Suitability of Silicon Photomultiplier Devices for Readout of a Scintillating Fiber Tagger Hodoscope

I. Senderovich and R.T. Jones

February 15, 2007

Abstract

The University of Connecticut nuclear physics group has produced a design for a scintillating fiber hodoscope to be used in the GlueX tagging spectrometer. The design calls for the readout of individual fiber pulses using silicon photomultipliers (SiPMs). These solid state devices are a relatively new technology and have different sensitivity to wavelength and operating conditions than do conventional photomultiplier tubes. Several SiPM devices have been obtained from industry and characterized in terms of their gain, timing, dark rate, and quantum efficiency. This report describes the methods used for these tests, the results obtained, and outlines the next steps to be taken in the prototyping of a fiber hodoscope with SiPM readout.

Photons with energies up to 12 GeV will be produced in Hall D in Jefferson Lab via coherent bremsstrahlung (CB) of electrons in a diamond crystal. This technique produces a broad radiation spectrum: the energy of the photon is not known *a priori*. Its energy, however, can be measured (i.e. it can be "tagged") by measuring the energy of the post-bremsstrahlung electron that produced it, which has a known initial energy equal to the beam energy: 12 GeV.

This tagging will be implemented by dispersing the post-bremsstrahlung electrons in a magnetic spectrometer and detecting them on the focal plane. This spectrometer is equipped with a broad-band detector array and a higher resolution detector "microscope" covering the narrow energy band of the primary CB peak. The latter detector is instrumented as a packed array of about 600 square 2×2 mm^2 scintillating fibers with silicon photomultipliers (SiPMs) used for photon detection. The following are the principal advantages of SiPMs over the more commonly used photomultiplier tubes (PMTs):

1. no high voltage required (bias voltages of order 50 V are sufficient)
2. response times are about a factor two faster
3. geometric cross section comparable to that of the fiber

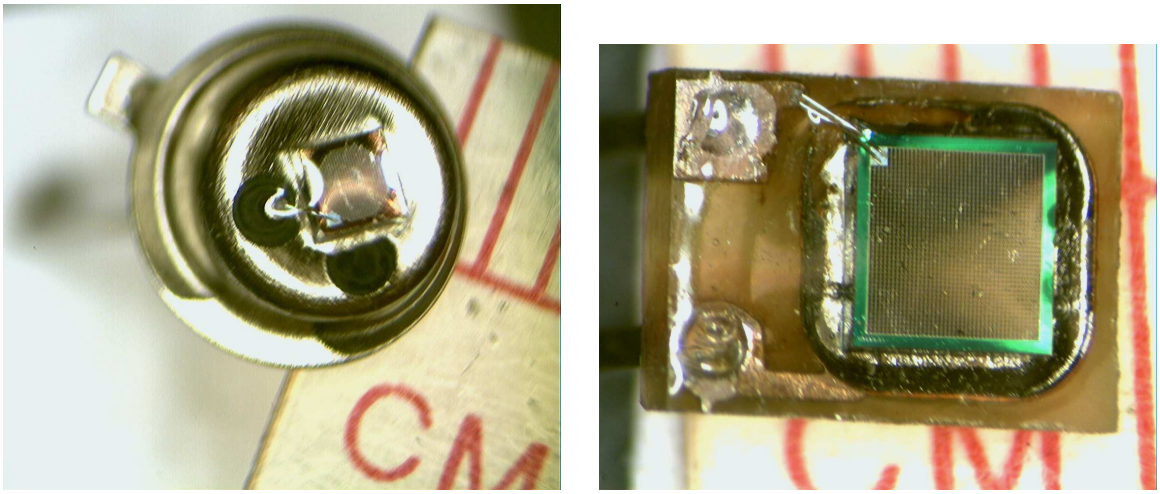

Figure 1: Photonique devices from left to right: SSPM-050701GR-TO18 (1 mm^2 - 556 pixels) and SSPM-0606BG4-PCB (4.4mm^2 - 1700 pixels).

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. The SiPM is are particularly interesting for scintillating fiber applications because of their small size and the good match between their active area and the cross-section of a fiber. However there are a number of areas of concern where SiPMs may prove to fall short of PMT performance: dark rate, sensitivity to operating conditions, cross-talk between pixels, and dynamic range. This report examines the requirements for these parameters of the focal plane fiber hodoscope readout and presents bench-top measurements of their values for two candidate SiPM devices.

Two models of SiPMs were acquired from the firm Photonique SA (Geneva, Switzerland), SSPM-050701GR-TO18 and SSPM-0606BG4-PCB. Photographs of the two devices are shown in Fig. 1. The first device has an approximately circular active area of 1 mm^2, and is contained inside a metal TO18 package connected to two external leads. The second device has a square active area of 2.1×2.1 mm^2, and is mounted on a small printed circuit board with two external leads. Fig. 2 shows the two devices connected to a compact preamplifier circuit that was produced for use with the SiPMs by the manufacturer. The preamplifier serves to buffer the SiPM signal. It has a transimpedance gain 3 kΩ and pulse rise [fall] times of 2 [25] ns.

