

I. Introduction

GlueX is a nuclear physics experiment which aims to discover how quarks and gluons are confined within hadrons. Quarks and gluons are elementary particles that are the building blocks of heavier particles like protons and neutrons. Protons and neutrons are the lightest members of a large family of composite particles called hadrons that make up the atomic nucleus. The physical theory that governs hadrons, quantum chromodynamics (QCD), tells us that hadrons can be grouped in a hierarchical table that is analogous to the periodic table of the elements. The hadron that plays the role of hydrogen in this table is known as a meson, which consists of a quark and an anti-quark bound by an elastic string of gluons. According to QCD, this string should be able to vibrate, giving rise to more energetic particles in the same family of the table. Mesons with vibrating glue are referred to as exotic mesons. Experimental evidence is now mounting for the physical existence of exotic mesons. The goal of the GlueX experiment is to produce exotic mesons in large numbers in order to unambiguously identify them, map their spectrum and measure their properties.

The GlueX experiment will produce mesons by colliding high-energy particles of light (photons) with a liquid hydrogen target. At the U.S. Department of Energy's Thomas Jefferson National Accelerator Facility, high-energy electrons from the accelerator are directed onto a piece of diamond crystal. As these electrons pass through the diamond, they lose some energy by radiating high-energy photons in a process known as bremsstrahlung. If the diamond is oriented such that the electrons

travel nearly parallel to the planes of atoms in the crystal, then the photons produced are polarized in the direction perpendicular to the planes, and the process is called coherent bremsstrahlung. Among the myriad particles produced by the collisions of polarized photons with protons in the liquid hydrogen target are the mesons of interest to the experiment. To separate these from all of the other radiation produced in the target, it is essential to measure the energy and time of each photon in the beam. The “photon tagger” accomplishes this by measuring the energy of the electrons as they exit the diamond. Electrons coming out of the back of the diamond radiator are deflected by a magnetic field into an array of optical fibers that glow (scintillate) when high-energy electrons pass through them. The scintillation produced by these electrons is then carried by waveguides to solid state light detectors called silicon photomultipliers (SiPMs).

The goal of this research project is to design and prototype electronics for the readout of signals from the SiPMs. The primary function of the readout electronics is to amplify the signals for subsequent digitization. Their second function is to provide the bias voltage that is necessary for each SiPM’s operation and to monitor critical environmental parameters. The project entails the design and layout of circuit boards to accomplish the above functions, and the production and testing of prototypes.

II. Nuclear Physics Background

Particle Physics

Matter is made up of tiny particles called atoms. For quite some time, it was believed that atoms were fundamental, indivisible particles. However, J.J. Thompson's experiments with cathode rays in the 1890s led to the conclusion that electrons, small negatively charged particles, were a component of all atoms. He proposed that the electrons existed inside a sea of uniform positive charge. This model, shown in Figure 1 became known as the "Plum Pudding" Model, because the electrons, "plums" were suspended inside a positive charge "pudding."




Figure 1: The plum pudding model of the atom, where electrons (yellow) are contained within a sea of positive charge (pink)¹

Rutherford's 1909 Gold Foil experiment called the plum pudding model into question, however, when he explained the deflection of alpha particles passing through a sheet of gold foil by predicting the presence of concentrated positive charge at the

¹ Figure adapted from http://en.wikipedia.org/wiki/File:Plum_pudding_atom.svg

center of the nucleus, surrounded by mostly empty space, and an appropriate number of orbiting electrons to balance the positive charge. Figure 2 shows Rutherford's basic concept of the atom.




Figure 2: Rutherford's concept of the atom. Positive charge is concentrated in the center, surrounded by empty space and orbiting electrons.

Rutherford's 1918 discovery that the positive charge of any atom is always approximately equal to an integer multiple of the positive charge of hydrogen, coupled with the fact that the same is not true for atomic mass, suggested that in addition to the positive charge in the center of the atom, there may be neutral particles as well. He coined the term proton to describe the positively charged particles, and the term neutron to describe the negatively charged particles.

Though Rutherford's 1918 analysis forms the basis of the modern model of the atom, many new things have been discovered since then. Quantum mechanics predicts

that electrons don't move in discrete orbits around the nucleus, but rather exist in a probability cloud given by solutions of the Schrodinger equation. It also explains the observation that certain quantities, called quantum numbers, obey particular sets of rules. Some examples of quantum numbers include electric charge, which always exists in integer multiples of the electron's charge, and spin, which always exists either in odd integer multiples of $\frac{1}{2}$, or as integers, depending on the type of particle.




Figure 3: The Modern Concept of the Atom, in Particular, Helium. Inset: The Helium Nucleus, Showing Two Protons and Two Neutrons

The standard model of particle physics is a theory which explains how protons and neutrons themselves are composite systems composed of more elementary particles called partons. Figure 3 shows an accurate depiction of the modern concept of an atom, but the standard model suggests that there is even more structure within the nucleus of an atom than just the grouping of protons and neutrons. Indeed, within the

protons and neutrons are several different species of partons known as quarks and gluons.

	I	II	III	
mass→	2.4 MeV	1.27 GeV	171.2 GeV	0
charge→	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0
spin→	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
name→	u up	c charm	t top	γ photon
Quarks	4.8 MeV	104 MeV	4.2 GeV	0
	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	0
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
	d down	s strange	b bottom	g gluon
Leptons	<2.2 eV	<0.17 MeV	<15.5 MeV	91.2 GeV
	0	0	0	0
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
	ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	Z⁰ weak force
	0.511 MeV	105.7 MeV	1.777 GeV	80.4 GeV
	-1	-1	-1	±1
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
	e electron	μ muon	τ tau	W[±] weak force
				Bosons (Forces)

Figure 4: Table of the quarks, leptons, and bosons - the particles of the standard model



Figure 4 shows a table of the particles incorporated within the standard model.

At present, all known nuclear properties and phenomena can be explained as interactions among these particles.

The standard model explains the preponderance of protons and neutrons in the present universe as the most stable form that quarks and gluons can assume at low temperatures. Combining two up quarks and a down quark yields a proton, while combining two down quarks and an up quark yields a neutron.

Similar to the protons and neutrons they constitute, quarks have a variety of properties. Figure 4 shows the most familiar of these, including mass, charge, and spin.

Less familiar properties include intrinsic parity, C-parity, and isospin. Intrinsic parity is a property which describes whether the quantum mechanical wave function of a particle

has reflection symmetry. The wave function of a particle with negative parity, if reflected about the origin, will experience an overall change in sign from the original wave function. A particle with positive parity would not experience this change in sign.

The parity of a particle or a system of particles is conserved for all interactions with similar reflection symmetry.

C-parity arises from the existence of antiparticles for each of the quarks and leptons in Figure 4. Antiparticles have the same mass and total angular momentum as their corresponding particles, but have the opposite charge.??????????

Isospin is another property of particles which describes the relationship between multiplets of particles which behave similarly in certain interactions. Though not at all related to real angular momentum or spin, mathematically, isospin behaves the same.

Certain groups of particles have a particular isospin, and each particle in the group has a unique projection of that isospin along a virtual axis. This allows similar particles to be described as various unique isospin states of a single particle. It turns out that only up and down quarks have isospin, with up having isospin $+\frac{1}{2}$, and down having isospin $-\frac{1}{2}$.

A fourth property unique to quarks is that of color charge. Despite its name, color charge has little to do with color in the optical sense, and little to do with

electromagnetic charge. Each type of quark, up, down, strange, charm, top, and bottom, comes in three color varieties. Appropriately, these varieties are named red, green, and blue. Antiquarks come in antired, antiblue, and antigreen varieties. The reason color charge is described as a charge is because just as electromagnetic charge leads to electromagnetic interactions, color charge leads to another type of interaction, strong force interactions.

For a particle to experience an interaction with any other particle, the two particles must both have a charge appropriate to the force by which they are interacting. Neutrons, which have no net electromagnetic charge, don't significantly interact with electrons and protons by the electromagnetic force. However, neutrons and protons are made up of quarks which have color charge, meaning that neutrons can interact with protons via the strong force. This explains why neutrons can be bound within the nucleus, even though they are not electromagnetically attracted to protons.


The strong force is carried by exchange particles known as gluons. As shown in Figure 4, the gluon is just one of four force carrying particles currently known to exist. The photon, which we observe as light, is the electromagnetic force carrier, and the Z^0 and W bosons carry the weak force. A fifth hypothetical particle called the graviton (not shown in Figure 4), has been postulated to carry the force of gravity. These four fundamental forces explain all known particle interactions. Of these forces, gravity and electromagnetism work over extremely long ranges compared to the strong and weak forces. The strong and weak forces are the only forces that are relevant on the size scale of the atomic nucleus.


The weak force, as one might expect from its name, is more than 10 trillion times weaker than the strong force. Among other phenomena, the weak force is responsible for the process of beta decay – where a neutron decays into a proton, an electron, and an electron antineutrino. While much more can be discussed about the weak force, further discussion would be beyond the scope of this paper.



The strong force, as its name suggests, is the strongest of the four fundamental forces. It acts only among and between quarks, antiquarks, and gluons, since these are the only particles which have color charge. Appropriately, the theory which describes the strong force is called quantum chromodynamics, or QCD. QCD explains that gluons simultaneously carry both color and anticolor, meaning that they can interact not only with color carrying quarks and antiquarks, but also with themselves. Additionally, the fact that color and anticolor are carried by gluons means that quark color changes during any exchange of gluons. Since exchange of gluons is what holds quarks together inside the nucleus, it would seem that protons and neutrons would come in several varieties with different net color, perhaps even changing color with time. Experiment shows, however, that this is not the case. Protons and neutrons, along with a group of other composite particles known as hadrons, exist only in color neutral configurations.





Figure 5: Quark color combinations that form color neutral hadrons

A color neutral configuration, one with no net color charge, can be formed in two ways, as shown in Figure 5. First, a red quark, a blue quark, and a green quark may be bound together to make a color neutral, three quark hadron. Second, a quark and its complementary antiquark may bind together to make a color neutral, two quark


hadron. The first of these two types of hadrons is called a baryon, and the second is called a meson. 

