Prototype Scintillating Fiber Tagger Microscope Design and Construction

I. Senderovich, C.R. Nettleton and R.T. Jones

June 20, 2008

Abstract

The University of Connecticut nuclear physics group has produced a design for a scintillating fiber microscope to be used in the GlueX tagging spectrometer. The design calls for the readout of individual scintillating fibers using silicon photomultipliers (SiPMs). These relatively new solid state devices have a different sensitivity to wavelength and operating conditions than do conventional photomultiplier tubes and require careful study. Three SiPM devices were obtained from commercial manufacturers in Europe and Japan, and characterized in terms of their gain, timing, dark rate, and quantum efficiency. The most suitable unit was further tested for its sensitivity to variations in bias voltage and ambient temperature. The group has developed a design for the readout electronics that provides individual remote control over the bias voltage on each SiPM and continuous monitoring of temperature and supply voltages. A small number of prototype boards are being constructed to be used to read out a prototype microscope consisting of 30 fibers and readout channels, 5% of the final channel count. Fiber polishing and gluing techniques have been developed for construction of the prototype. The prototype will enable all of the critical components of the detector design and fabrication techniques to be tested in a real electron beam before the full-scale detector construction begins.

1 Overview

Photons with energies up to 12 GeV will be produced in Hall D at Jefferson Lab via coherent bremsstrahlung (CB) of electrons in a diamond crystal. This technique produces a broad radiation spectrum with pronounced peaks at discrete energies that correspond to coherent scattering of the radiating electron from sets of planes in the diamond crystal. The photons in the coherent peaks are linearly polarized, with polarization selected by the orientation of the crystal. Fig. 1a shows the bremsstrahlung beam spectrum before (upper curve) and after (lower curve) collimation. As can be seen, the peaks are strongly enhanced by the collimation.

The energy of a beam photons is tagged by measuring the post-bremsstrahlung electron energy in a magnetic spectrometer (tagger). Since the tagging takes place prior to the collimator, the tagging counters must deal with the full beam rate (upper curves in Fig. 1. These rates are high enough that they place limits on the minimum segmentation required in the tagger focal plane that are more stringent than the requirement of tagger energy resolution. As can be seen in Fig. 1b, restriction of the rates to 3 MHz per channel, a conservative upper limit for scintillation counters,

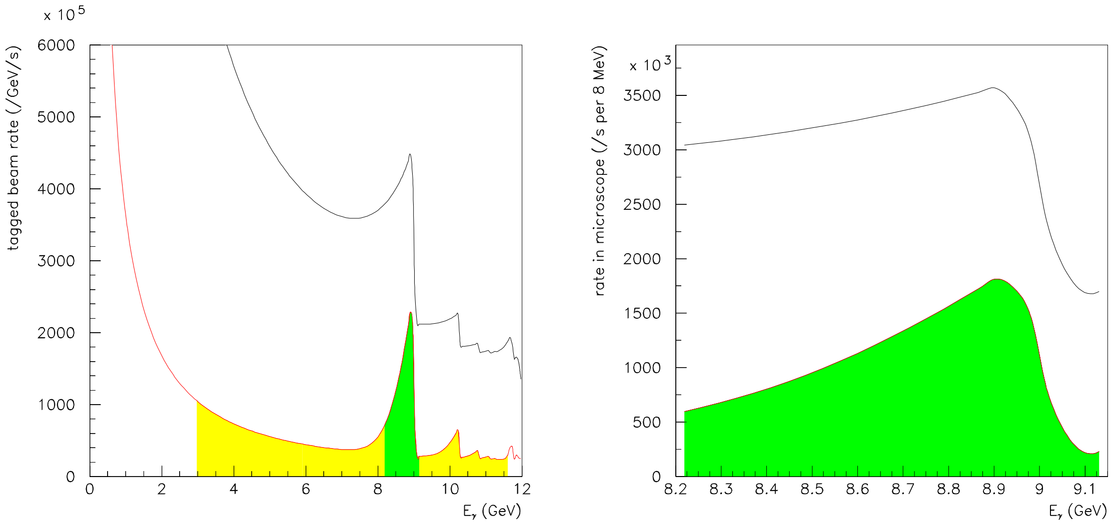

Figure 1: Coherent bremsstrahlung beam spectrum from 12 GeV electrons at the current foreseen for full-intensity running with the Hall D tagged beam. In the left panel, the upper curve shows the full beam spectrum seen in the tagger, while the lower curve shows the part of the beam that passes the collimator. The broad shaded zone shows the region tagged by the broad-band tagger hodoscope. The microscope subtends the narrower shaded zone. In the right panel the same information is shown for the narrow energy zone subtended by the microscope, rescaled to show the total rate at full tagger intensity per microscope readout channel of width 8 MeV.

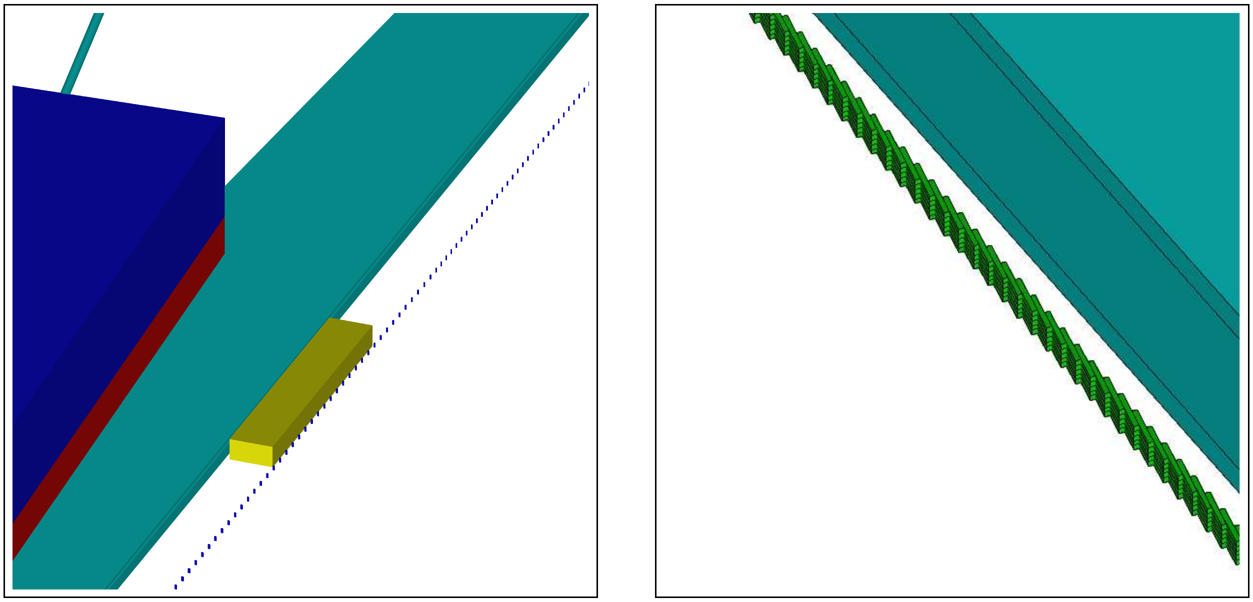

Figure 2: Layout of the tagger focal plane (left panel) showing the box containing the microscope counters, and the array of individual counters comprising the broad-band hodoscope. The right panel shows a close-up of the microscope fiber array with the box removed.

requires that each channel subtend no more than 8 MeV in the region of the primary coherent peak.

It is not necessary to instrument the entire focal plane at this fine level of segmentation because the polarized component of the beam that is of paramount interest to the GlueX experiment is restricted to the region of the primary peak extending from 8.4 to 9.0 GeV in photon beam energy. Unpolarized tagged photons outside this region are also of interest in Hall D, but such running does not require the narrow collimation used for polarized beam. Removing the narrow collimator results in significantly improved tagging conditions, and should be done whenever linear beam polarization is not required. Tagging with the broad-band hodoscope will still be possible during running with the collimator, albeit under highly non-ideal conditions with a tagging efficiency of only 15-20%.

For these reasons, the tagger focal plane has been divided into two independent sets of counters: the microscope for use with polarized beam running with a collimator, and the broad-band hodoscope for crystal setup and unpolarized beam running. The two arrays are shown in Fig. 2. The microscope is composed of a rectangular 120 x 5 array of square square 2×2 mm^2 scintillating fibers glued in a staggered arrangement as shown in Fig. 2b. The back end of each scintillating fiber is coupled to a clear acrylic fiber light guide of the same dimensions, which carries the scintillation light to a SiPM photodetector. Each scintillating fiber and light guide is optically isolated from its neighbors by two layers of cladding and an outer opaque coat.

The microscope tagging counters must be capable of running at rates up to 5 MHz without significant pile-up or loss of gain. They must have a reasonably high efficiency to avoid extra background in the experiment caused by untagged beam photons, and have sufficient time resolution to cleanly distinguish between adjacent pulses in the electron beam time structure. They must have sufficiently low background rates that extra tags caused by accidental coincidences with tagger background do not appreciably dilute the experimental signal. Detailed specifications corresponding to these requirements are the subject of the next section.

Phototubes are the conventional choice for tagging counter readout. However, the compact form-factor of the scintillating fibers, combined with a number of advantageous features of SiPMs has led to their adoption for the readout of the Hall D tagger microscope. The following are the principal advantages of SiPMs over the more commonly used photomultiplier tubes (PMTs) in this application.

1. No high voltage is required (bias voltages of order 50 V are sufficient).
2. Signal rise times are about a factor two faster than ordinary PMTs.
3. No sensitivity to stray magnetic fields.

In contrast to standard solid-state photon detectors, the SiPM has most of the desirable features of PMTs: gains of order 10^6, rise times on the order of a few ns, and high signal/background in photon-counting applications. However SiPMs have only recently become commercially available and their use for applications such as ours has been limited. This report reviews the careful characterization of some candidate devices undertaken by the Connecticut group. Additionally, matters of scintillator selection, coupling to sensor, its bias voltage control and feedback from readout circuitry are discussed.

2 Requirements

The following list summarizes the tagged photon beam requirements for Hall D which have a direct impact on the design of the tagging microscope. A photon beam meeting these specifications is needed in Hall D in order to ensure that the GlueX experiment can reach its physics goals.