1 Requirements

The fiber hodoscope consists of many identical readout channels, each connected to a single scintillating fiber through a clear acrylic fiber light guide. Each scintillating fiber and light guide and is optically isolated from its neighbors by two layers of cladding and an outer coating of opaque material. A tagging electron travels axially down the full 2 cm length of a scintillating fiber deposits 4 MeV of energy in the fiber. This results in 1600 scintillation

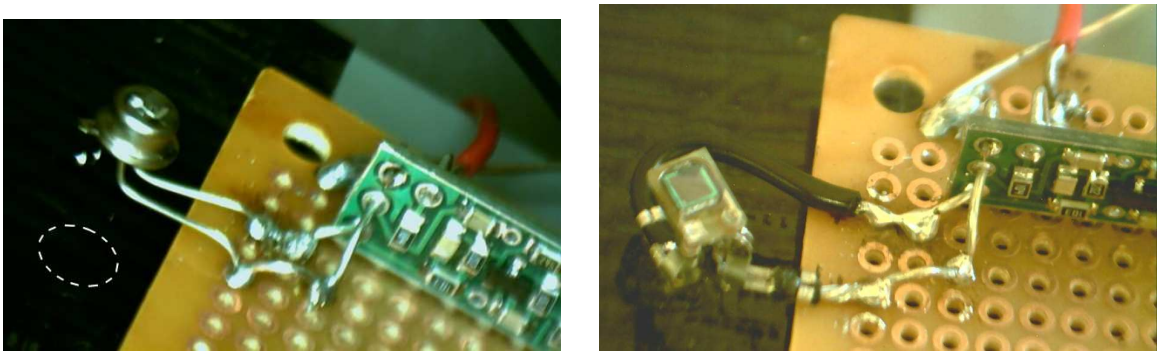

Figure 2: SiPM devices from Fig. 1 mounted with their amplifier.

photons within the forward capture cone of the fiber. Assuming that 80% of these are delivered to the SiPM active surface and a conservative estimate of 15% for the efficiency of the SiPM leads to an estimate of 190 for the average SiPM pixel count per event. Monte Carlo simulations have shown that an efficient detection threshold corresponds to 40% of the average pulse height, or 80 SiPM pixels. This implies a requirement for the dark current that the rate of spontaneous pulses containing 80 or more pixels must 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.

The large photon yield in these pulses results in a very soft requirement for the SiPM detection efficiency. In the wavelength range of fast-green scintillator (BCF-20), the SiPM manufacturer's stated detection efficiency is 15% at 22.5°C for the 1 mm device, 21% at 22.5°C for the 2 mm device, both well within the required range for the microscope readout. The plots of detection efficiency versus wavelength are provided by the manufacturer under specific conditions of temperature and bias voltage. Apparently these devices are sensitive to such operating conditions, which in turn imply specific tolerances on the conditions in the tagger hall where the detectors must operate for extended periods without significant maintenance. The design approach taken in this area is to measure the dependence of the gain, efficiency and dark rate on operating conditions and then invert the results to produce tolerances on environmental factors that are consistent with robust and stable tagger operation.

To the extent that the individual pixels of the SiPM device are independent detectors, the statistical fluctuations of the pulse height are given by the Poisson distribution. Under this assumption, the chances that 80 thermal avalanche hits will occur within a 100 ns interval to produce a spontaneous pulse over threshold is entirely negligible. However it is expected that some fraction of the time an avalanche in one pixel will induce an avalanche in another pixel, either through electrical or photon cross-talk effects. Fluctuations in the size of these multi-pixel dark rate events may define the minimum feasible operating rate for the tagger. It is a requirement for the success of the SiPM readout that this rate be low enough to permit tagging at 1 nA electron beam current.

The dynamic range of the SiPM is determined by the number of pixels illuminated by the fiber. Dynamic range is not a critical parameter for tagging counters because they function

as digital devices. Taking into account the needs of initial set-up and gain equalization, a dynamic range of 300 would be sufficient. At high rates the available dynamic range can be reduced by pile-up effects arising from the recovery time of individual pixels after they discharge. Early SiPM devices had pixel recovery times in the range 100-200 ns. Tagging counters must function at rates up to several MHz per channel. At 5 MHz and a single-pixel recovery time of 200 ns, the available dynamic range is reduced by a factor of 0.63 relative to the dynamic range available at low rates. Photonique now claims that they can produce SiPM devices with pixel recovery times less than 20 ns, in which case pile-up effects should be negligible even at the maximum rates seen in the tagger.

2 Test Stand

Characterizing these prototypes requires a testing environment in which fast light pulses of controlled amplitude can be generated and in a low-background environment. A test stand has been constructed for this purpose. It is essentially a hermetically sealed chamber with a pulsed LED in one end and a mount for a photodetector in 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. 3 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. 4.

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

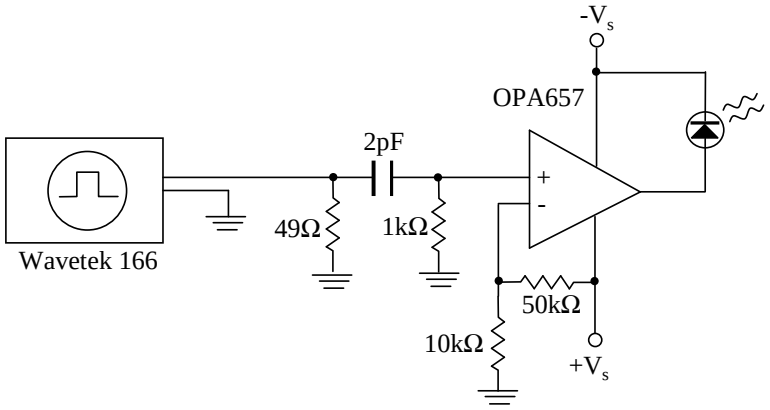

Figure 3: Diagram of the LED pulser circuit. The input square wave from the function generator is differentiated and used to drive the high-frequency operational amplifier OPA657. The value of V_s is 5 V.