An interesting consequence of the color neutrality of hadrons arises when color neutrality is considered along with the fact that the strong force carriers, gluons, carry color charge. When quarks inside a hadron are separated for any reason, rather than the force between them decaying by the inverse square law like in electromagnetism, the gluons between two separated quarks form a tube,  ~~aligning so as to keep a net neutral color, while at the same time maintaining~~ a constant force between the quarks.

When two quarks are separated far enough from each other that this gluonic flux tube  ~~can no longer be maintained~~, a new quark/antiquark pair will spontaneously  form in the middle of the flux tube to maintain the color neutrality two, now separate, quarks. This behavior is known as confinement, and it prevents quarks from ever being observed as isolated objects.

The gluonic flux tube which holds together quarks and antiquarks can,  in theory, be excited and begin to vibrate,  just like a plucked string or a rubber band. ~~In the case of mesons,~~ mesons with vibrating gluonic flux tubes are referred to as hybrid or exotic  mesons. It is these exotic mesons that the GlueX experiment  hopes to learn more about.

GlueX

The GlueX experiment is a high energy nuclear physics experiment which will be conducted at the Thomas Jefferson National Accelerator facility in Newport News, VA, 

~~which is known as JLab for short.~~ The purpose of the GlueX experiment is to map the spectrum of hybrid mesons. In order to produce hybrid mesons GlueX will use a process known as photoproduction, which involves bombarding hadrons with high energy photons.

To produce these high energy photons, GlueX will make use of JLab's 12 GeV electron accelerator. The photons will be produced by passing the electron beam through a thin diamond wafer. As the electrons pass through this wafer, they experience a quick deceleration, during which they radiate lost kinetic energy in the form of high energy photons. This process is known as bremsstrahlung radiation, ~~and~~ since the diamond wafer, ~~like all diamonds,~~ has a regular crystal structure, the radiated photons ~~in this case~~ are ~~both coherent and~~ polarized. Hence, the process is known as coherent bremsstrahlung radiation. These bremsstrahlung photons ultimately collide with a liquid hydrogen target, ~~where~~ producing a ~~huge number of particles,~~ among which GlueX hopes to find hybrid mesons.

In order to be able to identify hybrid mesons among the radiation produced in the hydrogen target, it is essential to know the energy of the bremsstrahlung photons. Direct measurement of the energy using some type of detector would destroy the photons in the measurement process. Therefore, photon energies are determined indirectly by measuring the amount of energy lost by the electrons as they pass through the diamond radiator. A device known as the photon tagger is responsible for separating the electrons from the bremsstrahlung photons, and then measuring the post-radiation

energy of the electrons. Design of electronics for this photon tagger was the focus of my research.

III. Functional Description of the Photon Tagger

After passing through the diamond radiator, the electrons are deflected out of the beam of radiated photons by a strong magnetic field. The permanent magnet which produces this field is called the tagger magnet.

The separation of electrons from photons is possible because of the Lorentz force, Equation (1), which describes how particles behave in a magnetic field.

$$\vec{F} = q\vec{E} + q\vec{v} \times \vec{B} \quad (1)$$

In the case of the tagger magnet, there is no electric field \vec{E} , so the first term of Equation (1) is always zero. The bremsstrahlung photons in the tagger magnet field have no electric charge q , so total force acting on them, \vec{F} is zero, and they pass through the field unaffected. The electrons, however, have a nonzero charge q , a velocity \vec{v} , and are moving through the nonzero tagger magnet field \vec{B} . Therefore, they experience a deflection, separating them from the bremsstrahlung photons. Since the electric charge q is the same for all electrons, and the tagger magnet's field \vec{B} is uniform, the only variable affecting how much an electron is deflected is the electron's velocity \vec{v} .



This magnetic deflection force on the electrons always acts perpendicular to the direction of motion, thanks to the cross product in Equation (1), meaning that the electrons will follow a circular path through the magnetic field. To determine the radius of the deflection, it is necessary to consider the centripetal force the deflection produces. Centripetal force is given by Equation (2), where m is the electron's relativistic mass (assumed to be constant over the expected electron energy spread), and r is its radius of curvature:

$$\vec{F} = \frac{m\vec{v}^2}{\vec{r}} \quad (2)$$

Substituting Equation (2) into Equation (1), it is possible to solve for the radius of curvature:

$$\vec{F} = \frac{m\vec{v}^2}{\vec{r}} = q\vec{v} \times \vec{B}$$

$$\vec{r} = \frac{m\vec{v}^2}{q\vec{v} \times \vec{B}} = \frac{m\vec{v}}{q\vec{B}} \quad (3)$$

Equation (3) shows that the radius of curvature will be greater for electrons with greater velocity. This means that electrons passing through the tagger magnet's field become spatially separated based on their energy.

To exploit the spatial separation of electrons based on their energies, an array of scintillating fibers is located in the path of the electrons, just outside the tagger magnet. These fibers are 2 mm x 2 mm in size, and roughly 2.5 cm long. They are arranged in a

5x100 fiber grid such that the electron energy spectrum maps along the longer side of the grid. The 5 fiber height of the grid was chosen to capture electrons that may enter the tagger with a small vertical displacement above or below the horizontal axis. Though vertical displacements convey no information about electron energies, they are useful to calibrate the alignment of the scintillating fibers with the incoming electrons. Figure 6 shows a drawing of the array of scintillating fibers.



Figure 6: Drawing of the 5x100 scintillating fiber array

When an electron passes through a scintillating fiber, the fiber scintillates, releasing a flash of light due to **coulomb interactions with the electron**. This scintillation is captured and carried by a matching 5 x 100 array of clear optical waveguides to the tagger electronics, where the scintillation is converted to an electrical pulse, amplified and sent off to be recorded for later analysis. The electrons themselves are discarded into a ~~concrete~~ beam dump after passing through the scintillators. Figure 7 shows a functional drawing of the **photon tagger**, including several representative scintillators and waveguides.






Figure 7: Functional drawing of the photon tagger, showing a group of 25 scintillating fibers, a few representative waveguides, and the tagger electronics array

As shown in Figure 7, the tagger electronics consist of three unique circuit boards: the amplifier board, the control board, and the backplane. Each set of these three boards is responsible for monitoring twenty-five waveguides for scintillation. The board which directly interfaces with the twenty-five waveguides is the amplifier board. It contains photosensors, called SiPMs, to detect scintillation and convert it to an electrical pulse, amplifiers to boost the SiPM signals, and summing circuitry to combine signals that come from the each fiber column, since the vertical displacement of electrons conveys no information about electron energies.

The amplifier board is connected by way of a ~~DIN 41612 Eurocard~~ connector to a backplane, which supplies power, and also provides outputs for signals from the amplifier and summing circuitry. Attached to the backplane through another ~~DIN 41614 Eurocard~~ connector is a digital control board, responsible for adjusting amplifier gain and also monitoring temperatures and voltages in the board array. This digital control board interfaces with a computer by a standard Ethernet interface to allow for easy monitoring and control of the tagger.

IV. Design of the Tagger Electronics

Digital Control Board

Component Selection

The digital control board was the first board designed for the photon tagger. The first step in its design was to determine exactly what it needed to be able to do, and select the most appropriate components to provide the necessary functionality.

Since one of the primary tasks of the digital control board is to adjust SiPM bias voltages, a multichannel digital-to-analog converter (DAC) capable of outputting a wide range of voltages had to be selected. For this purpose, we selected the Analog Devices AD5535 DAC. Though expensive, it has 32 independent channels, each able to output between 0 and 200 V, and source up to 700 μ A of current. The DAC's ball grid array (BGA) package helps keep it from taking up excessive space on the board, and it's serial

programming interface allows it to be easily controlled by a microprocessor or a field programmable gate array (FPGA.)

A field programmable gate array is a hardware device that can be rewired on the fly. Using ingenious combinations of field effect transistors and logic gates, an FPGA electronically rewires itself at power on to create a custom hardware based logic device. Hardware based solutions in general provide more speed and stability than software based designs. While the digital control board doesn't require extraordinarily fast clock speeds, the negligibly higher cost of an FPGA was still justifiable to produce a board can operate stably and be ready for possible future applications.

We selected the Spartan-3A XCS50A FPGA in a very thin quad flat pack (VQFP) package to use on the digital control board. The VQFP package minimizes the size of the FPGA to a mere 1.6 x 1.6 cm. Despite this small size, however, the FPGA has 50,000 internal logic gates, and supports a total of 68 input/output channels. Though 50,000 internal logic gates is considered few for an FPGA, software analyses with Xilinx's own programming platform suggested that for our needs, 50,000 would be more than enough.

With the FPGA selected, we next had to select a component to allow the FPGA to report back to the master computer. While many means of communication were considered, Ethernet was determined to be the most practical. The wide availability of cheap Ethernet switches allows all the digital boards to plug into a nearby switch with a single wire uplinking the entire array of tagger electronics boards to the main computer.

The perfect Ethernet controller for this role was the Silicon Laboratories CP2201. The CP2201 is a 10-Base T Ethernet controller in an incredibly small QFN package that is a mere 5 mm x 5 mm. To achieve this small size, the chip uses a multiplexed address/data bus, which conveniently also minimizes the number of I/O pins that have to be used for controlling the CP2201 on the FPGA. To physically connect the CP2201 to an Ethernet cable, we selected a Conn Pulsejack J0012D21. This jack has a 1:1 inductive coupling, perfect for short Ethernet wire runs, while also keeping cost to a minimum.

While these components take care of the basic control and communication features of the digital control board, the board is also supposed to be designed with certain monitoring capabilities. In particular, the digital control board needs to be able to monitor its own voltage levels and temperature, as well as the temperature inside the tagger darkbox. To accomplish voltage monitoring, we selected the Analog Devices AD7928 12-bit analog-to-digital converter. This 8 channel ADC is able to measure voltages between 0 and 5V, making it perfect for monitoring most of the voltages on the digital control board. In addition, the AD7928 uses the standard Serial Peripheral Interface (SPI) bus, which minimizes the number of PCB traces that are necessary to connect it to the FPGA.

