1-11-2017

Design and Performance of SiPM-Based Readout of PbF₂ Crystals for High-Rate, Precision Timing Applications

J. Kaspar
University of Washington

A. T. Fienberg
University of Washington

D. W. Hertzog
University of Washington

M. A. Huehn
University of Washington

P. Kammel
University of Washington

Follow this and additional works at: https://uknowledge.uky.edu/physastron_facpub
See next page for additional authors

Part of the [Astrophysics and Astronomy Commons](#), and the [Physics Commons](#)

[Right click to open a feedback form in a new tab to let us know how this document benefits you.](#)

Repository Citation

Kaspar, J.; Fienberg, A. T.; Hertzog, D. W.; Huehn, M. A.; Kammel, P.; Khaw, K. S.; Peterson, D. A.; Smith, M. W.; VanWechel, T. D.; Chapelain, A.; Gibbons, L. K.; Sweigart, D. A.; Ferrari, C.; Fioretti, A.; Gabbanini, C.; Venanzoni, G.; Iacovacci, M.; Mastroianni, S.; Giovanetti, K.; Gohn, W.; Gorringer, Tim; and Pocanic, D., "Design and Performance of SiPM-Based Readout of PbF₂ Crystals for High-Rate, Precision Timing Applications" (2017). *Physics and Astronomy Faculty Publications*. 652.
https://uknowledge.uky.edu/physastron_facpub/652

This Article is brought to you for free and open access by the Physics and Astronomy at UKnowledge. It has been accepted for inclusion in Physics and Astronomy Faculty Publications by an authorized administrator of UKnowledge. For more information, please contact UKnowledge@lsv.uky.edu.

Design and Performance of SiPM-Based Readout of PbF₂ Crystals for High-Rate, Precision Timing Applications

Digital Object Identifier (DOI)

<https://doi.org/10.1088/1748-0221/12/01/P01009>

Notes/Citation Information

Published in *Journal of Instrumentation*, v. 12, P01009, p. 1-20.

© 2017 IOP Publishing Ltd and Sissa Medialab srl

After a 12-month embargo period from the publication of the Version of Record of this article, everyone is permitted to use, copy, and redistribute this article for non-commercial purposes only, provided that they adhere to all the terms of the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported license: <https://creativecommons.org/licences/by-nc-nd/3.0>

The document available for download is the authors' post-peer-review final draft of the article.

Authors

J. Kaspar, A. T. Fienberg, D. W. Hertzog, M. A. Huehn, P. Kammel, K. S. Khaw, D. A. Peterson, M. W. Smith, T. D. VanWechel, A. Chapelain, L. K. Gibbons, D. A. Sweigart, C. Ferrari, A. Fioretti, C. Gabbanini, G. Venanzoni, M. Iacovacci, S. Mastroianni, K. Giovanetti, W. Gohn, Tim Gorringer, and D. Pocanic

Design and performance of SiPM-based readout of PbF₂ crystals for high-rate, precision timing applications

**J. Kaspar^{a,*}, A. T. Fienberg^a, D. W. Hertzog^a, M. A. Huehn^a, P. Kammel^a, K. S. Khaw^a,
D. A. Peterson^a, M. W. Smith^a, T. D. Van Wechel^a, A. Chapelain^b, L. K. Gibbons^b,
D. A. Sweigart^b, C. Ferrari^{c,d}, A. Fioretti^{c,d}, C. Gabbanini^{c,d}, G. Venanzoni^c,
M. Iacovacci^{e,f}, S. Mastroianni^e, K. Giovanetti^g, W. Gohn^f, T. Gorringer^h, D. Pocanicⁱ**

^aUniversity of Washington, Department of Physics, Box 351560, Seattle, WA 98195, USA

^bCornell University, Department of Physics, 511 Clark Hall, Ithaca, NY 14853, USA

^cLaboratori Nazionali Frascati dell' INFN, Via E. Fermi 40, 00044 Frascati, Italy

^dIstituto Nazionale di Ottica del C.N.R., UOS Pisa, Via Moruzzi 1, 56124 Pisa, Italy