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. 4. A typical output waveform from the HPD in response to a small input pulse is shown in the figure. A screenshot of the data acquisition system is shown in Fig. 5.

3 Light Sources

Three LEDs with different emission spectra were calibrated with the HPD: blue (QT Optoelectronics MV5B60), yellow (Fairchild MV8304) and red (Fairchild MV8104). The emission ranges published by the manufacturers for these devices are indicated by the horizontal bars in Fig. 7.

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 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

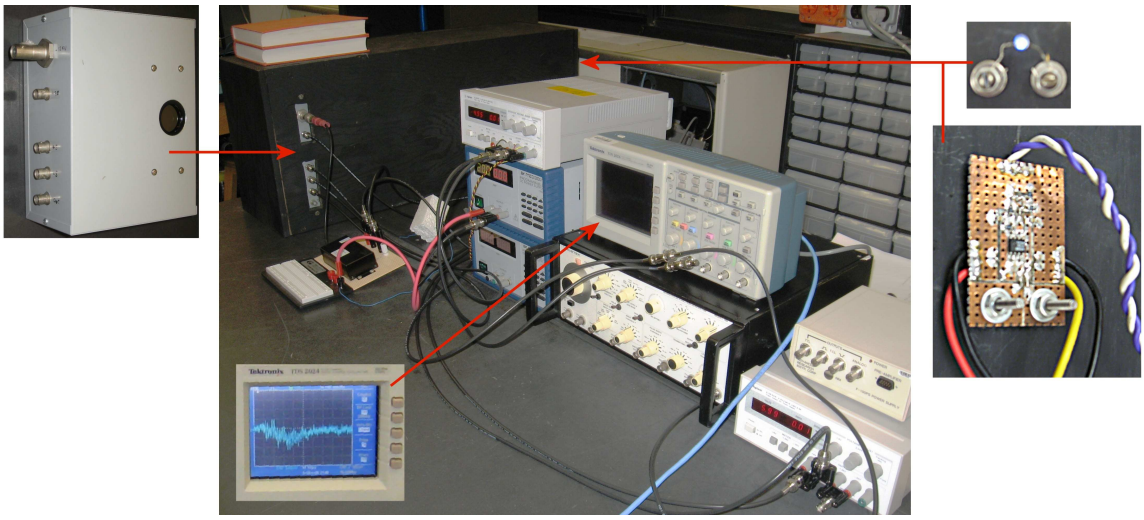

Figure 4: Test stand during measurement of the HPD response. The black box in the back of the central photograph is the dark chamber used to test the photodetectors. The HPD unit mounted on the wall of its electronics enclosure is shown separately on the left. The LED pulser is shown on the right with its pulse-shaping amplifier circuit.

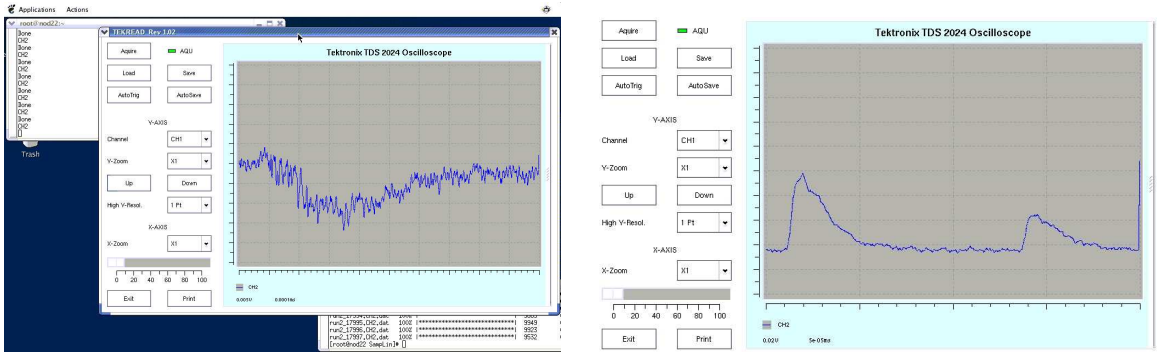

Figure 5: Screenshot of the data acquisition software. On the left is a screenshot of a captured waveform from produced by the HPD in response to a small LED pulse signal. A similar screenshot on the right shows a triggered waveform from a SiPM.