For temperature monitoring, we selected the Analog Devices AD7314 temperature sensor. Though it would have been possible to dedicate a channel on the ADC to reading out a thermistor, a digital temperature sensor like the AD7314 is not very expensive, takes up minimal space on the board, and is compatible with the same SPI bus interface used by the ADC. Since SPI devices can be daisy chained together, the

number of FPGA I/O ports needed to readout SPI devices is not changed by the addition of an SPI temperature sensor. In addition, a separate temperature sensor left an extra channel on the ADC that is used to readout a thermistor on the distinctly non-digital amplifier board.

In order to connect the digital board to the backplane and the amplifier board, we chose a 48-pin DIN 41612 Eurocard connector. There were many reasons for selecting the DIN 41612 board interconnect system, including the fact that the connectors are keyed to prevent backwards insertion, housed in sturdy plastic to prevent accidental misalignment and bent pins, tightly fitting enough to hold the amplifier board upside down beneath the backplane, and able to safely isolate high voltages even beyond the potential 200V associated with each DAC output channel.

PCB Layer Selection

Before any schematics were drawn or any traces were laid, I first did an audit of the datasheets of the components that had been selected for the digital control board to determine what voltages they needed to operate. With this information in hand, I determined the best way to provide these voltages, either by bringing them in directly by way of the backplane, or by using voltage regulators on the digital control board to produce them from other available voltages. The voltage requirements, and how to provide them, are shown in Table 1.

Table 1: Voltage Requirements of the Digital Control Board

Voltage	Current Required	Used For	How to Supply
+5V	21.7 mA max	DAC internal amplifiers, analog ADC subcomponents	Direct by way of backplane
+3.3V	335 mA max	DAC internal temperature diode, AD7314 power, ADC logic, CP2201 power, auxiliary FPGA power, EEPROM power	Produce from +5V using voltage regulator
+2.5V	1 μ A max	ADC reference voltage	Produce from +5V using voltage reference
+1.2V	<10mA typ.	FPGA internal logic	Produce from +5V using voltage regulator
-5V	3.5 mA max	DAC internal amplifiers	Direct by way of backplane
+210V max	???	DAC high voltage supply	Direct by way of backplane

The rationale for deciding to supply only +5V, -5V, and +210V (max) to the digital control board from the backplane was as follows: There are not enough free pins on the Eurocard connector to supply all of the voltages directly. In addition, since +1.2V and +2.5V are used each by only one component, it would be wasteful to purchase separate power supplies for those voltages. While +3.3V is used by many components, it can be easily generated from +5V using a voltage regulator, and introduces a lesser noise risk than there would be from producing +5V with a +3.3V supply using a switching regulator. For the same reason, it is better to supply -5V directly, and obviously, +210V (max).

Once the power supplies were selected, the next task was to determine how to distribute the power to components on the board. Typically, PCBs supply power to components by means of copper planes inside the PCB. Connections to an electrical ground are also made using copper planes. Especially on digital boards with rapidly fluctuating logic cells, this helps to stabilize supply voltages by keeping an entire plane of charge available right near the components rather than attempting to suck charge through a small trace from a distant power supply. Also, large copper planes help shield against crosstalk between traces on opposite sides of the PCB.

Since adding internal layers to a PCB ultimately increases the fabrication cost, I had to determine which voltages were essential to supply by a plane, and which could be supplied to components using traces. Since +5V and +3.3V are both used by multiple components spread all around the board, it made sense to dedicate two planes to these voltages. An additional two planes were to be dedicated to both a digital and a separate analog ground. This layer arrangement resulted in a 6 layer PCB, which is the maximum number of layers that is considered standard (and hence cheap to produce).

However, this still left three voltages to deliver without planes. I determined that it was acceptable to deliver the +2.5V and the +1.0V (max) lines by traces, since the current demands would be fairly stable. The +1.2V line, however, would be powering digital FPGA logic, and would be more problematic to deliver by a trace. The ultimate solution to the power plane dilemma didn't come until much later, when most of the components had been placed into the PCB layout. I was able to place the FPGA in a

location away from all components requiring +5V, and create a +1.2V island within the +5V power plane. The final layer stackup looked like this:




Figure 8: Digital Control Board Layer Stackup

Wiring the DAC

The DAC was the first component to be “wired.” The term wired, in the context of this paper, refers to the process of drawing up schematics and PCB layout files using a computer aided design program called Altium Designer.

More than any of the other digital control board components, the DAC needed to be placed in a position that would facilitate easy routing of PCB traces to the backplane and ultimately the amplifier board. Therefore, logically, I placed it at the bottom of the PCB, as close as possible to the Eurocard connector.




Figure 9: The yellow square (U3) is the Analog Devices AD5535 DAC. To facilitate routing of the its many outputs, its lower edge was placed only 360 mil away from the Eurocard connector (P2).

The decision to place the DAC so close the Eurocard left insufficient space to for vias to move any of the bias voltage traces to the back side of the board. This was not necessarily a problem, though it did impose a requirement that other traces connected to the Eurocard had to be either be routed around the Eurocard to row “a” (see Figure 9), or sent in on the back layer. Fortunately, there was ample space on the rest of board to allow for passing out these outgoing signals to the back layer.

While routing the bias voltage traces was the greatest challenge of the DAC wiring, there were also several other supporting components that had to be wired in the region surround the DAC. A host of power conditioning circuitry had to be placed

around the DAC, consisting of decoupling capacitors for power stability, as well as diodes to help with power supply sequencing. For decoupling, I selected both 10 μF and 0.1 μF capacitors to place near each power input to the DAC. The high voltage input has only one decoupling capacitor near the DAC (labeled C2). This capacitor is 0.1 μF . However, there is a 10 μF electrolytic decoupling capacitor located farther away from the DAC in the upper right corner of the board. The rationale for placing this capacitor farther away is that there simply was not enough room in the region of the DAC for such a large, through hole, capacitor. A smaller surface mount device could not be used because no capacitors with 10 μF capacitance and a 210V rating exist. Since a 10 μF electrolytic capacitor has a fair amount of inductance anyway, placing it farther away shouldn't make it a whole lot worse than if it were placed closer to the DAC.




Figure 10: Close-up of the DAC and its Power Circuitry

A few more components worth noting in the area of the DAC are resistor R1, and resistors R9-R11. Resistor R1 is a current limiting resistor attached to the DAC's internal temperature sensing diode. +5V is internally applied to the anode of this diode, and it has a typical diode drop of 0.65V. The resistor is a 270K resistor, which, given a maximum of +4.35V will allow through a safe 16 μ A of current. Since the diode drop is temperature dependent, reading out the voltage at the diode/R1 junction using ADC allows the board to monitor the DAC's internal temperature.

Resistors R9-R11 provide a means by which the board may test the basic functionality of one channel of the DAC. Resistors R9-R11 function as a voltage divider, reducing the DAC's output voltage to a range readable by the 0-5V ADC. Since full scale on the DAC ranges from 0-200V, the divider needed to divide the voltage approximately by a factor of 40. Picking exact resistor values for this voltage divider was difficult, because it needed to be stiff enough so that the leakage current of the ADC would not perturb the voltage level in the divider, but use a small enough amount of current so that the channel would still be usable as a SiPM bias voltage source. In addition, given that the DAC could output as many as 200V the divider has to be able to withstand high voltages without exceeding the breakdown voltage of the resistors.

I found that there were no small surface mount resistors capable of withstanding voltage drops on the order of 200V, and therefore decided to use three resistors rather than the typical two. This way, the voltage drop across any one resistor is far less than 200V. I then calculated the current that would be used by the voltage divider at all possible DAC output voltages. Knowing the current requirements of different

configurations of the divider allowed me to select the one that would still leave sufficient current left to use the channel as a SiPM bias voltage.

As shown in Figure 11 below, the current required increases dramatically as the divider is changed to provide greater precision in the measured voltage. Since the DAC is capable of outputting 700 μA maximum per channel, the divider had to be designed so that it never would create a current draw of more than 700 μA . The 0.01V precision line in Figure 11 does indeed never exceed 700 μA . Precision of 0.01V at the readout junction of the divider corresponds to 0.4V precision in measurements of the DAC voltage, which we determined to be an acceptable error for testing the functionality of a DAC channel. To achieve this precision, I used two 200K resistors in series, followed by a 10.2K resistor to ground. Readout occurs between the second 200K and the 10.2K resistor, providing the 40:1 resistance ratio that is needed. See Figure 12 in Wiring the ADC for a schematic.