1. complete coverage within coherent peak 8.4 - 9.0 GeV
2. tagging time resolution better than 200 ps r.m.s.
3. tagger segmentation finer than 10 MeV across coherent peak
4. capable of operating up to 250 MHz total rate within coherent peak
5. capable of tagging efficiently at beam currents down to threshold of 10^{-3} of nominal intensity for absolute normalization measurements
6. average tagging efficiency at least 70% with collimation

One of the principal requirements for this instrument is a time resolution less than 200 ps. This is necessary in order to ensure the clear identification of an electron with a particular electron beam bunch. The bunch period is 2 ns. With decay time (τ) of 2.7 ns for the BCF-20 fast-green scintillator, this requirement places a lower limit on the number of detected photons (n) produced in a tagging counter by a single passing electron. A crude estimate for attainable time resolution from a scintillator pulse is $\tau/\sqrt{n}$, which suggests at least 180 photons. This figure is attainable if the scintillating fiber is oriented so that a tagging electron travels axially down the scintillating fiber. If the path length of the electron in scintillator is 2 cm then it deposits 4 MeV of energy. This results in 1600 scintillation photons within the forward capture cone of a double-clad fiber. Assuming that 80% of these are delivered to the SiPM active surface and a conservative SiPM photo-detection efficiency (PDE) of 15% leads to an estimate of 190 SiPM pixel count per tagging event. It may be possible to nearly double this yield by applying a reflective coating to the upstream end of the fiber and capturing the light that goes the wrong way up the fiber, but this is not needed in order to meet the time resolution requirement for the microscope.

The scintillating fibers shown in Fig. 2b are 2 cm long and oriented with their axis parallel to the nominal rays of tagging electrons exiting the spectrometer in the region of the coherent peak. Beam divergence and multiple scattering lead to an angular spread in the incidence angle of electrons entering the fibers, but this spread is small compared to the angular acceptance of a $2 \times 2 \times 20$ mm^3 fiber, as shown in Fig. 3. The 2.86° half-angle acceptance of the fiber is defined as the incidence angle of a ray that enters one end of the fiber at the center and exits through the downstream end of the fiber, regardless of the azimuth. From the figure it is seen that the alignment of the fibers with electron trajectories remains well within the acceptance across the entire detector if the fibers are simply glued parallel to each other.

Examination of Fig. 1b suggests that the last requirement in the above list cannot be attained with the chosen collimator; even at the peak the tagging efficiency does not exceed 55%. However, the vertical subdivision of the microscope fiber array can be exploited to overcome this limitation. If the fiber array is correctly centered about the plane of the electron beam in the tagging spectrometer, then momentum conservation at the radiator combined with the focusing properties of the spectrometer guarantees that only the central row of fibers in the microscope sees the electrons that tag photons which pass the collimator. The microscope is designed so that each horizontal row of fibers can be turned on and off in the readout, so that maximum tagging efficiency can be obtained by turning on only the central row of fibers, once the vertical alignment of the counters with the plane of the electron beam has been properly set up.

Monte Carlo simulations have shown that energy from a single electron is divided between two adjacent fibers 15% of the time, and that a single-fiber detection threshold at 1.5 MeV leads to essentially full detection efficiency. For SiPM readout, this corresponds to a threshold of 80 SiPM pixels. This in turn implies a requirement that the rate of spontaneous pulses ("dark rate") containing 80 or more pixels be small compared with the minimum nominal rate in the tagging counters of 1 KHz per channel, corresponding to running at 1 nA electron current.

It will be shown below that these requirements do not push any of the limits on SiPM performance. Devices currently available possess very low dark rates for this application and a have more than sufficient PDE given the expected light yield from the fibers.

Secondary performance parameters like pixel cross-talk and after-pulsing are also important considerations. The former may exacerbate the effect of the dark rate by producing a larger number of pixels per spontaneous discharge, increasing the probability of "pileup" beyond the detection threshold. The latter stretches the pulse, creating problems for leading-edge timing in the TDC.

Although tagging counters function as digital devices, their dynamic range is nevertheless important. There must be enough pixels for the initial set-up and gain equalization. If they begin to saturate before they have detected the number of photons required for the the 200 ps time resolu-

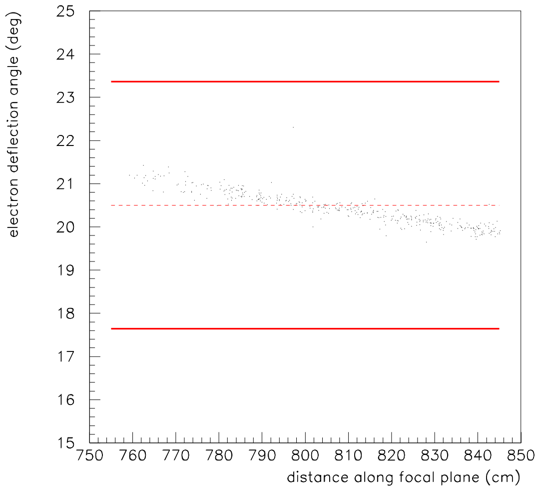

Figure 3: Dispersion of the bend angle of the tagging electrons in the tagger dipole field as a function of the impact location z along the focal plane in the region covered by the microscope. The electron energy is approximately linear in z. The points are an unbiased Monte Carlo simulation of all electrons that exit the tagger vacuum box in the region subtended by the microscope. The simulation incorporates the spot size and angular divergence of the electron beam, and multiple scattering in the diamond radiator, in the vacuum exit window, and the plastic scintillator. The horizontal lines show the acceptance half-angle of the fibers aligned so that the nominal electron ray at the center of the microscope is parallel to the fiber axis. All of the fibers are assumed to be parallel.

tion, the time resolution will be degraded. Pixel occupation, given the average recovery rate, and reduced count in signal due to saturation must all be considered.

Lastly, stability of the device with respect to the above parameters must be understood. It is well-known that the gain and dark rate of SiPMs are sensitive to variations in environmental parameters, including bias voltage and temperature drifts. A control system is required to monitor readout board voltages and temperature, and adjust the SiPM bias voltage accordingly to compensate for their variations. Also inherent performance characteristics may vary considerably from one device to the next. The control system must be capable of adjusting the bias voltage on each individual SiPM detector such that all performance parameters are kept within the necessary bounds. Modest cooling of SiPM devices may also prove necessary.

This design for the tagger microscope and associated readout electronics entails a number of novel features which should be tested in a prototype before the full-scale detector is constructed. A small-scale prototype microscope is under construction by the Connecticut group that incorporates all aspects of the final detector, from the scintillating fibers to the readout electronics. The following sections describe the studies that have been performed so far, in working out the details of the prototype design and fabrication techniques.

3 SiPM Characterization

3.1 Test Stand

Characterizing SiPMs requires a testing environment in which fast light pulses of controlled amplitude can be generated in low background. A test stand has been constructed for this purpose. It is essentially a hermetically sealed chamber with a pulsed LED on one end and a mount for a photodetector on the other. These are separated by a distance of order 1 m to reduce the solid angle subtended by the LED at the detector surface in order to test the detector at low light intensities. A series of filters with varying transmission factors were also prepared so that the light intensity could be varied without changing the LED output pulse shape.

The pulser circuit shown in Fig. 4 was fabricated for the test stand. The output pulse magnitude is controlled by the amplitude of the square wave from the function generator that is supplied as input, with a threshold near 4 V. The quiescent output state of the OpAmp is -4 V which is below threshold for the LED. When a rising edge is received from the function generator, the pulser circuit output swings toward +4 V at its maximum slew rate of 2 V/ns, saturating at the point where the LED draws the maximum output current of 50 mA. If the injected square-wave amplitude is close to the minimum of 1 V, a small output pulse is produced that never reaches saturation and has a width of about 2 ns FWHM. Increasing the square wave amplitude initially increases the output pulse amplitude until it reaches the level where saturation occurs, which depends on the LED threshold. Increasing the square wave amplitude further produces output pulses of increased duration with a roughly constant brightness during the pulse. At the maximum square wave amplitude of 8 V, the output pulse duration was measured to be 6 ns. This design can drive LEDs with thresholds in the range 2-6 V, which includes most commercial devices. Photographs of the circuit and the LED excited by a high-repetition rate pulse train (in order to make its output visible) are shown at the right in Fig. 5.

A Wavetek 166 50 MHz function generator and power supplies of appropriate ratings were procured for these measurements. A data acquisition system consisting of a 2 GHz digital sampling

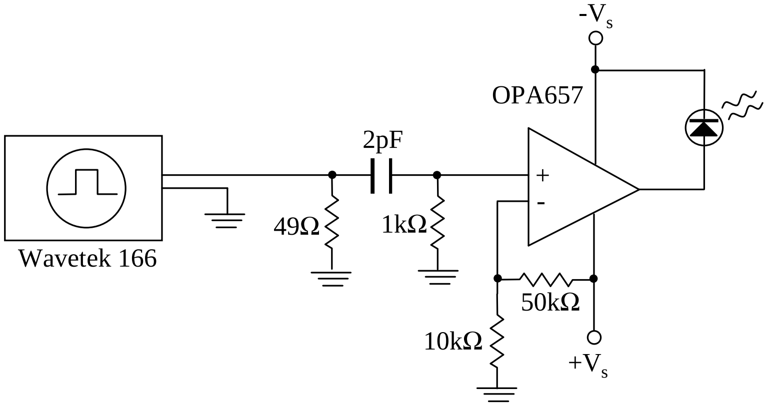

Figure 4: Diagram of the LED pulser circuit. The input square wave from the function generator is differentiated and used to drive the high-frequency operational amplifier OPA657. The value of V_s is 5 V.