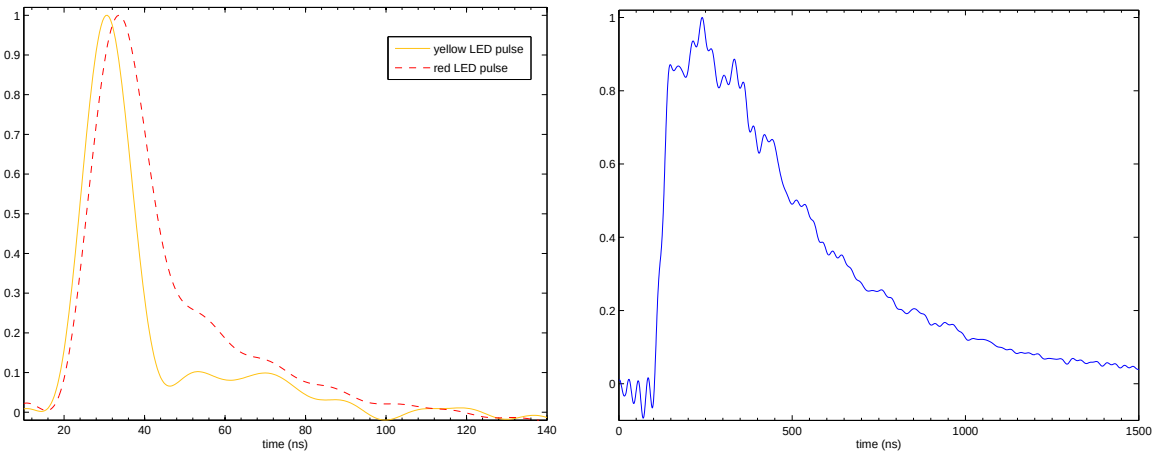

Figure 6: Output light pulse shapes of the LEDs extracted from analysis of the HPD pulse waveforms taken at the maximum amplitude of the pulser. The widths of the primary peaks from the red and yellow LEDs are comparable to the 6 ns width of the current spike from pulser. The extended tail on the blue LED output is apparently a feature of the material used for these super-bright blue LEDs.

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 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. 6. 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 consistent with decay lifetimes comparable to that of the scintillation in BCF-20 (half-life 3 ns).

4 Photon Detection Efficiency

The photon detection efficiencies of the HPD, SSPM-05∼ and SSPM-06∼ weighted by the emission spectrum of the yellow LED are 8.7%, 29% and 24% respectively, based upon the manufacturer's specifications. The low efficiency of the HPD in the yellow is due to the fact that it is designed for maximum efficiency in the blue. The expected efficiencies of the SiPMs relative to the HPD in the yellow are therefore 3.4 and 2.8. The manufacturer-supplied detection efficiency curves of both SiPMs are shown in Fig. 7.

5 SiPM Measurements

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

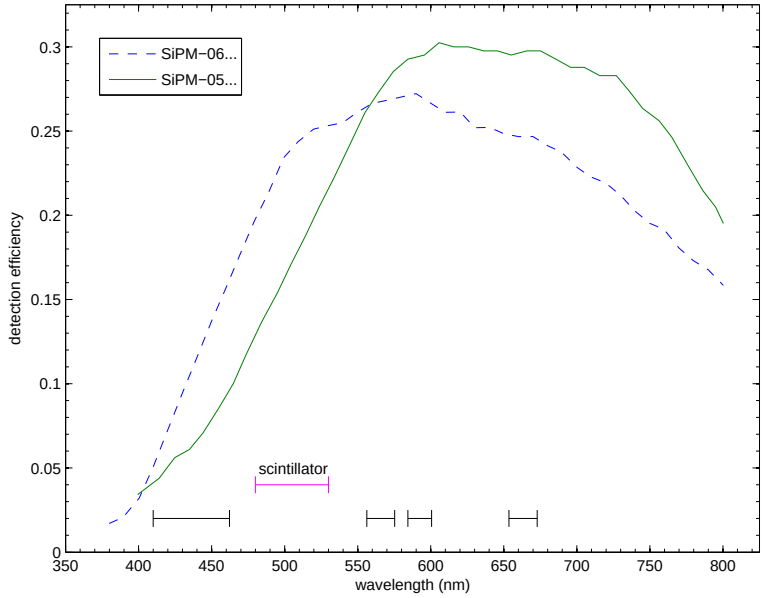

Figure 7: Photon detection efficiency of the the SiPMs as a function of wavelength, as stated by the manufacturer. The emission ranges (one standard deviation) of blue, green, yellow and red LEDs are shown as horizontal bars.

(see Fig. 5) were recorded in a 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. 8 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 3.2, very close to the expected value 3.4 based on the respective published detection efficiencies for these two devices. Analysis of the first run showed a dark rate of 9.8MHz, 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 claims a single-pixel recover time of 15 ns, which is much faster than is seen in the literature for past generation devices. The TO18 specifications do not give a number for pixel recovery time, but one may guess that the fast recovery of the newer device has been obtained by reducing its gain (1.5×10^5 instead of 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

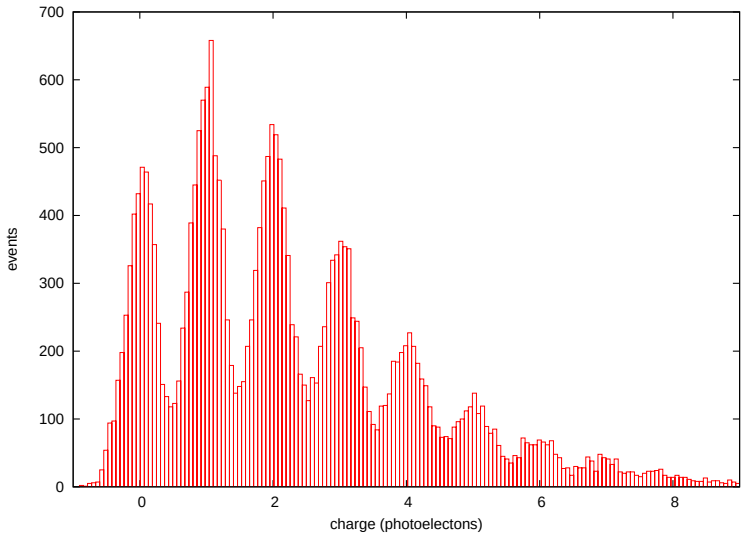

Figure 8: Histogram of the output charge from the 1 mm SiPM when illuminated by small pulses from the yellow LED. The average number of detected photons is about 2 per event. The charge was integrated over a 100 ns gate timed relative to the LED pulser leading edge. The charge was normalized by dividing by the average charge per pixel, which was measured to be 0.21 pC for this device under the measurement conditions.