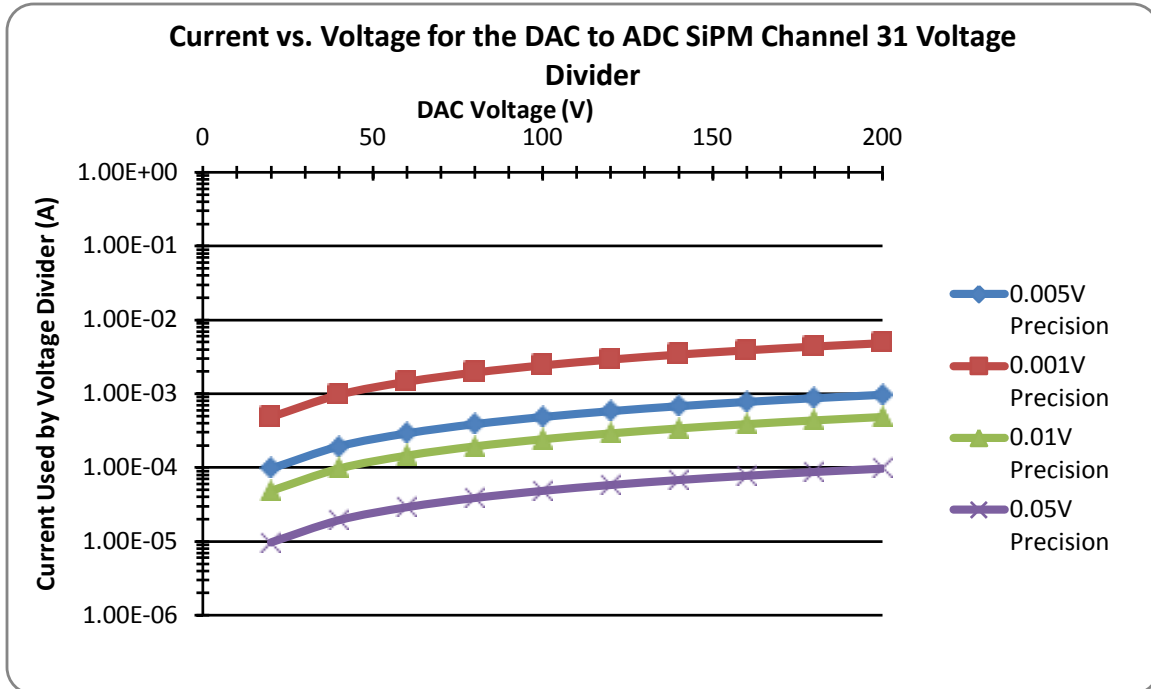


Figure 11: Current/Voltage Relationship for the DAC to ADC Voltage Divider. Precisions indicated in the key refer to the voltage drop caused by the ADC's leakage current at the readout junction of the divider. Multiply by 40 to obtain the potential error in the measured DAC channel voltage.

Wiring the ADC

Compared to the DAC, wiring the ADC was quite simple. The ADC required several power connections all 5V or less, decoupled with similar 0.1 and 10 μ F capacitors to the DAC. The lower voltage of these power connectors meant that no large electrolytic capacitors were needed.

The ADC also required several connections to the FPGA, and of course, inputs of the various voltages that it is to measure. Figure 12: Schematic of the ADC Wiring Figure 12 below shows the basic wiring of the ADC.




Figure 12: Schematic of the ADC Wiring

A description of the VIN measurement channels follows in Table 2.

Table 2: Description of Voltages Measured by the ADC

ADC Channel	Voltage Measured	Notes
VIN0	-5V Supply	Since the ADC has a 0-5V range, a voltage divider is used to bring the -5V up above ground to be compatible with ADC.
VIN1	DAC thermal diode	
VIN2	+3.3V	
VIN3	+5V	Though it at first seems impossible for the DAC to measure the voltage of it's own power supply, it actually can since the DAC uses a separate 2.5V reference for voltage comparison.
VIN4	Amplifier board RT1	This is passed through the Eurocard to the amplifier board, where it measures the voltage drop across a thermistor, RT1.
VIN5	+1.2V	

VIN6	Amplifier board VCC	This is passed through the Eurocard to the amplifier board, where it measures the VCC power supply for the SiPMs.
VIN7	DAC channel 31 output	See Wiring the ADC for details.

The voltage divider on channel VIN0 was not nearly as difficult to design as the voltage divider for the DAC on channel VIN7. With the DAC channel, there was a strict 700 μ A current limit, and a range of voltages exceeding two orders of magnitude. The -5V to ADC voltage divider was designed with the assumption in mind that the -5V supply should be able to source enough current to compensate for any current used by the divider, and that the range of voltages needing to be measured would be on the order of only a few percent above or below the nominal -5V. Therefore, I selected a 33.2K and a 100K resistor for the divider, and placed the divider between -5V and +5V (rather than -5V and ground), producing a readout junction voltage of +2.5V when both +5V and -5V are at their nominal values. While at first glance this appears not to provide an independent readout of the -5V supply, it in fact does, since +5V is also separately read out on channel VIN3.

Wiring the Ethernet Controller

The Silicon Laboratories CP2201 Ethernet controller is available from the manufacturer in two variants. One, a 9 mm x 9 mm QFN package, has a total of 48 pins, including independent 8-bit addressing and data lines to the FPGA. The other variant is a 24 pin QFN package with a footprint size of merely 5 mm x 5 mm. The smaller variant

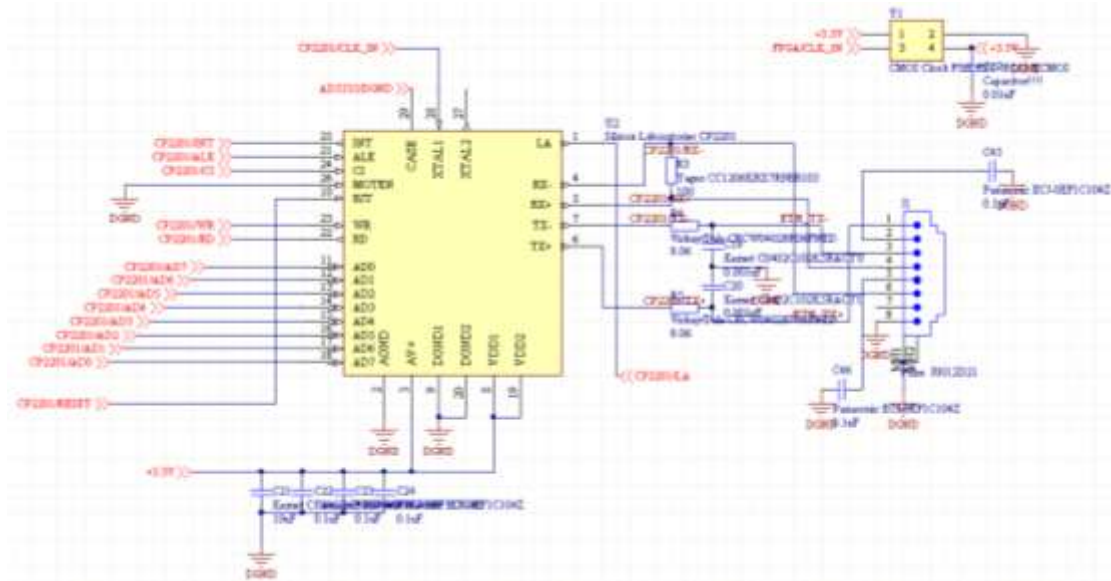


Figure 14: Schematic of the Ethernet Controller and Ethernet Jack

The above schematic details how the Ethernet controller is connected. Pins AD0 through AD7 are the 8-bit multiplexed address/data bus described previously; they connect to the FPGA. The INT (interrupt request), ALE (address line enable), CS (chip select), and RST (reset) pins also connect to the FPGA and help to facilitate communication over the address/data bus. The pins along the bottom of the component symbol, AGND, AV+, DGND, and VDD are appropriately decoupled power inputs and grounds, and the RX/TX pins on the right side of the symbol are where the controller connects to the Ethernet jack. Between the Ethernet controller and the Ethernet jack are several resistors and capacitors which are specified in the Ethernet jack's data sheet. These help ensure the signals produced by the controller are consistent with the 10BaseT Ethernet specification.

Pins XTAL1 and XTAL2 on the Ethernet controller are used for clocking the device. The data sheet specifies that pin XTAL1 may be used either as an input from a

CMOS clock, or a crystal oscillator. In the latter case, pin XTAL2 is used as an inverting driver. Originally, I selected the latter option, but in order to share a single clock between the Ethernet controller and the FPGA, we eventually decided it was better to replace the crystal oscillator with a CMOS clock. This is discussed in more detail below* in _____.

Wiring the FPGA

As the heart and soul of the digital control board, the FPGA was a particular challenge to route. Of the FPGA's 100 pins, 74 are used in the control board's design. Since the signals from these pins go to components all over the board, it was important to select input/output pins appropriately to avoid excessively crisscrossed traces.




Figure 15: Close-up of the FPGA, Showing Front (Red) and Back (Blue) Layer Traces

Since the spacing between FPGA pins is less than 20 mil, there is precious little space for vias in the region immediately surrounding the component. To ease this lack of space, I made use of vias under the body of the FPGA. The bottom of the FPGA is non-

conducting, so this poses little risk. Most of the pins connected by under-body vias are either power or ground pins, and connect directly to planes within the PCB. Some are in fact signals, however, and fan out on the back of the PCB as shown by the blue traces in Figure 15 above.

The FPGA shares a CMOS clock with the Ethernet controller in order to save space on the board. Implementing a shared clock proved to be more difficult than anticipated. First, while the Ethernet controller can accept a signal directly from a crystal oscillator, the FPGA needs a CMOS compliant signal. Simply connecting a single CMOS clock to the clock inputs of both devices is not possible because the capacitance of the two devices in parallel distorts the signal beyond recognition. The easiest solution to this problem was to send the CMOS clock signal into the FPGA, and then have the FPGA output the signal to the Ethernet controller using one of its I/O pins.

Figure 16 shows a close up of the connection between the CMOS clock, the FPGA, and the Ethernet controller. The CMOS clock is component Y1. It has only 4 pins: 1 ground, 1 power, 1 enable, and 1 clock output. The hatched trace going from pin 3 on the clock to pin 27 on the FPGA (bottom) is what connects the two components. Another hatched trace, leaving from pin 37 on the FPGA, brings another FPGA generated clock signal to the Ethernet controller (U2).




Figure 16: Detail of CMOS Clock Connection to FPGA and Ethernet Controller

The idea of having the FPGA generate a separate clock signal for the Ethernet controller seems like a poor design choice, especially since the documentation for the Xilinx Spartan-3A 100-pin VQFP does not designate pin 37 as optimized for clock signals. However, clock skew introduced by the use of a non-clock pin on the FPGA should not cause any serious problems because of the parallel nature of the data connection between the FPGA and the Ethernet controller. The Ethernet controller communicates with the FPGA using a non-clocked parallel interface. Reads and writes are performed using the ALE (address line enable), WR (write enable), and RD (read enable) pins on the Ethernet controller, and since the interface is not serial, the information on the address/data bus changes only when either device requests it. The clock signal needed by the Ethernet controller is used only for clocking the Ethernet signal itself.