^eINFN, Sezione di Napoli, Ed. 6 Via Cintia, 80126 Napoli, Italy

^fUniversità di Napoli, Corso Umberto I 40, 80138 Napoli, Italy

^gJames Madison University, Department of Physics and Astronomy, 901 Carrier Dr, Harrisonburg, VA 22807, USA

^hUniversity of Kentucky, Department of Physics and Astronomy, 505 Rose Street, Lexington, KY 40506, USA

ⁱUniversity of Virginia, Department of Physics, 382 McCormick Rd, Charlottesville, VA 22904, USA

E-mail: kaspar@uw.edu

ABSTRACT: We have developed a custom amplifier board coupled to a large-format 16-channel Hamamatsu silicon photomultiplier device for use as the light sensor for the electromagnetic calorimeters in the Muon $g - 2$ experiment at Fermilab. The calorimeter absorber is an array of lead-fluoride crystals, which produces short-duration Cherenkov light. The detector sits in the high magnetic field of the muon storage ring. The SiPMs selected, and their accompanying custom electronics, must preserve the short pulse shape, have high quantum efficiency, be non-magnetic, exhibit gain stability under varying rate conditions, and cover a fairly large fraction of the crystal exit surface area. We describe an optimized design that employs the new-generation of thru-silicon via devices. The performance is documented in a series of bench and beam tests.

KEYWORDS: Silicon photomultiplier; SiPM; Electromagnetic calorimeter.

---

*Corresponding author.

Contents

1. Introduction	1
2. Application-Specific Design Constraints	2
3. SiPM Selection	3
4. Electronics Design	5
4.1 Spice Model	5
4.2 Final Design	6
4.3 Operational point	8
4.4 Production and Testing	10
4.5 SiPM bias load tester	11
5. Performance and Tests	13
5.1 Pulse Shape and Calibration Procedure	13
5.2 Timing resolution	14
5.3 High- and Variable-Rate Performance Tests	15
6. Conclusions	19

---

1. Introduction

The rapid development of, and subsequent improvements in, commercial pixelated silicon photodiode sensors operated in Geiger mode — silicon photomultipliers, or SiPMs¹ — has led to a vast array of end-use applications in the physical sciences and in medical imaging [1]. High-energy and nuclear physicists increasingly are turning to the use of SiPMs in applications where small photomultiplier tubes (PMTs) would have otherwise been used (e.g., Refs. [2, 3]). This is particularly well motivated when space for the light sensor is constrained — SiPMs and their readout boards can be quite compact — and also when the detector must reside in a high magnetic field environment, as SiPMs performance is unaffected by multi-Tesla fields.

SiPMs have high quantum efficiencies across a broad range of wavelengths, which are typically emitted by scintillation detectors or even Cherenkov light sources. Because of the avalanche properties’ nature of single pixels within a device, the intrinsic rise time is fast, making them ideal for

---

¹Many names exist, but the community is largely converging on SiPM as a useful mnemonic. The Hamamatsu products described herein are called Multi-Pixel Photon Counters, or MPPCs, by the company.

time-of-flight counters and other applications having a premium on timing performance. In many applications originally using SiPMs, the emphasis has been on few photo-electron (PE) counting situations, where the single PE sensitivity is critical and easily exceeds that of PMTs. Impressive few PE resolution is typically a feature of commercial catalogs documenting performance. In contrast, here we describe the use of SiPMs in the many hundreds to thousands of PE regime associated with large light output situations such as calorimetry. As described below, even here, the SiPM holds many advantages over a PMT.

The dream that a SiPM can simply replace a PMT at a fraction of the cost, on the other hand, is naive. While costs have indeed fallen considerably, great care must be taken to enlarge the photo-sensitive surface and maintain the intrinsic performance of the smaller, core devices, which are usually only a few square mm in size. At present, the user can only purchase pre-packaged arrays² from the manufacturer as single devices or build custom arrays from chips that can be mounted edge-to-edge with very little dead space between them. In either case, the sub-units, or channels, must be gain matched and summed to produce the output of an effective single element.