of 15 MHz under the same conditions. The total charge per pixel was measured to be $2.3 \times 10^5 e$, as compared with $1.5 \times 10^5 e$ in the specification. Based on the measured gain, the photon detection efficiency was determined to be 8.38 relative to the HPD at the wavelength of the yellow LED. This result is more than a factor 2 larger than the expected value of 2.76 taken from the manufacturer's photon detection efficiency curve for 22.5°C. This anomaly is still under investigation. One possible explanation is that multiple pixels may be firing in response to a single detected photon through some kind of cross-talk. This possibility is ruled out by the analysis presented in the next section. A second explanation is that the LED output has somehow changed since the HPD calibration was carried out, or the LED has moved so that it now illuminates the position of the detector more directly than it did before. This is one of the first questions to be addressed in the next phase of the hodoscope readout prototype project.

Fig. 9 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.6). 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. 10 shows background events for both the 1 mm and the 2 mm device. These waveforms exhibit considerable noise, but

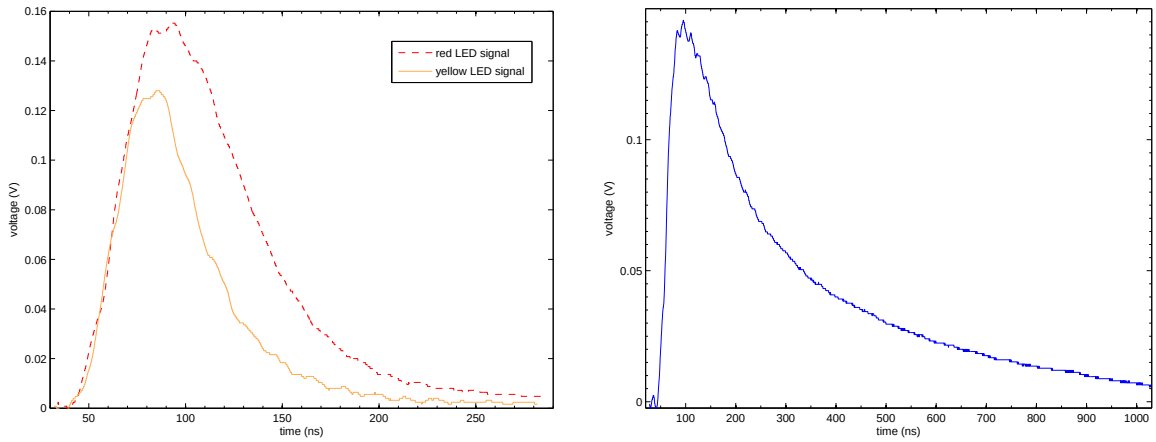

Figure 9: Typical output pulse waveforms from the 2 mm SiPM illuminated with the yellow and red LEDs (left panel) and the blue LED (right panel) at the maximum pulser amplitude. Note the different time scales in the two plots.

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

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

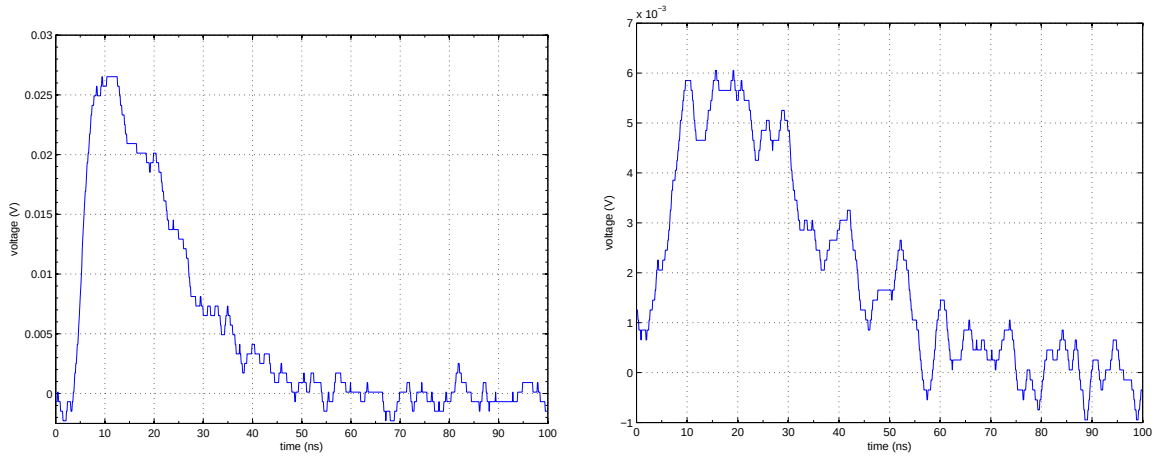

Figure 10: Single-shot output waveforms from thermal single-pixel events in the 1 mm SiPM (left panel) and 2 mm SiPM (right panel).

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 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 are shown in Fig. 11. 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. 12 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 why, 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 not 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. 11-12 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 ($\times 6$) 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 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 is 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.