Apart from the clock, another critical component of the FPGA system on the digital control board is the EEPROM. The Xilinx EEPROM stores the configuration data for the FPGA, and transmits that information to the FPGA at power on. Figure 17 shows the connections between the FPGA and the EEPROM.




Figure 17: EEPROM (U5) and the FPGA

The biggest challenge with EEPROM routing was trying to weave the traces through the jungle of other traces that connect to the FPGA. Unlike other many of the other components which could be connected to any of the 68 I/O pins on the FPGA, the EEPROM had to be connected to dedicated configuration pins on all four sides of the FPGA. Careful routing, and in some places multiple vias, were necessary to make all of the connections between the EEPROM and the FPGA.

Simply connecting the EEPROM to the FPGA was not all that had to be done, however. Since the EEPROM is soldered directly to the PCB just like all the other

components, there is no way to reflash it after the boards are manufactured unless an interface is provided on the PCB. Fortunately, this was not difficult to implement. The EEPROM uses a JTAG interface to communicate with the FPGA. Conveniently, JTAG is designed to allow easy connections from off the board using appropriate ribbon cables. With this capability in mind, I placed a 14-pin JTAG header on the board to allow for easy post-assembly reflashing of the EEPROM. This header also allows for hot reprogramming of the FPGA, bypassing the EEPROM entirely. The EEPROM and the FPGA are simply daisy chained together in a loop between the header's TDI and TDO pins. A close-up of the header is below in Figure 18. To see the header in context, look for it in Figure 9 on page 23, Figure 10 on page 24, or Figure 15 on page 32.

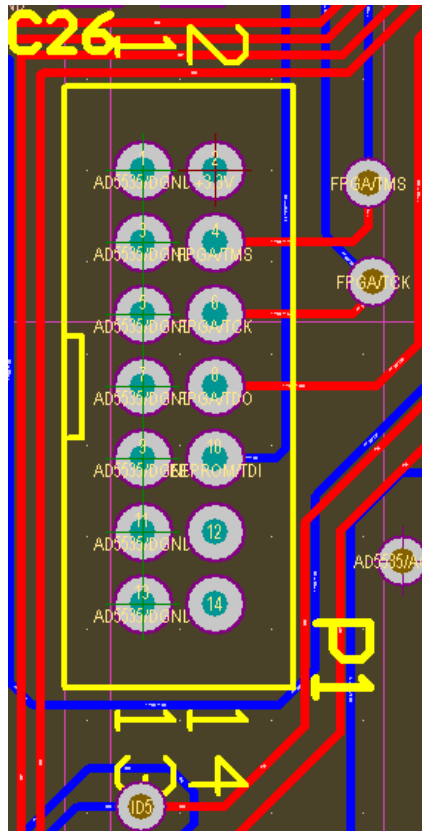


Figure 18: JTAG Header

One final noteworthy feature of the FPGA design is the use of unique location identifier traces that allow the FPGA to be aware of where it is positioned in the tagger circuitry array. An 8-bit active-low location ID bus runs between the FPGA and the Eurocard to the backplane. Once on the backplane, the bus connects to a set of jumpers which can be used to pull down any of the bits to ground in order to assign each card its own unique identification number. This number is transmitted within all Ethernet packets that leave the control board so that the main computer can tell the boards apart, and understand where they sit in the array. While it is possible to accomplish the same using the Ethernet MAC addresses of the control boards, we opted for a system of jumpers on the backplane so that a control board can be replaced on the fly without the

need to reconfigure a MAC/position mapping table. In Figure 19 below, pins 70, 71, 72, 73, 77, 78, 84, and 86 are used for the location ID bus.




Figure 19: Close-up of Location ID Pins on the FPGA

Backplane

As described in the Functional Description of the Photon Tagger, the backplane's primary purpose is to connect the digital board to the amplifier board, and to provide a means to route SiPM signals off the tagger electronics array for digitization. Figure 20 shows an overview of the backplane.




Figure 20: CAD View of the Backplane

PCB Layer Selection

The selection of layers on the backplane was driven by the power needs of the digital control board and the amplifier board. In order to accommodate as many voltages as possible without dramatically increasing production cost, we decided to make the backplane a 6-layer board like the others. This provided two outer signal layers and four internal power planes. The voltages assigned to the available planes were determined by the control board's and the amplifier board's current requirements. Voltages with high current loads were assigned to the planes, and voltages with lower current requirements were assigned to traces.

Since the digital control board's primary power source is the +5V power supply (see

Table 1 on page 20), I assigned +5V to one of the power planes. The digital ground from the control board was given its own plane as well. The analog ground from the digital control board was combined with the ground from the amplifier board into a single analog ground plane, and the final internal plane was reserved to carry the main power supply for the amplifier board, VCC. The effective layer stack up is shown in Figure 21.




Figure 21: Backplane Layer Stack Up

Not shown in Figure 21 is an additional opaque black FR-4 layer in the middle of the stack up, intended to aid with light sealing. Since the backplane makes up the top of the tagger microscope dark box, it had to be designed to prevent light from passing through. While the copper planes inside the PCB are adequately opaque for most purposes, Figure 21 reveals that there are holes in many of the planes in the area surrounding vias. Additionally, as Figure 22 shows, even where vias connect with internal planes, there is still room for light to pass through.

In order to aid with soldering and plating processes, vias are not connected with internal planes around all 360° of their circumference. Instead, there are four small

traces that make the electrical connection, leaving some areas with no copper to block light. Figure 15 shows an example of a via heat relief. The center black circle is the via hole, through which light is of course able to pass. Less expected are the black outer arcs which also transmit light. Gray gridded areas in this diagram represent copper. A layer of black FR-4 in the middle of the PCB stack up helps to block light that manages to pass through these heat reliefs.




Figure 22: A Typical Via Heat Relief

Eurocard Connectors

The digital control board and the amplifier board connect to the backplane by means of a DIN 41612 Eurocard connector. The backplane has receptacles for both of these connectors, and all the necessary routing to pass signals from the control board to the amplifier board, and vice versa. Routing the connections between the control board Eurocard connector and the amplifier board Eurocard connector was the biggest challenge of the backplane design.

In order to satisfy the size constraints of the tagger, the backplane had a fairly strict maximum depth of about 2 inches. This size constraint didn't leave much room for the 32 SiPM bias voltage traces that had to be routed between the two Eurocard receptacles. In order to squeeze so many traces into such a small space, I had to route traces on both the front and back of the backplane. I used 5 mil traces to maintain safe spacing between traces and Eurocard receptacle pads, and also maintained the 15 mil recommended spacing between the traces themselves. (Recall that these bias voltage traces have the capability to carry up to 200V.)

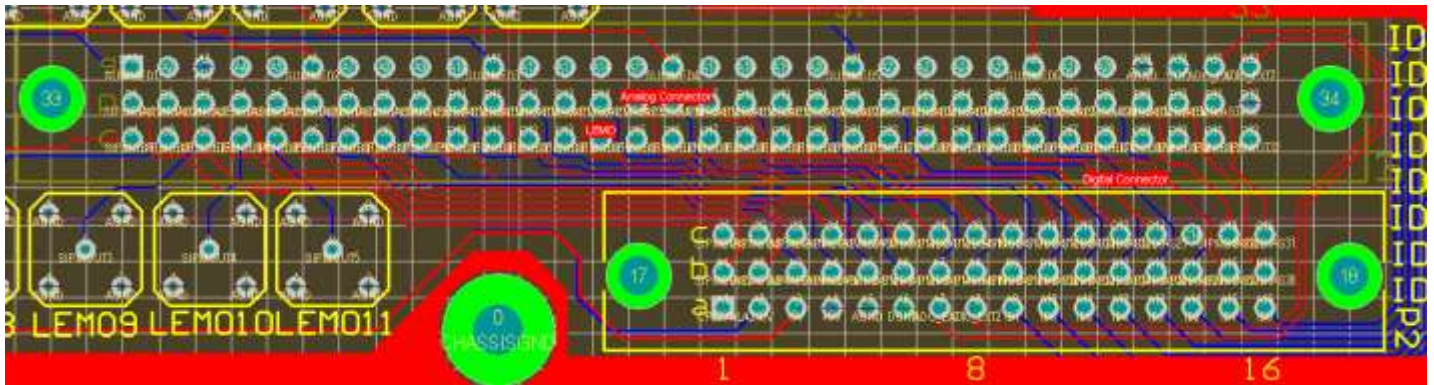


Figure 23: Close-Up of Eurocard Receptacle Routing

Figure 23 shows the intricate routing between the Eurocard receptacles. The large screw hole (labeled “0 CHASSIS”) to the left of the control board receptacle (P2) forced the control board receptacle to be offset quite a bit from the center of the amplifier board receptacle (P1, labeled in dark gold, indicating printing on the back side of the board). This made bias voltage trace routing very tight between the two boards.

Also worth noticing in Figure 23 are the two red traces that snake between screw hole 18 on P2 and screw hole 34 on P1. These traces connect the ADC on the control board to voltages on interest on the amplifier board.

Location ID Header

The location ID bus from the digital control board is brought to the backplane by pins A9 through A16 on the control board Eurocard connector. Traces routed on the back side of the backplane connect these pins to a 16-pin through-hole header, as shown in Figure 24.