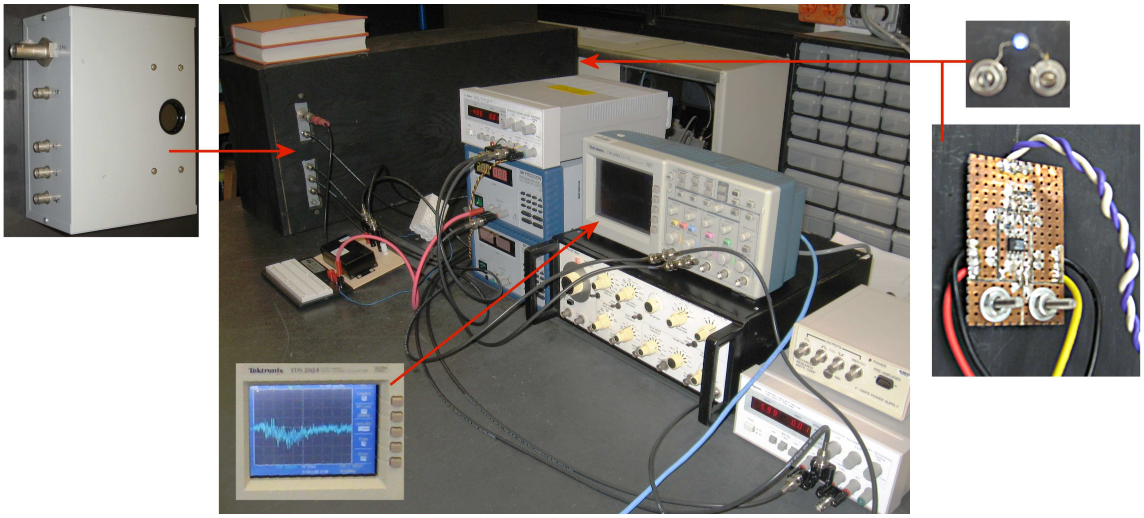

Figure 5: Test stand during measurement of the HPD response. The black box in the back of the central photograph is the dark chamber used to test the photodetectors. The HPD unit mounted on the wall of its electronics enclosure is shown separately on the left. The LED pulser is shown on the right with its pulse-shaping amplifier circuit.

oscilloscope with 200 MHz bandwidth (Tektronix TDS-2024) coupled to a PC was used to capture and store large numbers of waveforms from the detector. Analysis of these waveforms allowed the separation of signal and noise contributions to the detector response. Counting rates were also measured using analysis of captured waveforms.

The LED output pulse intensity was calibrated using a hybrid photodiode (HPD) as a photon counting standard. The HPD signal was amplified using an OpAmp circuit whose gain and impulse response function could be computed. The efficiency and intrinsic gain parameters of the HPD are specified by the manufacturer. Using this information, both the integral and the functional form of LED output pulse were extracted from the measured HPD output waveform. Dividing the measured photon count from the HPD by its sensitive area and quantum efficiency gives the absolute radiant intensity of the test stand LED pulser as a function of the amplitude from the function generator.

In addition to the 80 V required for the diode bias, the HPD also needs an external 12 KV high voltage supply. Electronics for high voltage, bias and signal amplification for the HPD have been incorporated into the test stand. A photograph of the test stand with associated electronics is shown in Fig. 5. A typical output waveform from the HPD in response to a small input pulse is shown in the figure.

The test stand has been fitted with a temperature-controlled photon detector mount. This "cold plate" has been fitted into the wall of the dark chamber. The cold plate is coupled to a heat exchanger on the outside of the box through a Peltier junction. The Peltier junction is regulated by a *programmable integral derivative* (PID) controller that senses the temperature using a thermistor attached directly to the detector mount. The controller is capable of regulating the temperature of the mount to within 0.1°C of the set point, which is variable over the range from $-20°$ to $+40°$C. The mount consists of an aluminum finger coupled to the SiPM with a drop of thermal paste. This setup will allow the SiPM temperature to be varied without affecting the operating temperature of the LED.

The set point of the temperature controller is adjusted by means of a dc voltage that is adjusted by means of a potentiometer . Additional potentiometers specify the PID control parameters, including integration time and sensitivity. The controller adjusts the temperature of the cold plate by sensing the resistance of a feedback thermistor attached to the plate and driving current in one direction or the other through the Peltier junction until the thermistor resistance matches a value that it looks up in a built-in lookup table for the given set point. The controller has been recalibrated using a thermistor that was itself calibrated using a standard laboratory thermometer.

3.2 Light Sources

Three LEDs with different emission spectra were originally procured and calibrated with the HPD: blue (QT Optoelectronics MV5B60), yellow (Fairchild MV8304) and red (Fairchild MV8104).

In order to understand the response time of the SiPM and amplifier system, it was necessary to first determine the duration and shape of the LED output pulse. It was assumed initially that the LED output intensity was proportional to the instantaneous current passing through the diode, but that turned out to be false in the case of the blue diode. Pulse shapes of all three LEDs were extracted from calibration data gathered using the HPD. The LED output pulse shapes were then determined by deconvoluting the measured HPD output waveform with the computed impulse response function of the HPD+amplifier circuit. The impulse response function was derived from a two-pole model of the HPD operational amplifier circuit. The gain-bandwidth product was

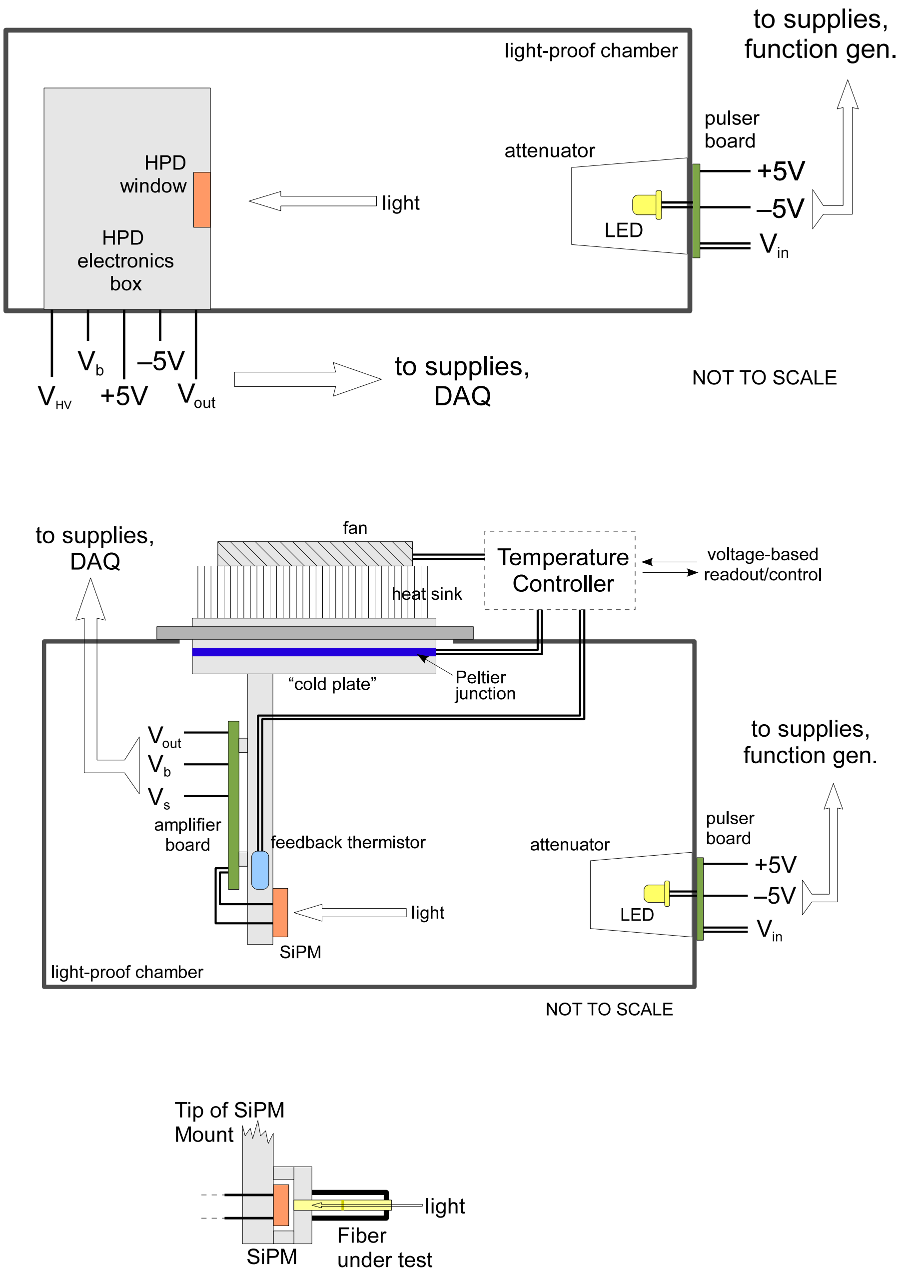

Figure 6: Dark box equipment scheme arranged for the HPD (upper) and SiPM (middle) with direct illumination by the LED pulser. The bottom figure shows the configuration used with the SiPM when illuminated by light transmitted through a fiber, used to measure the transmission of light through fiber glue joints.

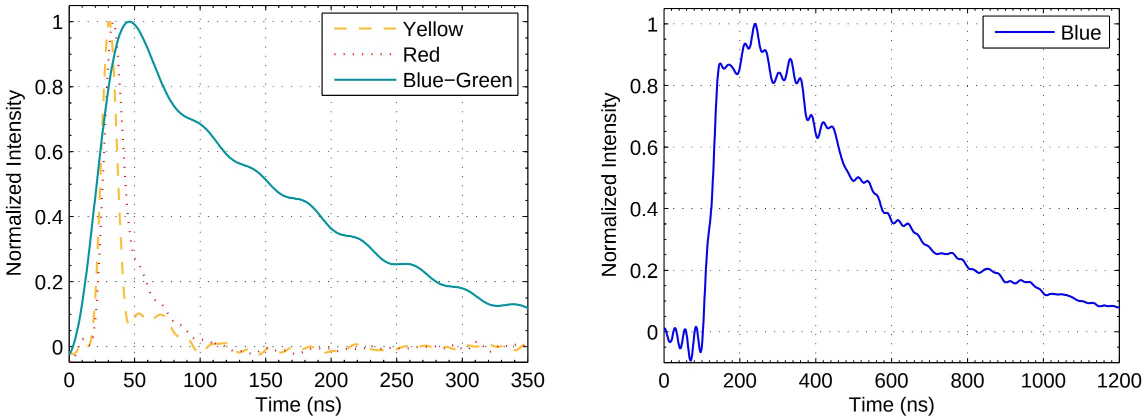

Figure 7: Output light pulse shapes of the LEDs extracted from analysis of the HPD pulse waveforms taken at the maximum amplitude of the pulser. The widths of the primary peaks from the red and yellow LEDs are comparable to the 6 ns width of the current spike from pulser. The extended tails on the blue and blue-green LED outputs are apparently a feature of the electro-optical materials used for these Super-Bright LED devices.

adjusted somewhat from the manufacturer's specification to improve the agreement between the model and the output waveform. The extracted LED pulse shapes are shown in Fig. 7. The ripples in the intensity functions come about from the tendency of a deconvolution to amplify small high-frequency components in the input waveforms, and are probably not real. One unexpected result from this analysis is that the blue LED has a significant long-lived tail in its output that lasts for hundreds of ns. The yellow and red LED pulse widths are comparable to the 6 ns FWHM of the pulser signal, fast enough to emulate the characteristics of a scintillation pulse in BCF-20 scintillator whose decay time constant is 2.7 ns.