7 Temperature Dependence

One important feature of the SiPMs that was not addressed by the above measurements is the sensitivity of the photon detection efficiency and the dark rate to the operating temperature and bias voltage. In preparation for systematic measurements of this dependence, the test stand is presently being fitted with a temperature-controlled photon detector mount. A 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 controller that senses the temperature using a thermister 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 $-20° - +40°$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 and a potentiometer. The controller adjusts the temperature of the cold plate by sensing the resistance of a feedback thermister attached to the plate and driving current in one direction or the other through the Peltier junction until the thermister resistance matches a value that it looks up in a built-in lookup table for the given set point. In order to check the manufacturer's set point voltage vs. temperature curve, a second thermister with a known R/T curve was coupled to the cold plate and a series of measurements taken. These measurements showed that the controller's calibration was not very accurate, so the entire calibration was repeated on the bench. The reference thermister R/T curve was measured over the range $0 - 50°$C using a standard laboratory thermometer immersed in a water bath as a reference. The reference thermister was then attached to the cold plate and used to

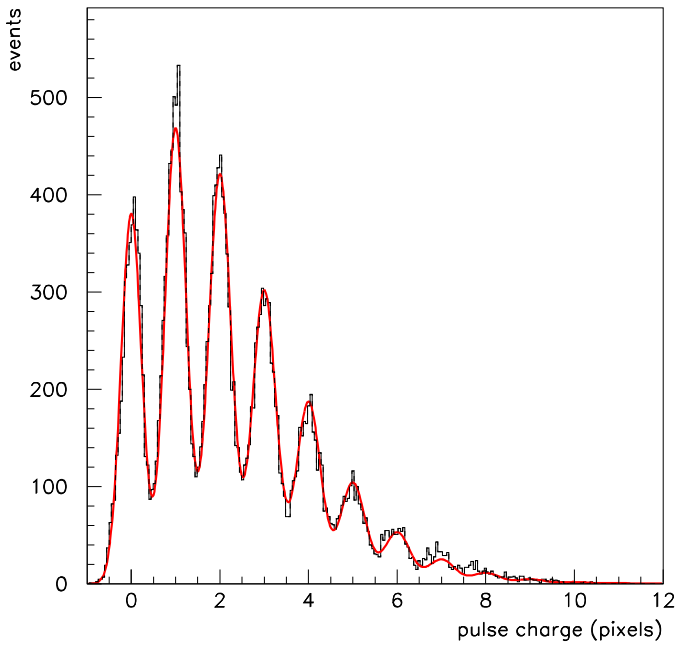

Figure 11: Charge distribution measured with the 1 mm SiPM with the yellow LED and the pulser set near to the minimum detectable level. The curve is a fit to the histogram using a multi-Poisson model which incorporates the parameters shown in Table 1.

Table 1: Parameters returned by the fit shown in Fig. 11.

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

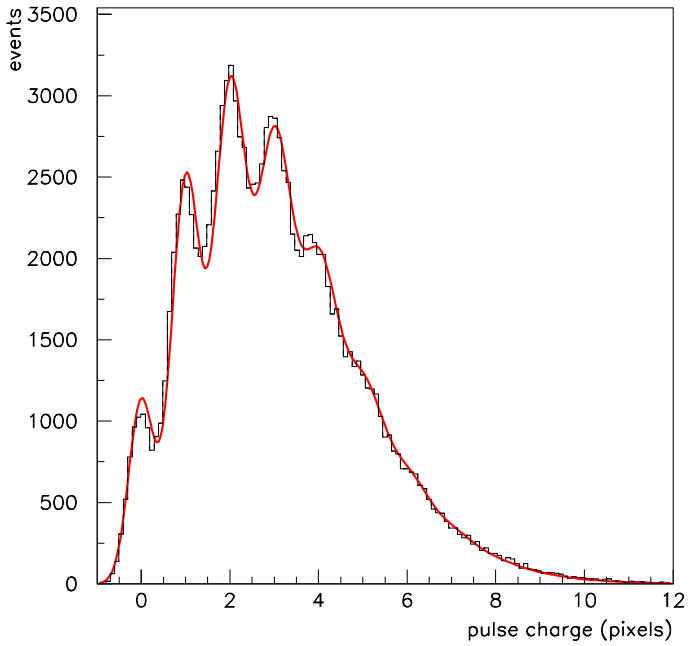

Figure 12: Charge distribution measured with the 2 mm SiPM with the yellow LED and the pulser set near to the minimum detectable level. The curve is a fit to the histogram using a multi-Poisson model which incorporates the parameters shown in Table 2.

Table 2: Parameters returned by the fit shown in Fig. 12.