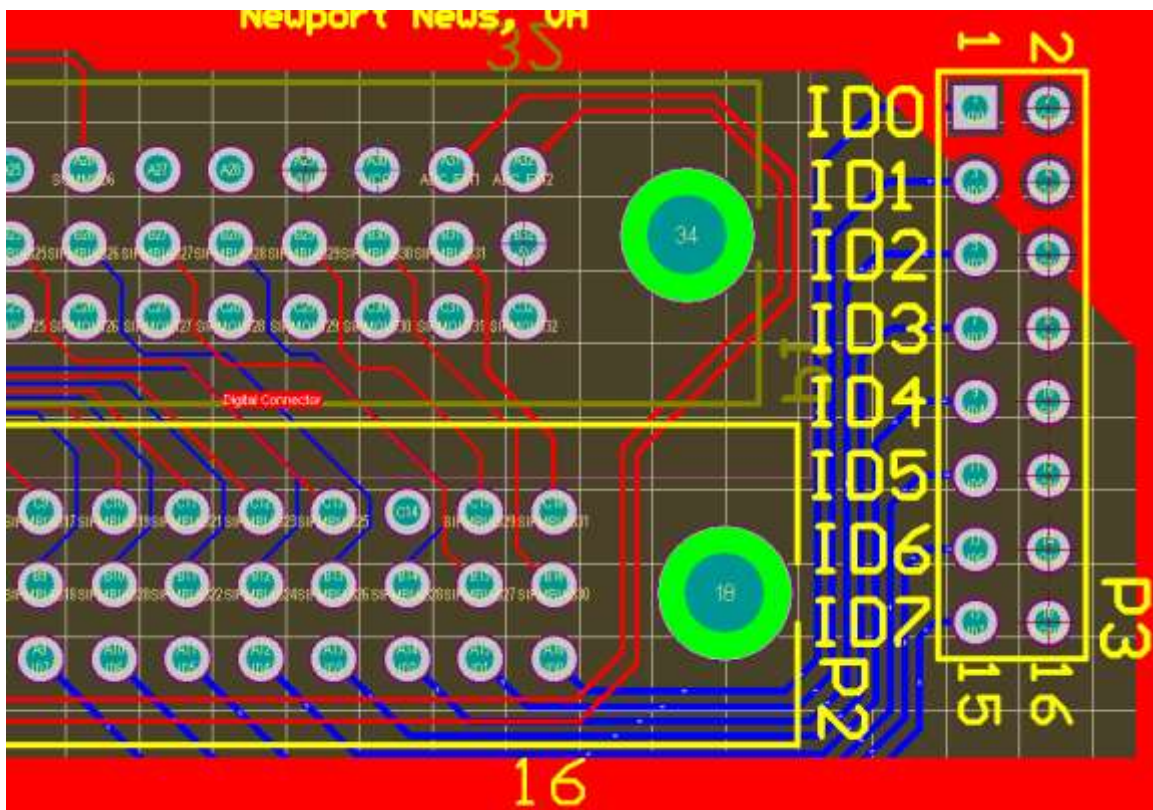


Figure 24: Close-Up of the Location ID Bus on the Backplane

To ensure that the backplane remains properly light-sealed, we used a through hole header to eliminate the need for hollow vias to tap into the internal ground plane.

Component through holes are filled with the component's own pins, as well as the solder used to attach the component. Vias, on the contrary, are typically left as open holes. While it would have been possible to manually plug vias with solder, it would just add an additional unnecessary step in the assembly process.

Ground Coupling Jumpers

To provide flexibility in grounding the control board and the amplifier board, two jumpers were added to the backplane. These jumpers allow for coupling of either the analog or digital grounds, or both, to the aluminum chassis of the tagger. The backplane is electrically connected to the chassis by means of the mounting screws that hold it onto the metal baseplate on the top of the tagger box. This chassis ground extends along much of the top surface of the backplane through a top layer hollow polygon (the large red area in Figure 20). Jumpers W1 and W2 allow easy connection of the chassis ground to either or both of the electrical grounds, as shown below. Since these jumpers may not be used, DGND and AGND are carried off the backplane separately.




Figure 25: Ground Coupling Jumpers

Power Connector and Fuses

The backplane is responsible for supplying power to digital board and the control board. To ensure safe and reliable power delivery, the backplane uses a 6-pin Molex KK® 6373 header for all power/ground connections. Though a slightly different design than the type of header used to supply power to 3.5" floppy disk drives, it is built around the same concept, with a friction ramp to secure the power cable in place. This type of connector has the advantage of being easily connected and disconnected without the need for extreme for or back and forth wiggling like the larger IDE hard drive style power connector. In addition, since the friction ramp is on only one side of the connector, it should be difficult or impossible to insert the cable backwards.




Figure 26: Molex KK® 6373 6-pin Power Connector

To prevent damage in the event of a short circuit, all of the incoming power lines are equipped with time-delay fuses. Selecting appropriate fuses was difficult because there is always a tradeoff between trip current, time to trip, and non-tripped resistance. A perfect fuse would have a low trip current, a short trip time, and a small non-tripped resistance, but in reality that is not always possible.

Table 3: Backplane Fuses

Power Supply Line	+5V, VCC	-5V	+210V
Fuse Selected	Tyco RUSBF120	Bourns, Inc. MF-MSMF014-2	Bourns, Inc. CMF-RL35-0
Mounting Type	Through hole	Surface mount	Through hole
Trip Current ²	2A	340mA	150mA
Hold Current ³	1.2A	140mA	75mA
Typ. Current Required ⁴	550mA	7mA max	3.2mA max
Non-Tripped Resistance	0.08Ω	6.5Ω	35Ω
Typ. Voltage Drop ⁵	0.044V	0.05V max	0.112V max

² Minimum current required for fuse to open the circuit.

³ Maximum current at which the fuse will stably leave circuit closed.

⁴ Current needed by digital control board and/or amplifier board.

⁵ Voltage drop from nominal values expected to be introduced by fuse in non-tripped state.

Time to Trip	0.5s at 8A	0.15s at 1.5A	0.15s at 3A
---------------------	------------	---------------	-------------

Table 3 gives an overview of the fuses selected for the backplane. These fuses were picked so that the time to trip during an overcurrent event of roughly one order of magnitude is 0.5 s or less. Where possible, through hole fuses were used to avoid the use of vias. The Bourns, Inc. MF-MSMF014-2 surface mount fuse does not require any vias because the -5V line is supplied to the digital control board entirely through a top layer trace. Though the +5V supply for the digital board is nominally the same voltage as VCC, the two were fused separately in case different voltages are needed. The 550mA current requirement represents the combined current requirements of digital control board and the amplifier board; each requires roughly half of this figure. Figure 27 shows a close-up view of the power connector and the backplane fuses.




Figure 27: Close-Up CAD Drawing of Backplane Showing the Power Supply Fuses. F1=-5V, F2=VCC, F3=+210V, F4=+5V

LEMO Connectors

The final important feature of the backplane is the LEMO connectors, which route signals from the amplifier board off the tagger electronics array to be digitized. In total there are eleven LEMO connectors, six of which route signals from the summed columns of scintillating fibers in each 5x5 fiber array, and five of which route signals from individual fibers in the first column of the 5x5 fiber array. The layout of the LEMO connectors on the backplane is shown in Figure 28. LEMO1 through LEMO6 are outputs of the amplifier board's summing circuits, and LEMO7 through LEMO11 are outputs of the first five amplifiers.




Figure 28: Close-Up of the LEMO Connectors

The specific part selected for the LEMO connectors is the LEMO EPA.250.00.NTN. Though expensive, this connector has many advantages over standard coaxial cable. In particular, cables that mate with the LEMO EPA.250.00.NTN are held in place by friction rather than any type of screw or twist mount system. This is important given the close spacing of the LEMO connectors on the backplane, as there is no room for fingers to screw or twist a connector into place. Figure 29 shows what the LEMO EPA.00.250.NTN looks like.




Figure 29: LEMO EPA.250.00.NTN Dimensions and Drawing

Amplifier Board

SiPMs on the amplifier board detect scintillation from the Bremsstrahlung electrons and output small electrical pulses. The amplifier board, appropriately, amplifies these signals so that they can be measured by an analog to digital converter located off of the tagger electronics array. The overall layout of the amplifier board is shown below.




Figure 30: CAD Drawing of the Amplifier Board

Hierarchy

As is evident from the appearance of the amplifier board in Figure 30, the board consists of several groups of repeated circuits. These circuits can be grouped into three hierarchical levels, as shown in Figure 31.

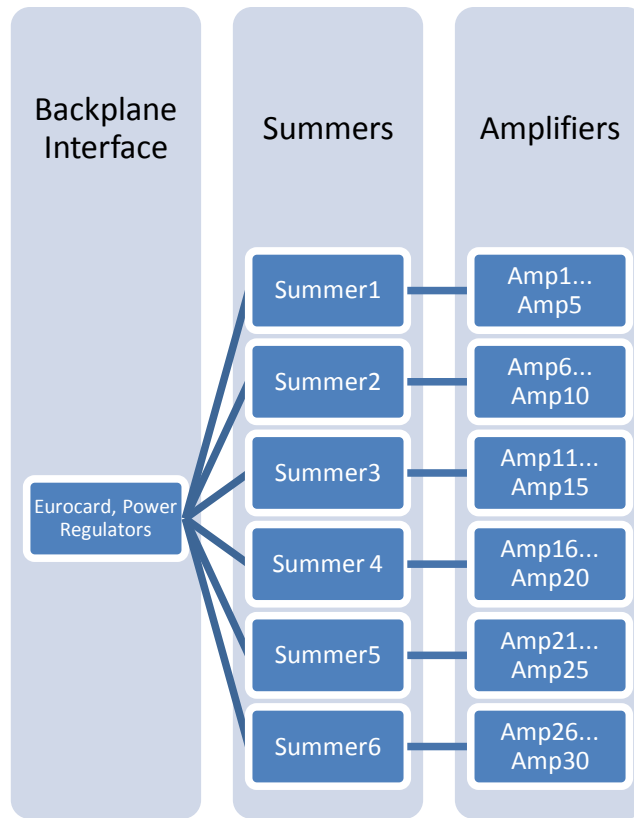


Figure 31: Hierarchy of the Amplifier Board Circuitry

Each of the 30 SiPMs is connected directly to a dedicated amplifier circuit. Each amplifier circuit outputs a signal into a summer. The summers add signals from adjacent groups of five amplifiers and output the signals to the backplane. Though not shown in Figure 31, each amplifier also outputs a signal directly to the backplane, bypassing the summing stage completely, for testing and calibration purposes.