Searching for a fast LED closer to BCF-20 in its emission spectrum, the group procured the blue-green Agilent (Avago Technologies) HLMP-CE30-QTC00. The manufacturer-supplied emission spectrum of this LED compares favorably to that of BCF-20, as shown in Fig. 8. The speed of this light source was analyzed using the HPD lineshape, using the same procedure as was used for the original set of LEDs. The deconvoluted shape of the blue-green LED signal is shown together with those of other LEDs in Fig. 7. Although the blue-green LED shows a long pulse decay time, it is nevertheless a factor 3 faster than the original blue one. This is fast enough, for the purposes of SiPM characterization described here, as long as the same integration gate is used for all measurements where an efficiency comparison is made. As far as spectral lineshape is concerned, this LED is an excellent substitute for actual BCF-20 pulses for these tests. The scan in temperature and bias voltage was performed with this LED.

3.3 Test Stand Stability

The reliability of the above-described test stand required verification. Temperature and voltage drifts affecting the LED light output as well as possible variation in ambient light were the principal concerns. Thus a simple stability test was performed: the stand was set to acquire data over the

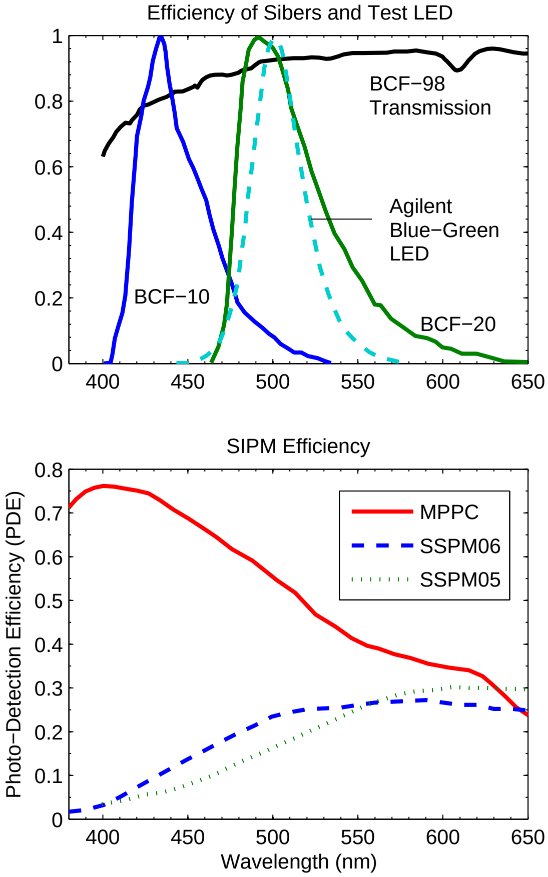

Figure 8: Emission spectra for the scintillating fiber candidates as well as the blue-green (Agilent) LED used in the pulser for SiPM testing (upper panel), and photon detection efficiency curves (lower panel) for the Photonique (blue curves) and Hamamatsu (red curve) SiPMs.

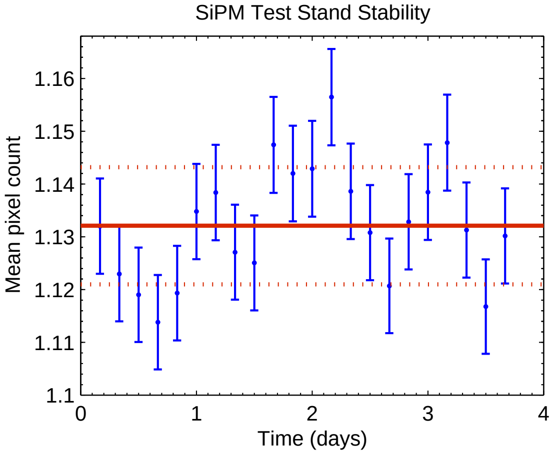

Figure 9: Average number of photons detected per pulse during a continuous 4-day run with the test stand. The mean detected counts are averaged over 4-hour periods for each data point. Error bars are statistical only. The solid curve shows the 4-day average, while the dashed curves show an upper limit on the systematic deviation throughout the run at 90% confidence level.

course of several days. The data was then examined in 4-hour slices against drifts in detection rates. Indeed, no statistically significant drifts in pulser performance was detected.

3.4 SiPM Measurements

There have been three SiPM models whose nominal performance characteristics showed promise and were procured by the group. Two models of SiPMs were acquired from the firm Photonique SA (Geneva, Switzerland), SSPM-050701GR-TO18 and SSPM-0606BG4-PCB. More recently the Hamamatsu MPPC S10362-11-050C was purchased and evaluated in the test stand.

The first SiPM tested was the SSPM-050701GR-TO18. Trials were conducted at the ambient temperature of 22° C and the recommended bias voltage of -40 V. Output waveforms (see Fig. 11) were recorded in two runs of 10000 triggers each. The first run was conducted with the light pulser disabled to measure the dark current. The second run was carried out with a small light pulse present to look for discrete photons and determine the SiPM's detection efficiency relative to that of the HPD. Fig. 12 shows a histogram of the integrals of the output waveforms measured in the second run. Discrete peaks are clearly seen, corresponding to the varying number of photons detected in each pulse. Once individual photon peaks can be seen, it is easy to calibrate the gain of the detector and amplifier. The SiPM detection efficiency relative to the HPD was measured to be 4.0, very close to the expected value 3.9 based on the respective published detection efficiencies for these two devices. Analysis of the first run showed a dark rate of 9.8 MHz, which agrees very well with the manufacturer's specification of 10 MHz at 22°C for this device.

The second SiPM tested was the SSPM-0606BG4-PCB. This device is a newer product, and features a few improvements over the previous device. First, it has a factor 3 larger active area that almost perfectly matches the size of the 2×2 mm^2 fibers to be used in the tagger hodoscope. Second, it has a factor 3 larger dynamic range, with 1700 pixels instead of 556. Third, the specification

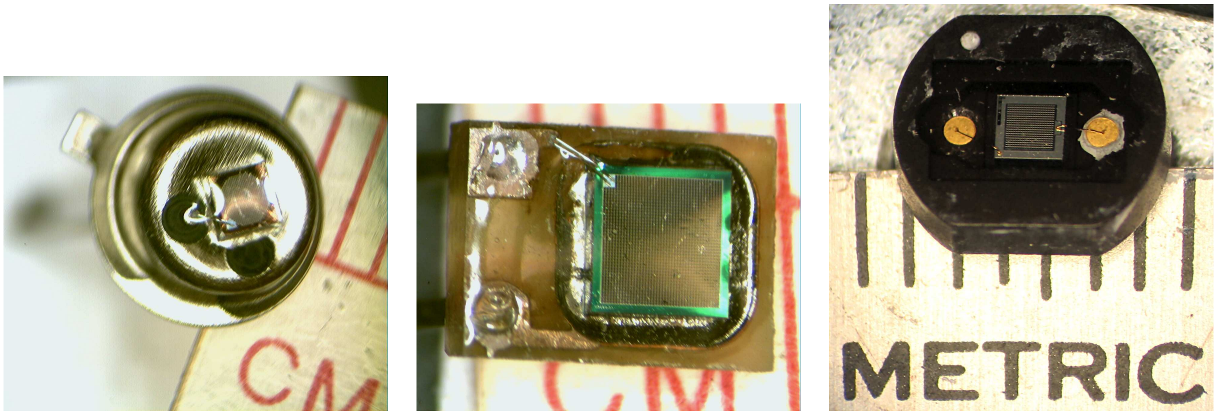

Figure 10: Devices from left to right: Photonique SSPM-050701GR-TO18 (1 mm^2 - 556 pixels), Photonique SSPM-0606BG4-PCB (4.4 mm^2 - 1700 pixels) and Hamamatsu MPPC S10362-11-050C (1 mm^2 - 400 pixels).

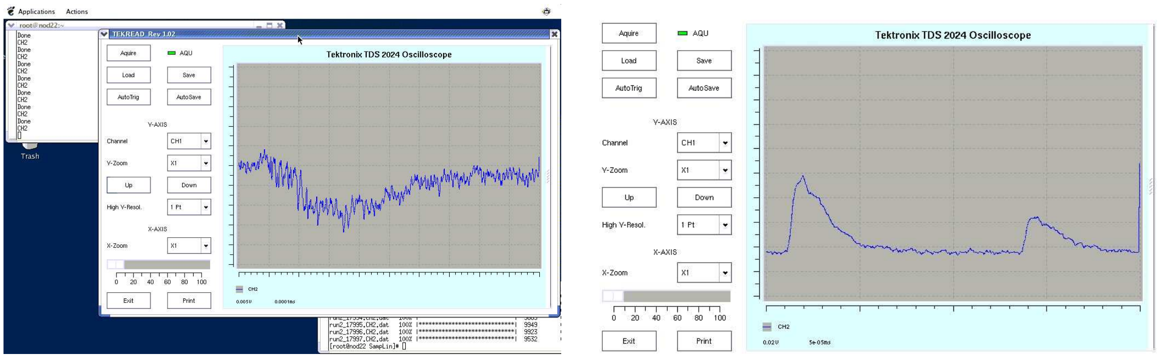

Channel CH1
Y-Zoom X1
High Y-Resol. 1 Pt
X-Zoom X1
0.02V 5e-05s

Figure 11: Screenshot of the data acquisition software. On the left is a screenshot of a captured waveform from produced by the HPD in response to a small LED pulse signal. A similar screenshot on the right shows a triggered waveform from a SiPM.

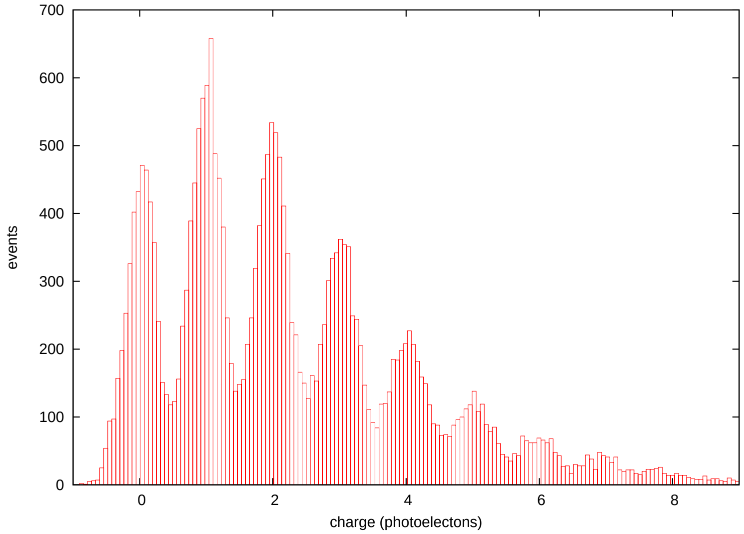

Figure 12: Histogram of the output charge from the 1 mm SiPM when illuminated by small pulses from the yellow LED. The average number of detected photons is about 2 per event. The charge was integrated over a 100 ns gate timed relative to the LED pulser leading edge. The charge was normalized by dividing by the average charge per pixel, which was measured to be 0.25 pC for this device under the measurement conditions.

claims a single-pixel recovery time of 15 ns, which is much faster than is seen in the literature for earlier generations of these devices. The manufacturer's specifications for the TO18 do not give a number for pixel recovery time, but one may guess that the fast recovery of the newer PCB device has been obtained in part by reducing its gain to 1.5×10^5 from 10^6 and bias voltage (20 V instead of 40 V) relative to the TO18 device.