The gain of a SiPM depends sensitively on the overvoltage above the Geiger-mode breakdown threshold, V_{bd} . Since V_{bd} depends on temperature, gain is intrinsically linked to bias voltage supply stability *and* to external temperature stability. The typical dependence of V_{bd} on temperature for recent Hamamatsu devices is 60 mV per °C at the recommended operational overvoltage. The gain sensitivity at the overvoltage operational point is $\sim 0.2\%$ per mV. Because of this, the mechanical housing, with some sort of cooling circuit or stabilization, and the bias control system, possibly with feedback based on online or offline SiPM temperature readout, must be carefully considered. For large-area, multi-channel devices, attention to the summing and pulse-shape electronics is critical and non-trivial. Finally, if the application involves high rate and potential pulse pile-up, special attention must also be paid to the bias voltage supply and its interface to the SiPM in order to provide an adequate current source to rapidly replenish the extracted charge from fired pixels.

Despite existing equivalent circuit models to guide the design, fine-tuning of the amplifier circuit owing to finite tiny capacitances, inductances, and resistance in circuit elements and board traces can be critical to achieve an optimized pulse shape. Our application goal for a light sensor requires a *PMT-like* performance with respect to pulse shape and two-pulse resolution as well as a *better-than-PMT* stability against high rate. In the report that follows, we demonstrate a design that meets these goals, which we believe will be of general interest to the community.

2. Application-Specific Design Constraints

For context, we briefly describe the application-specific requirements of the Fermilab Muon $g - 2$ experiment. Sixteen bunches of 3.1 GeV/ c polarized positive muons will be injected into the muon storage ring [4, 5] every 1.4 s. As the muons decay ($\gamma\tau_\mu = 64.4 \mu\text{s}$), the emitted positrons, having any energy up to ~ 3.1 GeV, will curl towards the ring interior where they might strike one of the 24 electromagnetic calorimeter stations that are located at regular intervals immediately adjacent to

---

²e.g., 4×4 , 8×8 , and 16×16 arrays of 9 mm^2 units are now available from Hamamatsu.

the storage ring volume. The magnetic field at the larger radius of the calorimeter is 1.45 T, and it drops to 0.8 T over the 225 mm width of the detector. The calorimeter is segmented into six rows and nine columns of $25 \times 25 \times 140 \text{ mm}^3$ lead-fluoride (PbF_2) crystals. Each crystal is read out by a single monolithic SiPM.

The calorimeter is required to provide information for determining the energy and time-of-arrival of each decay positron with the following emphases:

- The time resolution, extracted from a fit of the SiPM pulse, must be better than 100 ps for positrons having energy greater than 100 MeV.
- The calorimeter must be able to reliably resolve two showers by temporal separation with time separations greater than 5 ns, and it must accurately assign correct energies to the two pulses.
- The energy resolution of the reconstructed positron energy summed across calorimeter segments must be better than 5 % at 2 GeV.
- The maximally allowed, uncorrected, gain change is $\delta G/G < 0.04 \%$ over the $700 \mu\text{s}$ time period while the rate reduces by four orders of magnitude, starting with an initial rate that can exceed 1 MHz.
- The analog and digital signal chains must preserve the pulse shape of the Cherenkov light collection.
- The calorimeter, readout, and cabling all must reside adjacent to a highly uniform magnetic field, and the components must not perturb this field.

In Ref. [6], we documented many calorimeter-specific performance tests that have led to the final calorimeter design. That test employed an earlier version of the SiPM circuit design and used simplified digitization electronics to record the data. Here we describe the final SiPM design and performance and selected beam-test results using the custom 800 MSPS and 12-bit depth digitizers [7, 5], which will be used in the Muon $g - 2$ experiment.