PCB Layer Selection

Layers for the amplifier board were selected not only based on which power supply voltages were necessary in various areas of the board, but also on the requirement that all amplifier and summer outputs have a 50Ω impedance to the ground plane. To guarantee proper impedance and cross-talk free signals, the board was

designed with two ground planes – one on each side under the front and back outer signal layers. The two remaining internal layers were used for the amplifier power supply plane (VCC) and a split plane containing two transistor base voltages needed by the amplifiers. This layout is summarized in Figure 32.




Figure 32: Amplifier Board PCB Layer Stack-Up

Amplifiers

Except for which summing circuit they connect with, all amplifiers on the amplifier board are identical. They can be independently controlled, since each routes a separate bias voltage from the digital control board to a SiPM. They can also be independently read out, since each routes an output trace direct to the Eurocard connector.

The need for independent control and independent readout arises from the fact that no two SiPMs are identical. The optimum bias voltage for each SiPM varies even among SiPMs from the same batch. Independent bias voltages allow for sensitivity adjustments to compensate for variations in the SiPMs' performance. Independent readout of all thirty channels is intended for manual testing of each channel using an

The primary change made to the MATLAB model of the AMP_0604 circuit before implanting it into the schematics of the amplifier board was the addition of an extra transistor to support simultaneous output of the amplified signal both for independent readout and to a summing circuit. While it would seem that it should be possible to simply split the output trace into two traces, doing so allows amplifiers that share a common summing circuit to affect each other's individual outputs as well. (Regardless of the fact that the trace is split in two separate directions, both ends are electrically connected and can serve as a pathway for back feed from the summing circuit into the amplifier.) The additional transistor added to the amplifier circuit allows it to be safely connected to a summer without back feed of the summed signal into the amplifier circuitry.




Figure 34: Schematic of the Amplifier Circuit

In Figure 34, BIAS represents the bias voltage input from the digital control board. AMPREF represents one of the two transistor base voltages which is supplied by the split plane, and shared by all the amplifiers. SIPMOUT represents the signal sent to

the Eurocard connector for individual readout, and TOSUM represents the signal sent to the summing circuit.

Laying the circuit onto the PCB proved to be very challenging. In order to keep the amplifier board from being longer than the backplane, the amplifier circuit had to be placed in a space approximately the width of a single SiPM, 160 mil. Photonique's AMP_0604 was more than twice that width, and it had one less transistor than the design to be used on the amplifier board. In order to accomplish such a dramatic size reduction, I used 0201 (0603 metric) size components in the design. As indicated by their name, these components are 0.6 mm x 0.3 mm in size, and are barely visible without the aid of a magnifying glass. These tiny components, along with clever routing, led to an amplifier circuit design with overall dimensions 1.13 in x 0.18 in excluding the SiPM itself.

While the 0.18 in width of the summing circuit exceeds the 160 mil SiPM width, it does not pose a problem because certain components of the amplifier circuit are shared between adjacent channels. In particular, ground traces that run the full height of the amplifier circuitry for crosstalk resistance are shared between adjacent amplifiers. To see how this is possible, take a look at Figure 35 below. The yellow traces between amplifier channels and in the component key on the left are silkscreen which is printed on the PCB for visual reference. Right behind the yellow silkscreen dividers between AMP1, AMP2, and AMP3, there is a top layer (red) ground trace that jumps back and forth between vias and components in adjacent channels.




Figure 35: Close-Up of Amplifier Circuitry and Component Key

The blue traces in Figure 35 are bottom layer traces, and are from left to right across any amplifier channel, the bias voltage from the control board, the output to the summing circuit, and the output for independent readout.

Summing Circuit

Each of the six summing circuits on the amplifier board is designed to accept inputs from five amplifier channels. In Figure 36, the point at which amplifier signals are injected into the summing circuit is labeled "TOSUM." The signals go through a series of transistor stages, each boosting the summed signal slightly until finally they are output at the point "SUMMEDOUT." The "SUMREF" input, like the "AMPREF" input in the amplifier circuit, is a transistor base voltage provided by the split power plane, which can be adjusted as needed during the testing phase of production and fixed once performance is optimized. "GAINMODE" set by the DAC on the control board, adjusts the gain of the summing circuit to allow for adjustable sensitivity to photons. Transistor M1 is a field effect transistor which adjusts the amount of current allowed to flow from VCC through R10 and C8, thus adjusting how much amplification transistor Q4 can provide.

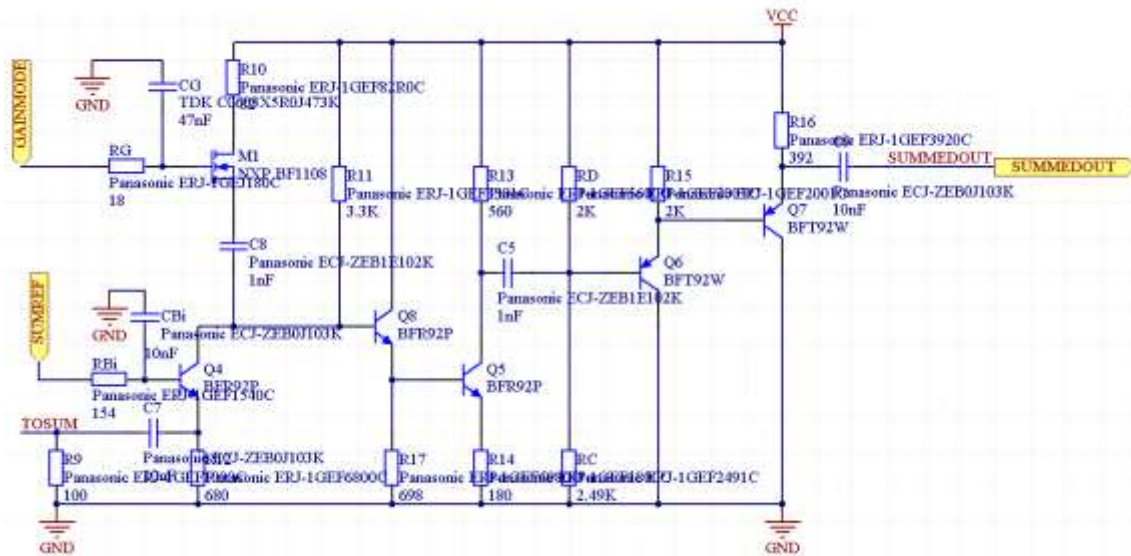


Figure 36: Schematic of the Amplifier Board's Summing Circuit

Figure 37 below shows a close-up of the layout of the summing circuit on the amplifier board. Though it still looks quite complicated, the layout of the summing circuit was dramatically easier than the layout of the amplifier, since the summing circuit had a more relaxed size constraint. In Figure 37, amplifier signals are input into the five vias just above the AMP1 through AMP5 silkscreen labels. The via labeled BASE1 (between the CBi and RBi labels), pulls the transistor SUMREF transistor base voltage from the gold plane layer, which can be seen underneath all the circuitry. The split in this plane, so that it can simultaneously carry the AMPREF and SUMREF signals, can be seen on the left side of the summing circuit.




Figure 37: Close-up CAD Drawing of Summing Circuit SUMMER1 on the Amplifier Board

V. Debugging the Tagger Electronics

Once all of the circuitry had been designed using Altium Designer, manufactured, and assembled, there came the task of testing and debugging. This stage helped reveal critical problems with the electronics so that they can be corrected for the final hardware revision.

Digital Control Board

Since the digital control board was produced before the other boards were even designed, it was the first board to go through the debugging process. The first problem we found with the digital control board was that it was not supplying the required +1.2V for the FPGA's internal logic. Instead, the +1.2V power plane was found to be at the same potential as ground. It turned out that the +1.2V voltage regulator simply wasn't connected to its +5V power supply. The pin was left floating in the schematics. To fix this, we placed a dab of solder between the regulator's power supply pin and its adjacent chip enable pin, both of which should be at +5V permanently.

Figure 38: Photograph of the +1.2V Voltage Regulator on the Digital Control Board, Showing the Solder Bridge Used to Connect the Floating Power Input Pin

The next problem we discovered was a miswiring of the SPI bus that connects the FPGA, ADC, and temperature sensor. The pin on the FPGA that was designated as SPI master-out-slave-in was connected to the SPI input of the ADC as it should have been, but also connected to the SPI output of the temperature sensor. The SPI bus

standard specifies that all slave devices should share common input and output traces which connect them back to the master device, in this case, the FPGA. High impedance logic and independent chip select lines are used to prevent slave devices from fighting over the signal lines. We corrected the miswiring by cutting the temperature sensor's DOUT trace, and soldering a small wire to connect it to the same trace as the ADC's output.

Figure 39: Photograph of the Temperature Sensor and ADC on the Digital Control Board After Rewiring the SPI Bus

A third mistake, we found while debugging the control board was that I had selected the wrong type of clock for our FPGA/Ethernet controller combo. Initially, I selected a crystal oscillator to serve as a clock for the Ethernet controller. Since the Ethernet controller had a built in inverting driver, it seemed like a great, inexpensive, and precise way of clocking the device. Later, however, we decided that to conserve on space, it would be convenient to share a single clock between both the FPGA and the Ethernet controller. Forgetting that I had selected a crystal oscillator, I connected the oscillator to both the FPGA and the Ethernet controller simultaneously. This caused a number of problems. First, the FPGA is configured to use 3.3V CMOS logic, which is not what the crystal oscillator's sine wave output provided. Additionally, the capacitance of the FPGA's clock input put was enough to distort the crystal's signal, preventing even the Ethernet controller from operating properly.

The solution to the problem with the crystal oscillator was to replace the oscillator with a CMOS clock. As mentioned in "Wiring the Ethernet Controller," the

CP2201 Ethernet controller can accept a CMOS clock input in place of a crystal oscillator. Though somewhat more expensive than a crystal, the CMOS clock output can be used by the FPGA as well as the Ethernet controller, satisfying the shared clock design requirement. To implement this change, I connected the output of the CMOS clock to the FPGA, and then ran a trace from another FPGA pin to the Ethernet controller. The FPGA is programmed to reproduce the clock signal on this second clock pin to clock the Ethernet controller.