The dark rate for this device was measured to be 8.9 MHz at room temperature of 22°C and a bias voltage of -20.5 V. This compares favorably with the manufacturer's specification of 15 MHz under the same conditions. The total charge per pixel was measured to be $2.5 \times 10^5 e$, as compared with $1.5 \times 10^5 e$ in the specification. The photon detection efficiency was for the two PCB devices tested were determined to be 2.6 and 3.0 relative to the HPD, somewhat different from each other and both lower than the factor 3.6 based on the PDE quoted by the manufacturer.

Fig. 13 shows typical single-shot output waveforms from the 2 mm SiPM illuminated by the three different LEDs. The 200 MHz bandwidth of the oscilloscope does not appreciably distort these waveforms. The rise times of the measured SiPM signals are approximately 10 ns for the yellow and 15 ns for the red LED, which are consistent with the respective LED pulse widths (see Fig.7). The intrinsic rise time of the SiPM is small enough to contribute negligibly to this width, in agreement with the manufacturer's specification of 2 ns for the rise time.

A direct check of the SiPM rise time is possible looking at the single-pixel thermal events which are responsible for the dark current. Fig. 14 shows background events for both the 1 mm and the 2 mm device. These waveforms exhibit considerable noise, but nevertheless single-pixel events are clearly distinguishable from electronic noise. The rise times of these traces are both about 3 ns, with approximately equal contributions from the SiPM amplifier and the SiPM itself. The 200 MHz

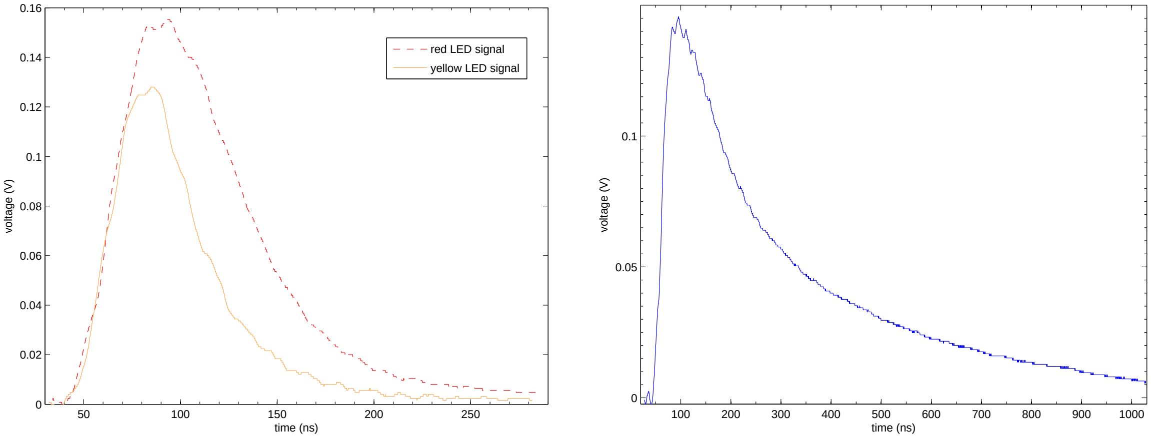

Figure 13: Typical output pulse waveforms from the 2 mm SiPM illuminated with the yellow and red LEDs (left panel) and the blue LED (right panel) at the maximum pulser amplitude. Note the different time scales in the two plots.

oscilloscope bandwidth limitation contributes about 1.2 ns to the total observed rise time, to be added in quadrature with the other sources.

The 2 mm PCB device has a very close match in its active area with the size and shape of the scintillating fiber and light-guide. Its dynamic range of 1700 is a factor of 5-6 larger than the mean photodetection yield expected for microscope pulses, unlike the factor close to 2 for the smaller 1 mm devices from Photonique and Hamamatsu that were tested. This dynamic range overhead is important to ensure that saturation effects in the photodetectors will not degrade the time resolution of the tagger. Even though the quantum efficiency of the 2 mm devices was measured to be somewhat lower than expected, it was still above 20%, a safe margin over the minimum specification of 15%. For these reasons, the group's attention has focused on this device for subsequent development.

3.5 Temperature and Bias Voltage Dependence

The test stand was used to measure the SiPM efficiency, gain, and dark rate as a function of temperature (T) and bias voltage (V_b). Several points in the T, V_b parameter space were sampled. Bias voltages of 19 V, 20 V and 21 V were tested, covering values in the neighborhood of the nominal operating range of 19.5 V - 20.5 V. In the degree of freedom of temperature, samples were taken at 3°C, 20°C and 35°C. A more detailed scan in temperature was performed at V_b =20 V, and a more detailed scan in bias voltage was carried out at 20°C. Figs. 15-17 show the results of these scans.

The data shows that both gain and detection efficiency improve with higher bias voltage and lower temperature. As expected, the dark rate increases with temperature and bias voltage. The tests confirm that at the recommended bias voltage (at or below 20.5 V) and at room temperature these devices satisfy the performance requirements for the tagger microscope. Their PDE, gain, and dark rate are all regular and smooth functions of temperature within 15°C of room temperature and of bias voltage within ±1 V of the recommended V_b. Tuneability of V_b within ±0.1 V is desirable

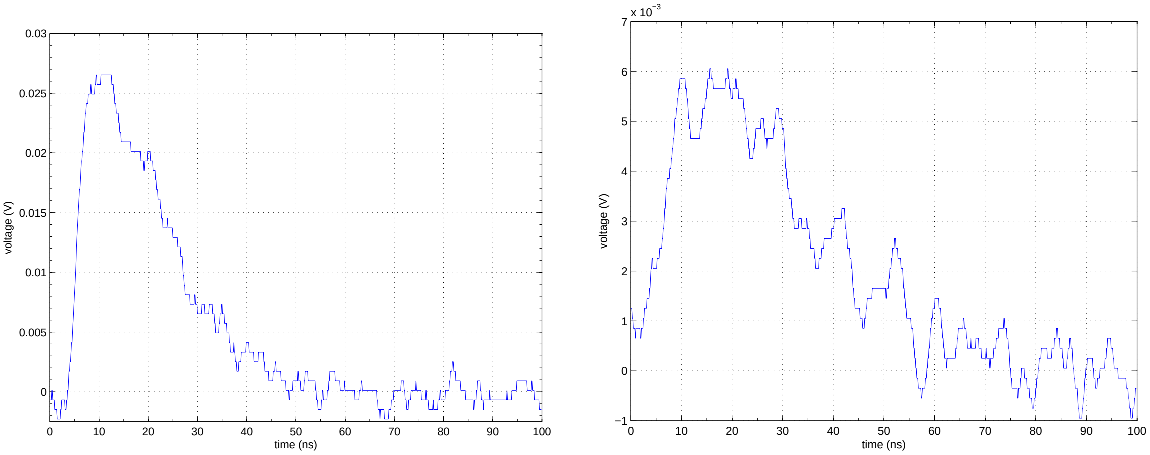

Figure 14: Single-shot output waveforms from thermal single-pixel events in the 1 mm SiPM (left panel) and 2 mm SiPM (right panel).

in order to match the operating characteristics between different channels and maintain a good degree of uniformity in the response across the microscope.

3.6 Cross Talk

The high gain and close packing of the structures in the surface of the SiPM make cross-talk between adjacent cells an important subject in this evaluation. The SiPM output signal is a sum of the currents of all of the pixels, so cross-talk in a SiPM is not the same as that in a mesh-anode phototube where the individual pixels are read out in individual electronics channels. Cross-talk in a SiPM comes about when a single detected photon results in the discharge of more than one cell. This is important because it degrades the dynamic range of the device. If the cross-talk varies across the surface of the SiPM, it can also degrade the signal resolution.

The fact that photons can excite current in the reverse-biased junction means that currents flowing across the junction also produce photons, and if these photons can reach adjacent cells then an avalanche can spread from one pixel to the next. The transient fields created by the discharge of one cell can also stimulate the discharge of neighboring cells by liberating electrons that are localized in shallow potential wells known as “traps” caused by defects in the crystal. These effects determine the upper limits on the bias voltage and the maximum gain possible. Above these limits a single detected photon or a discharge caused by a thermal excitation can spread over a region of cells, and at high enough gain it will propagate across the surface of the entire device. Conversely, the maximum stable operating bias voltage and gain are determined by the condition that the cross-talk be less than some upper limit. It was the discovery of means for controlling the cross-talk that was the critical step in the development of the silicon photomultiplier.

For the tagger hodoscope, the primary impact of cross-talk is through its effect on the dynamic range. Nevertheless it is an important feature of the SiPM and interesting to measure. The effects of cross-talk are seen in the deviation of the pixel multiplicity distribution from a Poisson distribution. For maximum sensitivity, one would like to study pulses containing only a few photons each. The

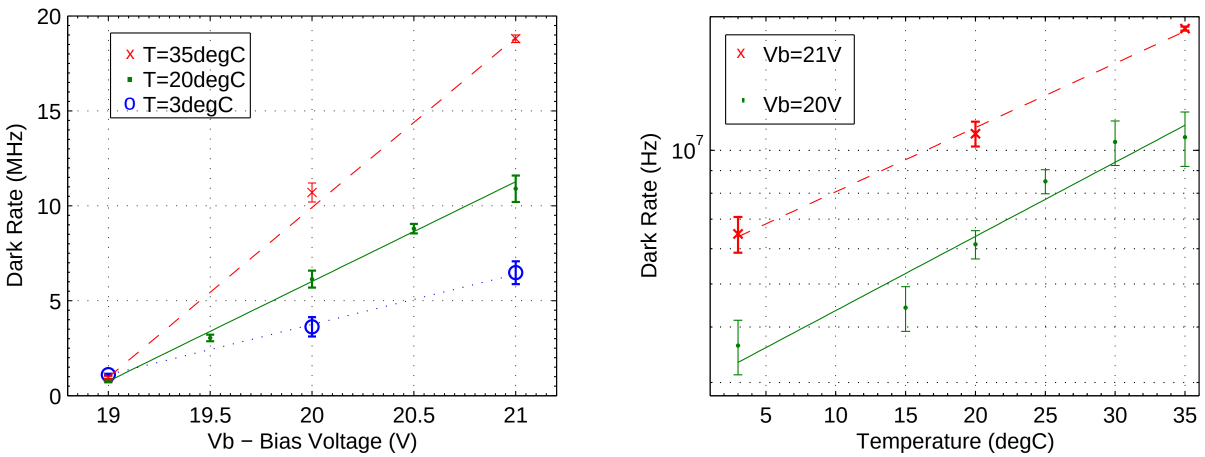

Figure 15: Dark rate plotted as a function of bias voltage at extreme and room temperatures and as a function of temperature at different bias voltages. Note the disappearance of temperature dependence at low bias voltage.