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

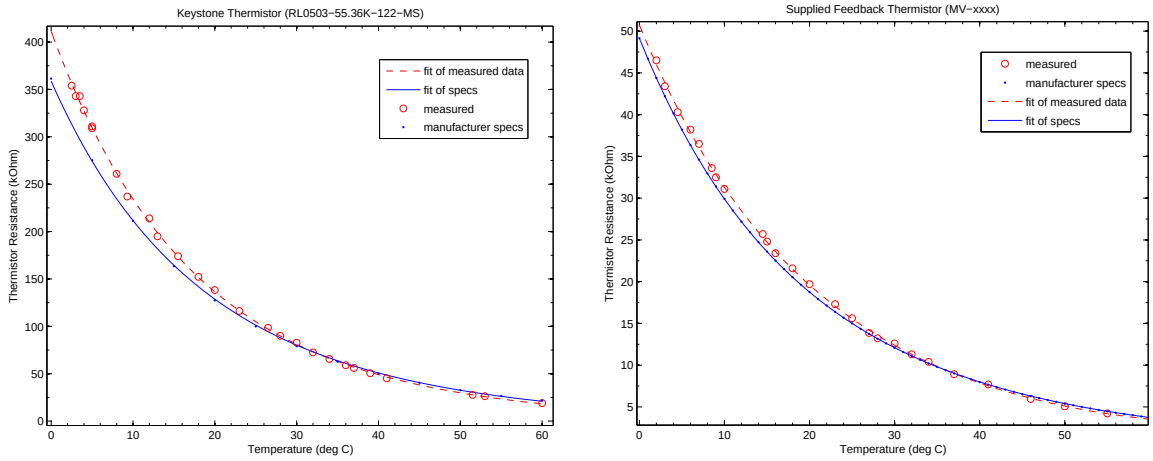

Figure 13: R/T curves for the reference thermister (left panel) and the temperature controller feedback thermister (right panel).

calibrate the thermal controller.

The R/T curves for both thermisters are shown in Fig. 13. Once the cold plate temperature could be accurately measured, the set-point V/T curve for the controller was recalibrated. The results of this calibration are compared with the manufacturer's V/T curve in Fig. 14. This calibration can be extrapolated below 0°C to reach lower temperatures. However it is intended that the SiPM measurements will focus on temperatures above the freezing point of water because there is little motivation to consider trying to operate these electronics below freezing in the humid laboratory environment.

Some modifications to the test stand had to be made in order to install the cold plate. The cold plate was fitted into the dark chamber wall and the perimeter carefully sealed to avoid light leaks. An aluminum post is mounted on the cold plate and extends into the interior of the dark chamber to make thermal contact with the detector, as shown in Fig. 15. An insulating thermal paste is used to ensure that there is good thermal contact between the post and the SiPM without risking an electrical short. The heat sink and fan attached to the Peltier junction on the exterior of the dark chamber provides a heat source or sink for the thermal mount. The chamber was fitted with additional electrical feed-throughs for the Peltier junction current and feedback sensor connections. Photographs of the thermal mount assembly are shown in Fig. 16. Measurements will begin as soon as the electronic check-out is complete.

8 Results and Recommendations

Two different SiPM devices produced by Photonique have been tested on the bench. Both devices were observed to have gain, efficiency, and dark rate performance parameters that are in agreement with, and in some cases better than, the specifications provided by the manufacturer. Either device would meet the requirements of the fiber readout for the Hall D tagger microscope. The smaller device has 556 pixels covering a circular area of

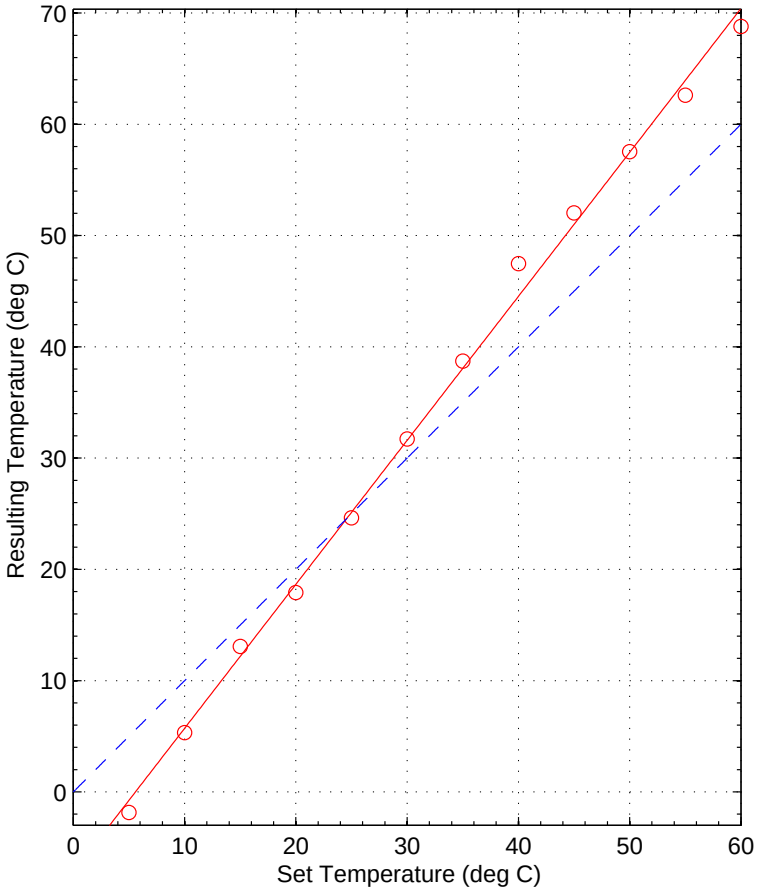

Figure 14: Plot of the measured temperature of the cold plate versus the set temperature according to the manufacturer's set-point V/T curve. If the original calibration were accurate then the measured points would fall on the dashed line.

Figure 15: The aluminum extension post designed to put the SiPM in good thermal contact with the cold pate. The feedback thermistor and the SiPMs are mounted on the post.

