Programming and utilizing the V1495 is one of the main portions of this thesis and therefore has an entire chapter dedicated to the user-defined code that was implemented to query the VFAT chips and store that information until it is requested by the Readout Controller.

1.4.1 Readout Controller (ROC)

Part of using the VME protocol is that one rarely directly accesses one of the VME modules on the VME plane; this project is no exception. One actually communicates to the individual modules, i.e. the V1495, via a Single-board Computer (SBC) known as the MVME6100. For our project we called this the Readout Controller.[6],[7]

The MVME6100 is a very powerful SBC with a high-performance MPC7457 PowerPC processor capable of dictating instructions to all of the modules contained within a single VME crate via the VME backplane at 320MB/s. It also has the capability of communicating to other computers outside of the VME crate via its ethernet ports. Programs for the V1495 are actually cross-compiled on a separate Linux-based host computer and then downloaded to the MVME6100 via this ethernet connection. The target OS that we utilize on the MVME6100 is actually an embedded version of VxWorks which is also used in many scientific as well as commercial application, i.e. Linksys Routers.

As with many of the modules discussed within the introduction, the reader is referred to other resources for specifics on how to utilize this module.

1.4.2 CODA Readout

Data from our cosmic particle experiments as well as HRRL experiments were taken and processed by a software package known as CODA. This software package was developed for use at Jefferson Lab. It contains many libraries that are useful for capturing and storing data that is recorded using the VME system.

For more information, visit the CODA resource website.[15]

Figure 1-6 below shows the typical setup of our cosmic particle detection runs. On the left side we can see the Region 1 detector laying flat on a stand. In the lower right side of the picture we can also see the VME crate containing both the V1495 as well as the Readout Controller.

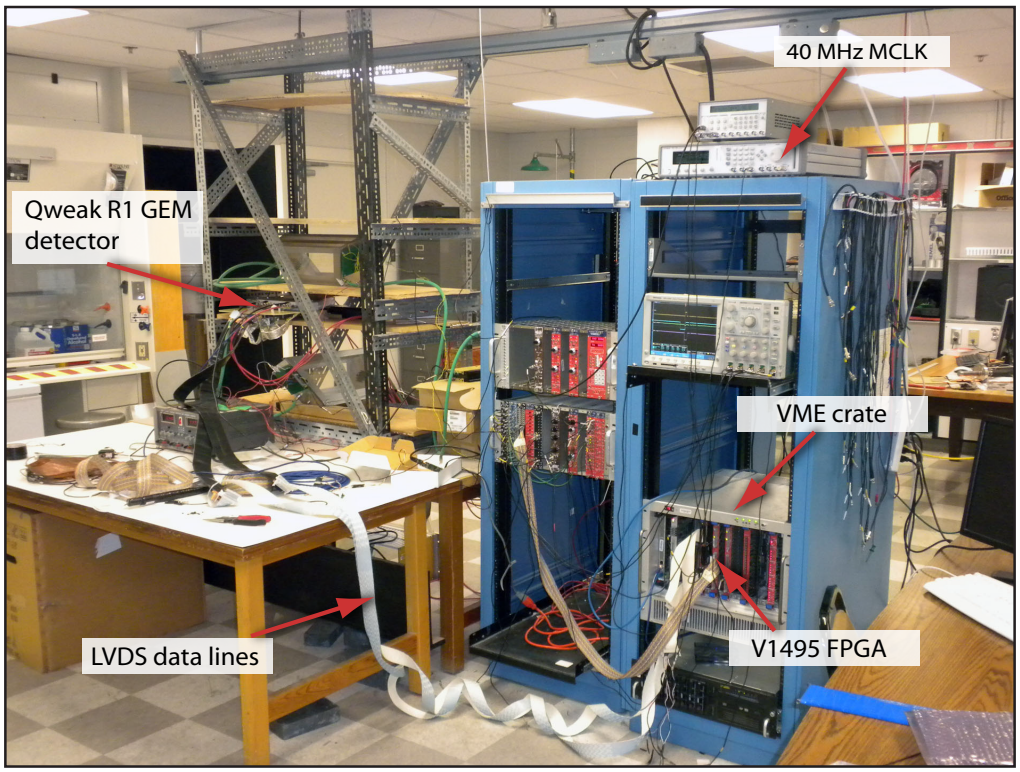

Figure 1-6: A typical setup of the Region 1 detector including VME crate