Figure 40: Photograph of the CMOS Clock (Immediately Below Label C17) Used to Replace the Crystal Oscillator on the Digital Control Board

After this change, the FPGA was fully operational, and ready for debugging of its own internal configuration. However, before debugging could begin, we ran into

another problem. The header that the FPGA's programming equipment attached to was not the correct header called for in the equipment's documentation. Though a 100 mil pitch header is much more common, the Xilinx Platform II USB FPGA configuration cable requires a 1 mm pitch header. Since the through holes in the board for this header cannot be moved, we were forced to stick with the incorrect 100 mil header for the prototype. We used an alternative "flying lead" adapter from Xilinx to connect signals from the Platform II cable to the individual pins on the 100 mil header. The design was updated to use the correct 1 mm header for future revisions.

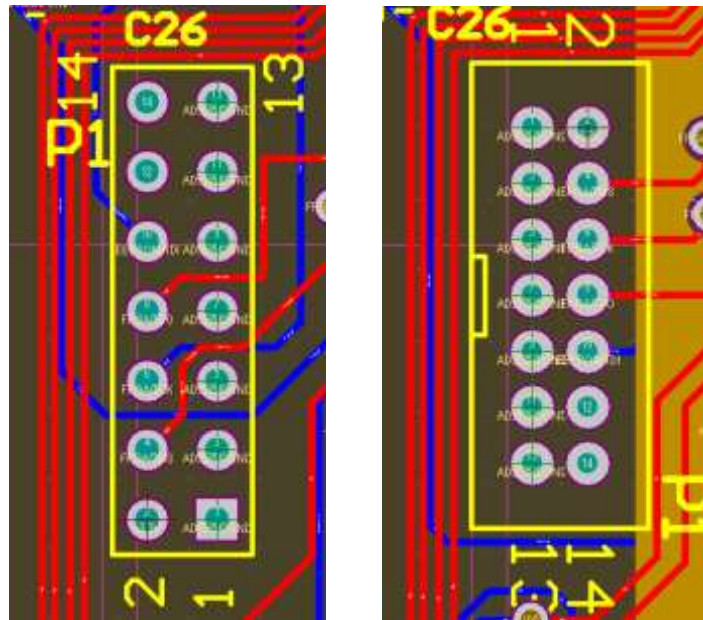


Figure 41: Close-Up CAD Drawings of the Footprint of the FPGA's JTAG Configuration Header and Surrounding Traces Before (Left) and After (Right) the Incorrect 100 mil Header was Replaced

Figure 42: Photograph Showing How the Flying Lead Adapter is Used to Connect the Xilinx Platform II USB FPGA Configuration Cable to the 100 mil Header on the Control Board

With the flying lead adapter in place, we were able to debug the FPGA's internal configuration, and reupload its program as needed from the Xilinx programming software running on a PC. Most of the work on debugging the FPGA involved getting the FPGA and the Ethernet controller to talk to each other correctly. All the details of this process are beyond the scope of this paper, but the process basically consisted of programming in VHDL code, downloading the program to the EEPROM/FPGA, and then testing using an oscilloscope to see if changes in the code were successful. To aid in monitoring particular internal signals in the FPGA, some were assigned to unused output pins and routed to a header on the edge of the board. This header was attached manually, only for the purpose of the debugging the FPGA, and is not part of the actual board design.

Figure 43: Photograph of the Diagnostic FPGA Signal Output Header Attached to One of the Prototype Control Boards

After several weeks of work, the FPGA was up and running, able to communicate with the Ethernet controller, the SPI bus, and all other components on the board without any problems. The next difficulty arose, however, when trying to send packets over the Ethernet wire. Though the FPGA was able to communicate with the Ethernet controller properly, the Ethernet controller was not able to transmit the packets into the wire. It turned out that I had ordered the wrong Ethernet jack for the board. The jack I had ordered, the Pulse J0011D21 had the same footprint as the Pulse J0012D21 that the bill of materials called for, but unfortunately, did not have the same pinout. Ethernet packets were being dropped before they ever left the board because they weren't being routed into the correct transmit/receive wires in the Ethernet cable. Fortunately, this

was easy to fix. We ordered the correct part, attached it in place of the wrong one, and the problem was solved.

The final problems we encountered on the digital board centered around the DAC. First, we found that diode D2, whose purpose it is to ensure proper power supply sequencing, was reversed. Though correct in the schematics, the footprint had the pin numbers backwards. This was causing a short between the -5V power supply and the analog ground. The built in ammeter on our -5V power supply was what first alerted us to this problem.

The last and more perplexing problem, is that on two of the three identical prototype control boards, the DAC was completely nonfunctional. As of March, 2010, we have been unable to determine the cause of the DAC failure. We are confident that it was not caused by the reversed diode, because the diode was corrected before the ever applying power to one of the two failed boards. We are also fairly confident that it wasn't caused by problems in the assembly process, because the assembly company's records show that they following all of the manufacturer's recommendations for soldering the DAC to the PCB. The failure also cannot be due to the design of the PCB, because one of the three boards works perfectly. This is a problem that we are still investigating, which we will be sure to solve before the final hardware revision.

Amplifier Board

Compared to the control board, debugging the amplifier board was a fairly simple process. Thanks to the MATLAB model of the amplifier and summing circuits,

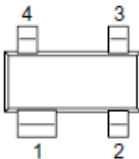
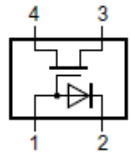
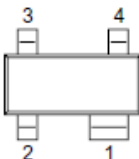
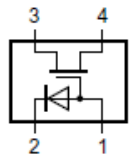
checking for errors simply involved probing different component pins to see if their DC levels matched those predicted by the model.

The first problem we noticed was that the two transistor base voltages “AMPREF” and “SUMREF” were both at ground, rather than their nominal values. We traced this problem back to the voltage regulating circuitry on the amplifier board, where we discovered the two identical voltage regulators used to produce the transistor base voltages had a footprint mismatch. On both of the regulators, I accidentally switched the output and the ground pins. Since the regulators are shunt-type regulators, there was no harm done by this, and a quick resoldering job fixed the issue on the prototype board. I updated the component footprint in Altium Designer for the final hardware revision.

Figure 44: Photograph of the Power Regulating Circuitry on the Amplifier Board, Showing the Resoldering of the Two Voltage Regulators to Resolve a Footprint Mismatch

Another problem with the amplifier board which we actually knew about from the time we sent the board out for assembly, is that one of the components we selected to use on the board was not available in the footprint that we selected. The NXP BF1108 MOSFET, which is responsible for controlling the gain of the summing circuit, is only available in North America as the BF1108R. The BF1108R is identical to the BF1108, except its footprint is a mirror image of the footprint that we selected while designing the board layout. To resolve this, I examined the data sheet for the BF1108 and determined how we could make the BF1108R function on the BF1108’s footprint. Fortunately, the BF1108 and BF1108R both have four leads, meaning that it’s possible to

rotate the component 180° and attach it by hand. While this alone still left a mismatch of pins, the particular design of the BF1108/BF1108R made it possible for this orientation to work. Since the BF1108/BF1108R is an FET transistor, it has a drain and a source which are interchangeable. The 180° rotation swapped the drain and the source of the FET, with no ill effects. The other two pins of the transistor were the base voltage and the cathode of an internal diode which was not used in our design. The rotation of the FET put the cathode pin where the base voltage pin should have been, and left the base voltage pin on a floating pad. Since we didn't plan to use the internal diode anyway, we simply shorted these two pins with a solder bridge, resolving the footprint mismatch.

Pin	Description	Simplified outline	Graphic symbol
BF1108 (SOT143B)			
1	FET gate; diode anode		 001aa1042
2	diode cathode		
3	source [1]		
4	drain [1]		
BF1108R (SOT143R)			
1	FET gate; diode anode		 001aa1043
2	diode cathode		
3	source [1]		
4	drain [1]		

[1] Drain and source are interchangeable.

Figure 45: Drawings and Pin Assignments of the BF1108 and BF1108R MOSFETs Taken Directly from the Manufacturer's Datasheet, Showing their Mirror Image Pinouts

Figure 46: Photograph of the Summing Circuit on the Amplifier Board, Showing the BF1108R Attached to the BF1108 Footprint

Once the BF1108R MOSTFETs were in place, debugging of the amplifier board was complete. Next, debugging began on the backplane.

Backplane

Since the backplane does little other than connect the control board and the amplifier board together, and route power connections from off-board power supplies, debugging it was a fairly simple process. We used a multimeter to test that power was being delivered to the correct pins on the control board and amplifier board Eurocard receptacles, and then plugged in the boards to see if they worked. We found no major problems with the backplane circuitry, except for when we attempted to hand solder the LEMO connectors onto the board. Unfortunately, we discovered that I had made the through holes for the LEMO connectors too small. The pins on the LEMO connectors simply would not fit through the holes that were in place on the board. The reason for this turned out to be that I had overlooked the thickness of the through hole plating, when designing the LEMO connector footprint.

Two possible solutions were considered to this problem. First, I considered the consequences of drilling the existing holes to make them large enough to accept the LEMO pins. Unfortunately, this would remove the plating on the inside of the through holes. This would not stop the outgoing signals from passing into the LEMO connector, because the end of the signal pin could be soldered to the metal pad carrying the signal on the back side of the backplane. However, this would stop the four ground pins on the LEMO connectors from making contact with the board's internal analog ground plane.

Having unshielded signals could lead to interference from outside electrical sources, or even cause signal reflection due to a mismatch in impedance. This is certainly not something we wanted to happen with these analog signals from the summing circuit.

To avoid these problems, we proposed simply filing the pins on the LEMO connectors to make them small enough to fit through the existing holes. Though a delicate process, it was successful, and we were able to attach all of the LEMO connectors to the backplane.




Figure 47: Photograph of the Completed Backplane, with LEMO Connectors Attached