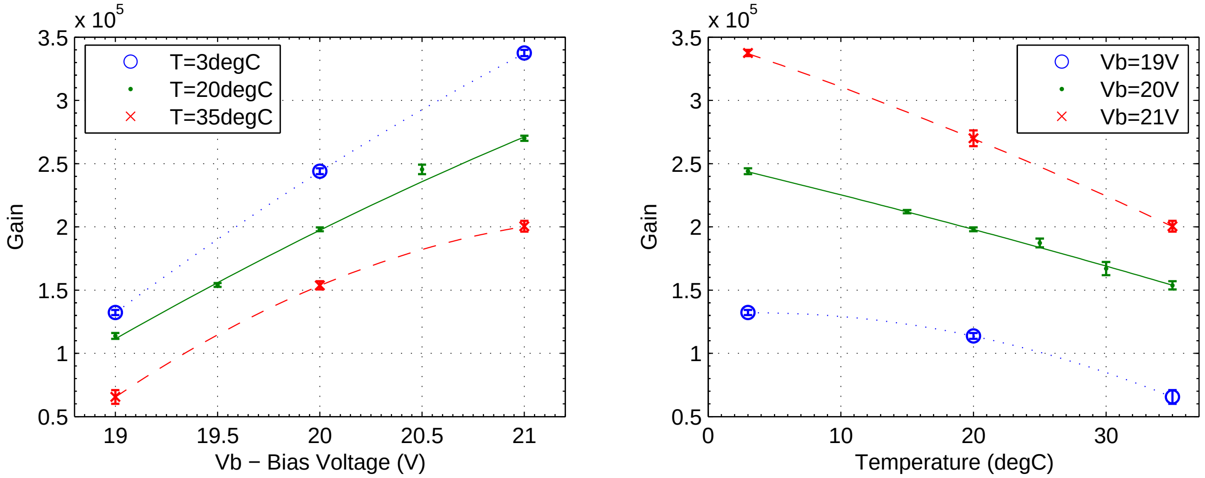

Figure 16: SiPM gain plotted as a function of bias voltage at extreme and room temperatures and as a function of temperature at different bias voltages.

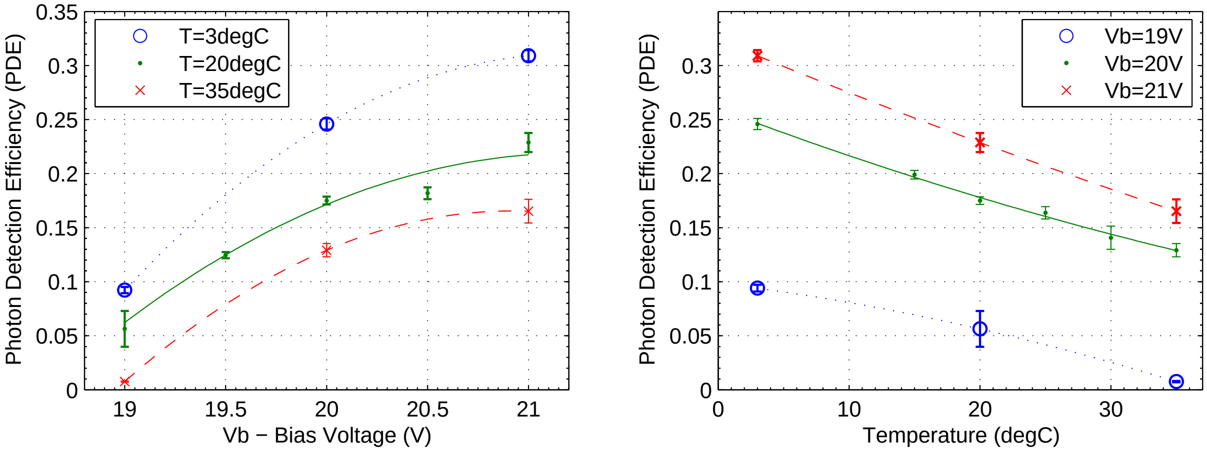

Figure 17: Photon detection efficiency as a function of bias voltage at extreme and room temperatures and as a function of temperature at different bias voltages.

pulser is not very stable when it is operated close to threshold, and it is essential that fluctuations in the pulse intensity not contribute to the statistics of the measurement, so the pulser was set to its maximum value where its fluctuations are only a few percent and the intensity reduced using a passive absorber between the LED and the detector. The resulting pulse-height spectrum for the 1 mm SiPM is shown in Fig. 18. The data have been scaled and shifted so that the peaks occur at integer values and the first peak is at zero, but otherwise the spectrum represents raw sums within a 100 ns gate following the pulse trigger. Fig. 19 shows the results of a similar measurement using the 2 mm SiPM. The pulse intensity had to be reduced for the 2 mm device because of its larger sensitive area.

The reader might wonder, if the oscilloscope is triggered on the LED pulser signal, why there should be a peak at zero photons in these histograms. The reason for this is that the oscilloscope trigger is derived from the sync from the function generator that fires the LED, and does not depend on the output from the SiPM. If the waveforms were self-triggered then the peak at zero would be missing. As a check that this peak is just the SiPM pedestal, a run was taken with the LED disconnected and only the peak at the position labeled 0 in Figs. 18-19 remained.

The curves in the two figures represent fits to a multi-Poisson model with Gaussian broadening of the peaks. The results of the fits are given in Tables 1-2. The goodness of fit of the model to the 2 mm SiPM data is actually somewhat better than it is in the 1 mm SiPM case; the larger χ^2 in Table 2 is simply a consequence of the larger (×6) statistics in the 2 mm sample. As with any real measurement, the most successful model will fail simple statistical tests in the limit of large statistics because of subtle systematic errors. The dominant systematic error in this case seems to be the nonlinearity of the SiPM + amplifier, as evidenced by the deviations of the higher-multiplicity peaks from their nominal positions. Notwithstanding these small systematic effects, the overall agreement with the model is remarkable. A simple Poisson model, in contrast to the multi-Poisson, fails completely to describe these spectra.

Setting the second parameter listed in the tables as p_2 to 1 corresponds to a simple Poisson model. Increasing its value above 1 allows each detected photon to fire a random number of additional pixels, each according to a Poisson distribution of mean $p_2 - 1$. This may be interpreted to mean that the cross-talk probability for the 1 mm device is 30% per pixel, and 9% for the 2 mm SiPM. The value of 30% is significant because it reduces the available dynamic range by 30%. The manufacturer does not specify this parameter, but instead gives a figure for the "excess noise factor", which is interpreted to mean the factor by which the width of the pulse-height distribution for monoenergetic pulses exceeds the Poisson limit of $\sqrt{n}$. For the multi-Poisson model this factor converges to about $1 + (p_2 - 1)/2$ for large pulses (> 10 photons). The measured values are in agreement with the manufacturer's stated limits of < 1.15 and < 1.10 for the 1 mm and 2 mm SiPMs, respectively.

The last SiPM studied was the Hamamatsu MPPC S10362-11-050C. Despite a 1 mm^2 active area, the device seemed promising due to its high quoted figures for its gain and detection efficiency, compared to the Photonique devices. The device has an especially high efficiency in the emission range of the BCF-10 scintillating fiber (Fig. 8). It was thought that the higher efficiency from this scintillator-detector pair might make up for incomplete coverage of the fiber waveguide cross-sectional area.

A statistical analysis analogous to the above was performed on the data from the MMPC test. It was immediately evident that the fitting model described above was not valid for this device. Individual peaks in the charge spectrum did not appear to be symmetric: they possessed a significantly longer left-hand tail. A new model was devised for the spectra from this device that produced satisfactory fits on the lower end of the spectrum. It was clear, however, that the multi-Poisson model could not describe the higher end of the spectrum. Essentially, the distribution did not fall off quickly enough in the limit of higher pixel count. Most pronounced at higher bias voltages, but present even down to threshold, this high-end tail seemed to arise from effects unrelated to primary count or cross-talk statistics. A visual examination of the wave-forms from the MMPC showed a significant number of saturated pulses coming from the simultaneous firing of all (or most) of the pixels in the device, and also significant after-pulsing. The latter is of particular concern in a tagging counter since they contribute directly to the tagging inefficiency. From these tests, it appears likely that these cross-talk effects are artificially inflating the published PDE value that is quoted by the manufacturer. In summary, it does not appear that these Hamamatsu MMPC devices are suitable for use in the tagger microscope.

4 R&D in Scintillation Fiber Coupling

The optics of the tagger microscope consists of a straight segment of scintillating plastic fiber 2 cm long glued to a clear waveguide 50-80 cm in length. The waveguide bends the light out of the plane of the spectrometer, and curves through an arc to direct the light onto the SiPM. There is a small air gap of approximately 1 mm between the waveguide and the active surface of the SiPM. The specified minimum transmission factor for the light within the acceptance of the scintillating fiber reaching the SiPM is 80%. Beside maximizing transmission, optical design must also minimize light leakage between adjacent fibers in order to keep optical cross-talk between channels to a minimum.

The primary source of transmission loss occurs at the joint between the scintillating fibers and the waveguides. Use of optical epoxy is one method to both limit the transmission loss and provide a stable bond between the scintillating fibers and the waveguides. The optical epoxy 20-3238 from the manufacturer Epoxies Etc., with an index of refraction of 1.54 and a tensile strength of 9,400 psi, has been evaluated for this purpose.