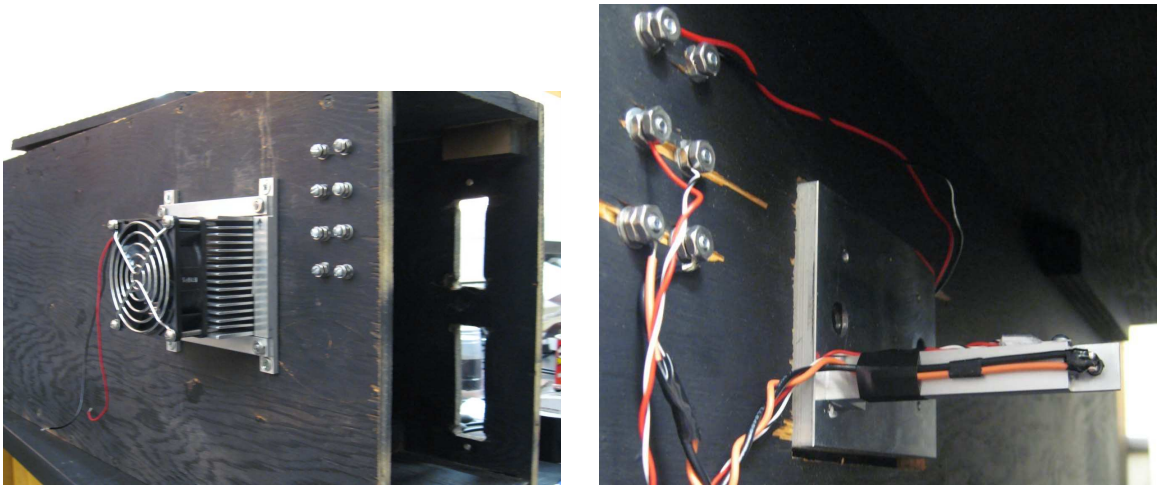

Figure 16: Photograph of the test stand with the thermal control system installed. The view from the outside of the heat sink and fan is shown on the left. The view from inside the chamber is shown on the right with the assembly shown in Fig. 15 installed.

approximately 1 mm diameter, so mating it to a 2 mm square fiber would necessarily involve some loss in intensity. A direct coupling would produce a photon count of 50 for the most-probable signal from a tagging electron in a microscope fiber, and a threshold signal of 20. This is close to the minimum level required for efficient operation. At these light levels, the limited dynamic range of 400 for this device is not an issue.

The primary concern with the choice of the smaller device for the tagger hodoscope is the time resolution that would be obtained from pulses containing only 50 photons. A time resolution of 200 ps is desirable in the tagger in order to ensure the clear identification of an electron with a particular bunch whose period is 2 ns. A crude estimate for attainable time resolution from a scintillator pulse is $\tau/\sqrt{n}$ where τ is the decay time of the scintillator and n is the number of detected photons. The value of τ for BCF-20 fast-green scintillator is 2.7 ns. This suggests that 180 detected photons would be desirable. A number close to half that value might be achieved by fabricating Winston cone couplers between the fibers and the 1 mm SiPM detectors, but a simpler option would be to chose the 2 mm detector.

The square 2×2 mm^2 device from Photonique is a particularly good match to the tagger hodoscope requirements. The manufacturer's curve for the detection efficiency of this device at the peak of the BCF-20 emission spectrum is 20%, in contrast to 15% for the 1 mm device. The most-probable signal of 250 detected photons expected for a tagging electron in a fiber is safely inside the linear region of the dynamic range of a device with 1700 pixels, even if the illumination factor varies by a factor of 3 from the center to the edges of the active area. The gain of 2 pC/pixel obtained with this device using the manufacturer-supplied amplifier circuit produces an output pulse height of 300 mV for a threshold signal of 100 detected photons in a fiber. The pixel recovery times are short enough to permit operation at rates a factor 3 higher than the maximum tagging rates intended to be used in Hall D. The detected photon count is sufficient to meet the requirement of 200 ps time resolution,

but will require confirmation in a beam test of the prototype detector.

The SiPM dark rate is dominated by single-pixel thermal events. Even though the total rate of 8.9 MHz per SiPM device seems high, the probability that a stream of such events produces enough pile-up to exceed a realistic threshold is negligible. For example, at 10 MHz of dark rate the probability of measuring a charge corresponding to at least 50 pixels within a 100 ns gate is 10^{-86}. Rather than pile-up, it is cosmic rays and electronic transients that are the primary sources of large spontaneous pulses. At nominal operating bias, these false triggers were observed to occur on the bench at a rate of order one to a few Hz. These will be completely dominated by beam-related room background from neutrons and X-rays under realistic tagger running conditions.

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

1. Repeat the detection efficiency measurements with the 2 mm SiPM to determine the source of the anomalously high efficiency that was measured previously using the yellow LED pulser.

2. Repeat the detection efficiency measurements with the second 2 mm SiPM sample and compare with the first. Obtain a fast green and fast blue LED spanning the peak in the BCF-20 emission spectrum and check the dependence of the efficiency on wavelength.

3. Vary the temperature and bias voltage of the SiPM and measure the dependence of the gain, detection efficiency and dark rate on the operating conditions. Use these data to determine the limits on the operating conditions that are to serve as requirements in the design of the SiPM "base" electronics and housing.

4. Design a 16-channel "base" board that houses the SiPMs together with their voltage regulators, preamplifiers, thermal sensors, etc.

5. Build a 3×5 scintillating fiber prototype array and couple it to a prototype base and data acquisition system for bench tests.

6. Test the prototype with electrons and check its stability and timing performance at high rates.