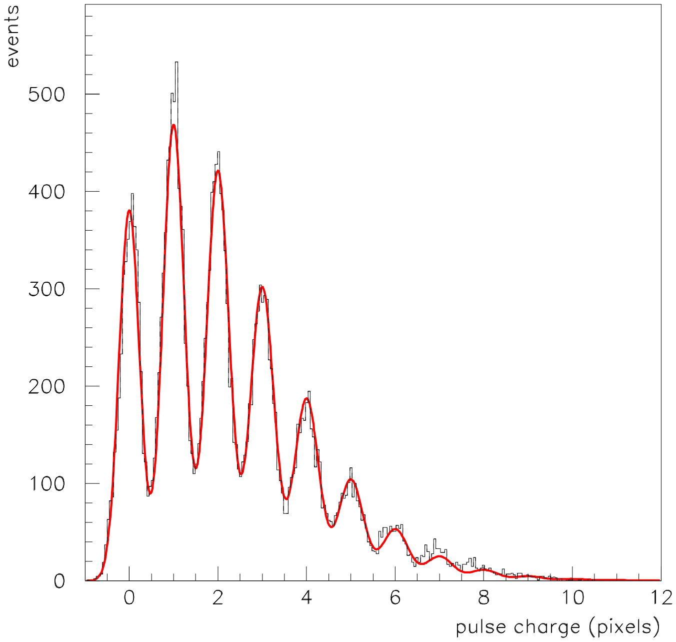

Figure 18: Charge distribution measured with the 1 mm SiPM with the yellow LED and the pulser set near to the minimum detectable level. The curve is a fit to the histogram using a multi-Poisson model which incorporates the parameters shown in Table 1.

Table 1: Parameters returned by the fit shown in Fig. 18.

fit parameter	name	value returned by fit
average photon multiplicity	p_1	1.739 ± 0.013
average pixels per photon	p_2	1.294 ± 0.008
rms width of pedestal (pixels)	p_3	0.230 ± 0.002
additional rms width per pixel	p_4	0.076 ± 0.003
quality of the fit	χ^2/dof	1.57

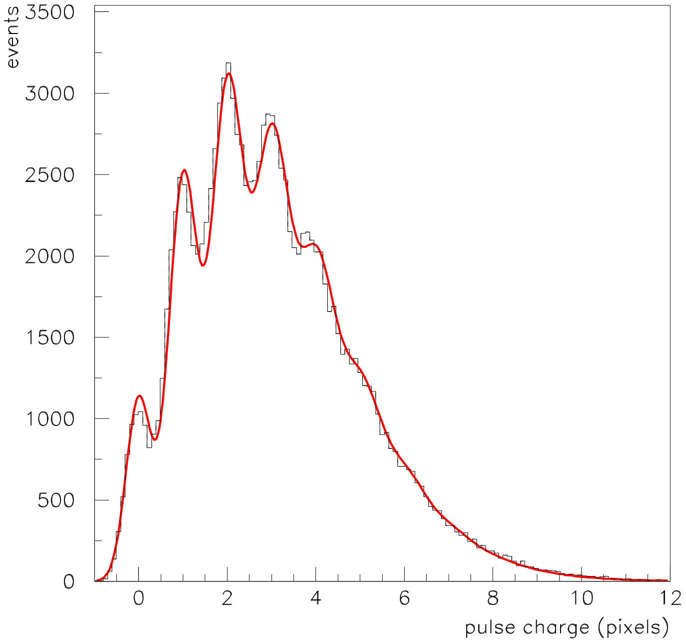

Figure 19: Charge distribution measured with the 2 mm SiPM with the yellow LED and the pulser set near to the minimum detectable level. The curve is a fit to the histogram using a multi-Poisson model which incorporates the parameters shown in Table 2.

Table 2: Parameters returned by the fit shown in Fig. 19.

fit parameter	name	value returned by fit
average photon multiplicity	p_1	2.782 ± 0.008
average pixels per photon	p_2	1.090 ± 0.003
rms width of pedestal (pixels)	p_3	0.284 ± 0.002
additional rms width per pixel	p_4	0.167 ± 0.003
quality of the fit	χ^2/dof	3.87

Before the epoxy was applied, the fibers were first carefully cut and the ends polished. To accomplish this task, several successive steps have been developed, each step adding an extra level of refinement to the final result. Fig. 20 shows photographs of a fiber at four major stages in the polishing procedure. Going through all of the steps, from original cleaving of the fiber to the final fine-grit finish takes about two hours of manual labor to complete for a single end of a fiber. Therefore one of the first questions addressed was to see if a rough polish is sufficient to prepare the ends of the fibers to be joined with glue. Several glue joints were made between segments of clear fiber, each made using ends prepared with a varying degrees of polish.

Fig. 21 shows two fibers mounted in the alignment jig for gluing (left panel) and after the glue has set (right panel). Transmission testing of each of the glued sample waveguides is being performed to determine if the additional polishing steps are necessary to ensure the required transmission rate. In addition, labor-saving means of cleaving and polishing the fibers are being researched.

Other sources of transmission loss include the attenuation of photons traveling through the waveguide, and reflection at the air gap between the waveguide and the SiPM. These sources produce transmission losses of 8% and 4%, respectively. A more significant concern related to the loss of photons along the waveguide is the possibility of their capture in adjacent fiber channels. To prevent this from occurring, the fibers will be coated with a reflective or absorbent material, only a few micron thick. One option is the EMA (extra mural absorber) coating offered by Bicron, the manufacturer of the fibers. Other options include vacuum deposition of opaque coating on the outer cladding of the fibers by a process known as sputtering. Sputtering involves creating a plasma of a specific material then injecting that plasma into a vacuum chamber containing the objects to be coated. The atoms in the cooling plasma then bond to any cool material, forming a thin coating around all exposed surfaces. This technique would permit the choice of a reflective material such as a material for the coating, which would have the additional feature that it would reflect light from the upstream end of the scintillating fiber. Scintillation photons within the critical angle of the fiber axis but moving away from the waveguide will be reflected back, nearly doubling the photon yield. Additionally, this coating may be less than 1 μm thick, compared to 10-15 μm for the EMA. The sputtering deposition could be performed after the glue junction is made, providing interchannel isolation in that region as well as along the length of the fiber, whereas relying on the manufacturer's EMA coating and may result in photons leaking into adjacent channels at the glue joint.

The use of either coating changes the optical properties of the waveguide, generally reducing the overall transmission in the fiber channel. For the EMA coating, the manufacturer claims that the effect of the absorbing layer on the outside of the double-clad fibers is to reduce slightly the acceptance of the fiber, without affecting the attenuation length. It would seem that this effect should be the same for either metallic or absorbing coatings. Investigations are underway to identify a facility where fiber samples may be submitted for sputtering.

5 Control Electronics

5.1 Design Concept

To satisfy the requirement of individual SiPM bias voltage control and monitoring of on-board voltages and temperature, the following scheme was devised. The SiPMs will be mounted in sets of 32 on circuit boards, which also house the analog amplifier and sum circuitry for each tagger

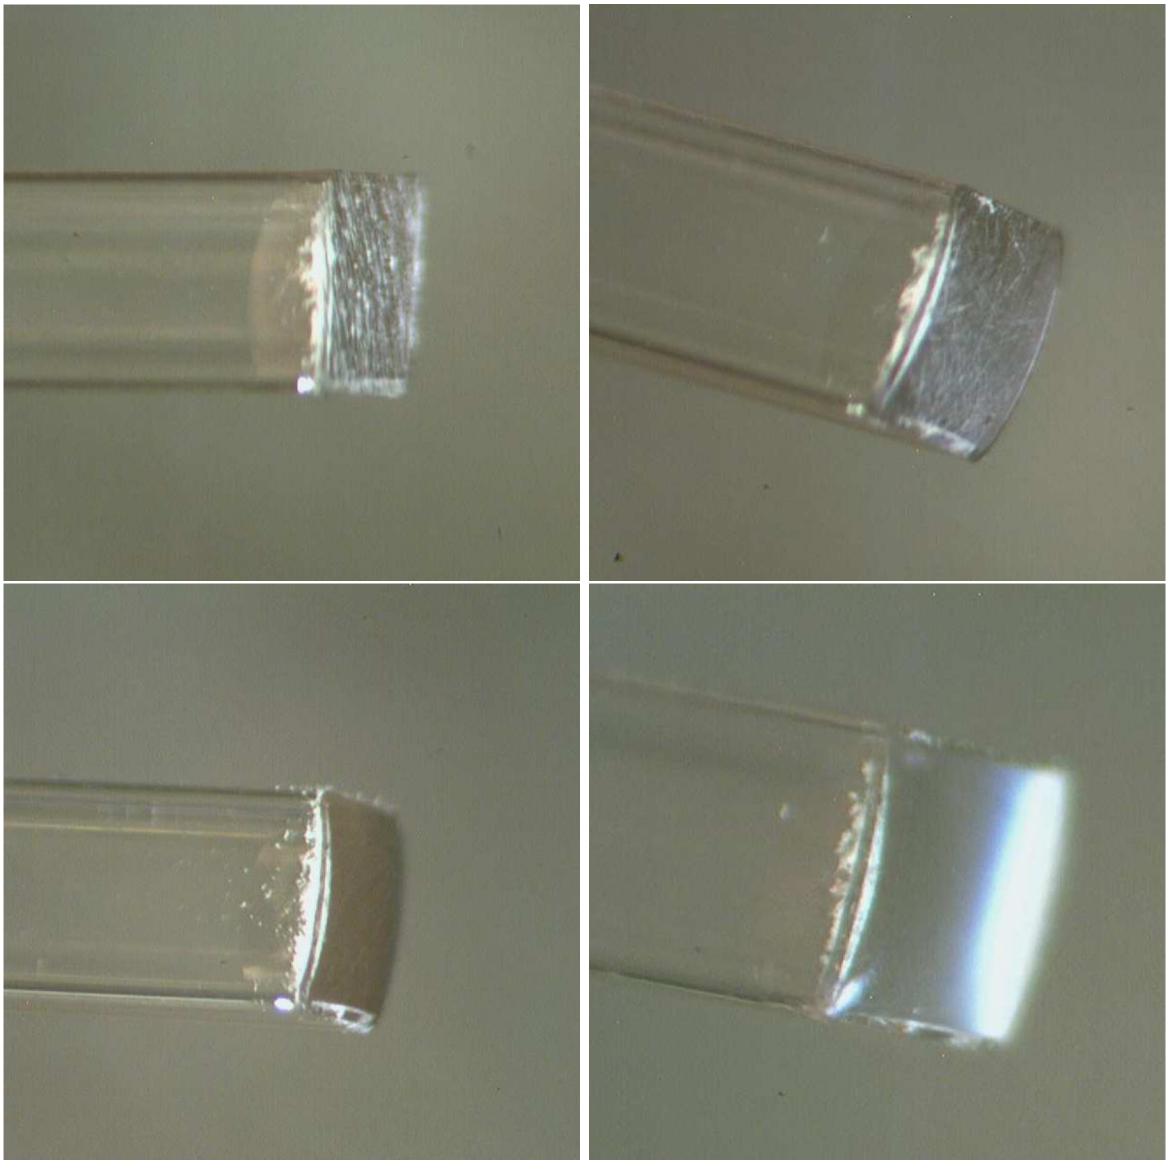

Figure 20: Photographs of the ends of clear plastic fibers of 2×2 mm^2 cross section after cleaving and varying degrees of polish, from rough polish with an emery board (upper left panel), to rough-grit polish (upper right panel), to medium-grit (lower left panel), to fine-grit (lower right panel).

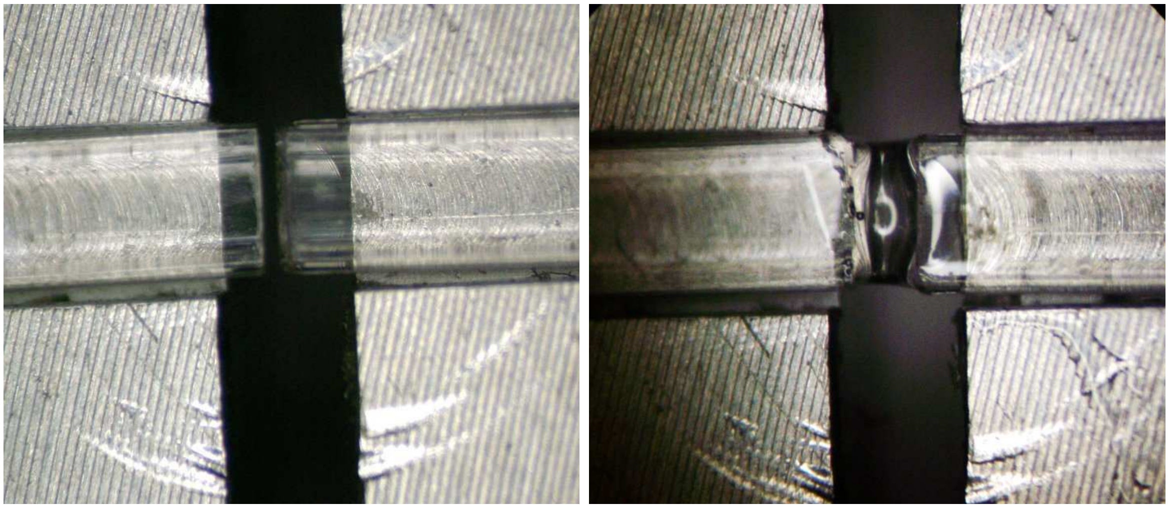

Figure 21: Photographs of an optical fiber joint before (left panel) and after (right panel) the glue has been applied. Note that the volume of glue applied must be carefully controlled to avoid the formation of a bulge at the joint that would interfere with the stacking of the fibers.

channel. A separate circuit board containing the associated digital control electronics will be paired with each analog amplifier board, connected through the main bus-board. The control boards will be equipped with a 32-channel DAC with an output voltage range up to 200 V, which more than covers the range in breakdown voltage foreseeable for an SiPM device. Each DAC channel is individually programmable with 14 bits of precision, which gives better than the required ±0.1 V control over V_b on each individual SiPM. The board will also be instrumented with a Temperature sensor and an 8-channel 12-bit ADC to monitor all the critical reference voltages on the board, as well as one channel of the DAC. Connection between the control boards and the outside world is over standard 10-baseT Ethernet. Ethernet was chosen for the following reasons:

1. ubiquity - Ethernet components are readily available on the consumer market.
2. high-level interface - Controllers are readily available that implement the low-level signaling and address filtering, presenting an asynchronous packet-level interface to the user's i/o controller.
3. robustness - The controller automatically negotiates the link parameters and detects errors, resending packets when necessary and performing integrity checks.
4. flexible addressing - It supports both one-to-one and one-to-many communication that is convenient for device initialization.
5. flexible interconnects - Ethernet cabling and switching provides an abundance of options for interconnects and no distance limits.

An Ethernet Controller chip (EC) will be placed on each control board. The core of the digital board design board will be an FPGA (Field Programmable Gate Array), a convenient and reliable choice for a micro-controller that will interface with all of the above components.

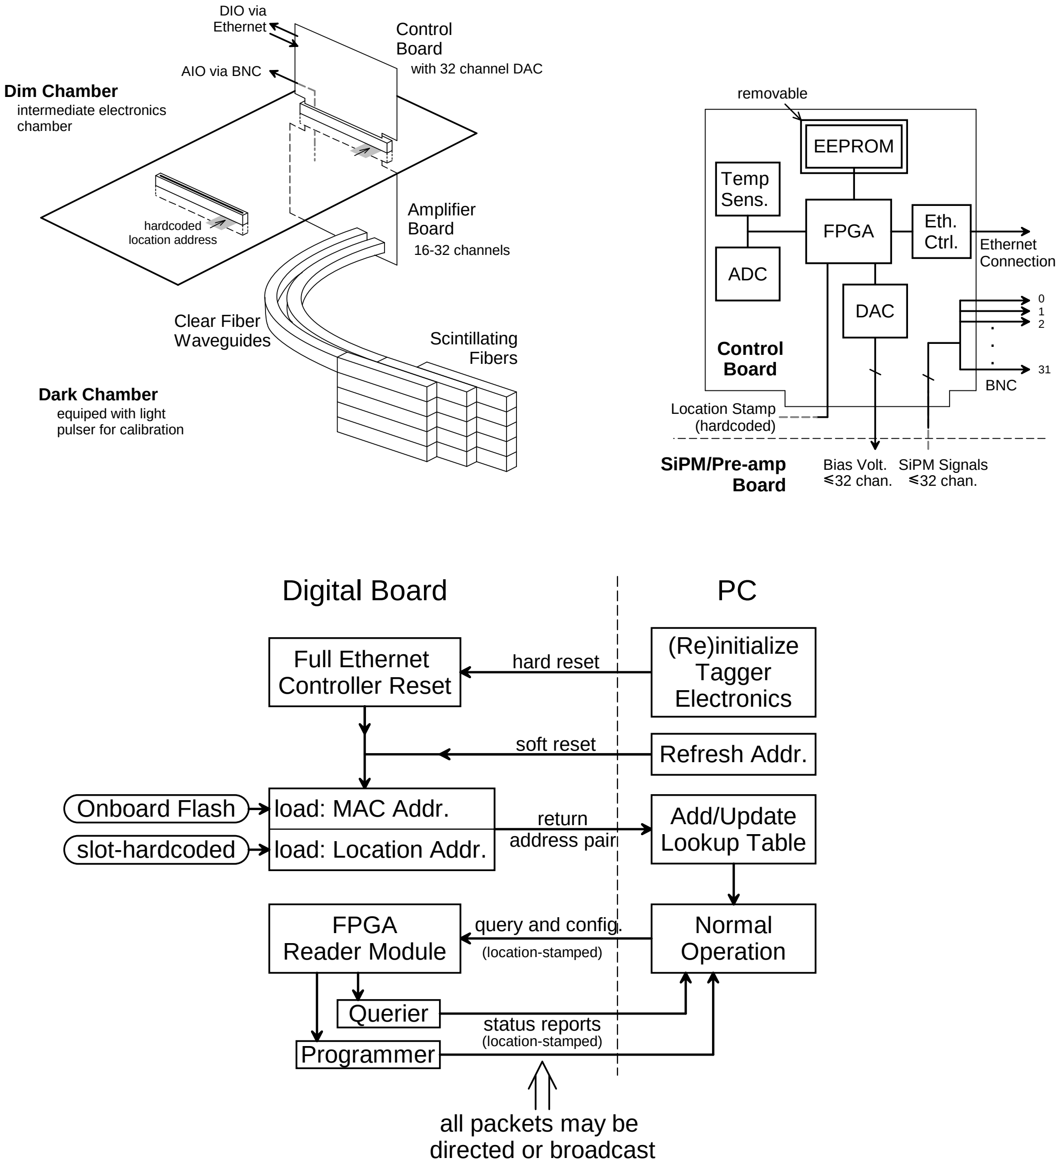

Figure 22: Clockwise: (1) the analog amplifier and digital control boards are shown schematically along with their input and output sources; (2) a detail of the digital control board design; (3) the tagger control flow along with the startup procedure.

Issues of mapping a MAC address to the corresponding control board (representing some known SiPM channel group) have been carefully considered. The current scheme mandates that each control board connection slot be hard-wired with an 8-bit location address, with which all communication packets must be stamped. Upon initialization or in response to a special "census" broacast, each control board sends a standard identification packet to the controlling PC, thereby announcing its presence and associating its hardware slot number to its unique MAC address. The map from SiPM position on a given readout board to the (x, y) index of the associated fiber must be recorded manually during detector assembly, and looked up in a table.

The above scheme requires an on-board EEPROM chip to provide non-volatile storage for the FPGA core. A non-volatile FPGA option was considered, but this option is not thought to provide any more flexibility. Upgrading the core would require a board by board reprogramming that is no less labor-intensive than swapping and rewriting EEPROMs.

5.2 Progress in Design Implementation

All of the chips needed for the above-described functional design have been selected. The core firmware for the FPGA has been written and "fitted" within the resources of the selected Xilinx chip. The design exhausts neither its "slices" nor its input-output blocks (IOBs). Significant numbers of both resources remain, so that future improvements to the functionality of the board should be possible in the form of firmware upgrades.

The group has now obtained access to the requisite licenses for the Altium Designer, the CAD software that will be used for the layout of the two circuit baords. A few prototypes of each board will be produced and populated for testing and eventual use with the prototype detector, whose construction is proceeding in parallel with the electronics effort.

6 Recommendations

The following recommendations are offered as a result of these studies.

1. Complete the design the 32-channel SiPM/amplifier board along with the control board and construct two prototypes of each.
2. Build a 6 × 5 scintillating fiber prototype array and a prototype dark box in which to mount it.
3. Test the prototype electronics boards, and chose one of each for mounting in the prototype detector.
4. Test the assembled prototype detector on the laboratory bench using a LED pulser as a light source.
5. Test the prototype detector during parasitic running with photon beam in Hall B at Jefferson Lab.