3. SiPM Selection

Because multi-particle pile-up separation is a dominant systematic effect in our high-rate application, we looked for devices delivering a narrow pulse shape without additional shaping or clipping. The type of silicon substrate and carrier mobility contribute to the pulse shape. Devices with better carrier mobility deliver narrower pulses but typically require a higher bias voltage [8]. Because the number of Cherenkov photons in our application is greater than that needed to fulfill our requirements on energy resolution, the absolute quantum efficiency of the SiPM was not among the critical parameters.

While the market has evolved considerably since we began our development, the Hamamatsu surface-mount 16-channel SiPM used here is a fairly typical modern device appropriate for our development. The 16 channels are packaged as a single 4×4 array of $3 \times 3 \text{ mm}^2$ discrete channels with

solder pads on the back.³ The net $12 \times 12 \text{ mm}^2$ active area is similar to that of a 19-mm diameter photomultiplier, such as the Hamamatsu R750. The $50 \mu\text{m}$ pitch selected yields a total of 57,344 pixels.⁴

In this newer generation of Hamamatsu SiPMs, the original poly-crystalline silicon quenching resistor was replaced with a Ni-based resistor⁵, which is less sensitive to temperature variance. The connection to the summing board uses through-silicon vias to avoid parasitic inductances of the finite length traces, which are in series with the output current pulse. Optical trenches between pixels minimize cross-talk, and, finally, the use of high-purity silicon wafers minimizes dark current and after-pulsing and increases the charge delivered by a fired pixel.

A photon striking a pixel can cause an avalanche, initiating current flow. This current is merged with that from the other struck pixels in the same channel. The 16 channels are independent, except that they share the same applied bias voltage. To create a large single active channel — a true PMT-like device — requires a custom summing board to add the individual contributions from the channels, preserving the attractive signal characteristics of the smaller units.

After an avalanche is initiated in a pixel, the quenching resistor arrests the flow of current by dropping the effective bias below the Geiger threshold. The pixel recovers with a time constant typically on the order of 10 ns. Those pixels that had not been struck remain ready for an event.

In general, good linearity is achieved when the number of pixels in the device, N_{device} , greatly exceeds the highest expected number that might fire for a given event. When two photons strike the same pixel, only a single avalanche can result; the response is intrinsically non-linear. In our application, a maximum of several thousand pixels will fire for the highest-energy positron shower. The $\sim 5\%$ pixel occupancy fraction is in the near-linear regime. For higher photon count applications, one can reliably estimate the number of pixels, N_{pixels} , that would have fired if the device had an infinite number of pixels available,

$$N_{pixels} = -N_{device} \ln \left(1 - \frac{N_{fired}}{N_{device}} \right). \quad (3.1)$$

This conversion requires only that the user has calibrated the output pulse integral/height to determine the fired pixel count, N_{fired} , a straight-forward procedure described in Section 5.1.

Hamamatsu SiPMs require a stable bias voltage in the range of 65 – 70 V. For our application, involving nearly 1300 elements, we grouped 12 – 16 SiPMs having common manufacturer specified breakdown voltages, to a single bias supply. A stable and nearly identical operating temperature is maintained on these grouped SiPMs by flowing forced outside air across the amplifier board’s heat sinks, which are enclosed in a common light-tight detector housing. Accommodation for temperature variances of the group can be made by small bias voltage adjustments if necessary.

---

³Model number S12642-4040PA-50.

⁴A 4×4 pixel sized area centered on each channel is used for the through-silicon via.

⁵The Ni fraction is tiny; we evaluated the magnetic perturbation of this device in our test magnet, and the effect of Ni creates a negligible perturbation.

4. Electronics Design

Several approaches can be used to read out the charge from 16 channels and sum them into a final pulse output. One can use a common shunt resistor to create a voltage amplifier, or construct a transimpedance (current) amplifier based on either discrete or integrated electronics.

The advantages of the voltage amplifier design are its simplicity, robustness, and ease of simulation, which make it particularly suitable for the characterization of an unknown SiPM. However the shunt resistor in the circuit is much larger than the transimpedance design. Therefore, the pulse width due to the RC time constant is wider. Furthermore, the bias voltage depends on the current drawn, making the gain dependent on the signal size. This may not present a problem for time-of-flight applications or for counting single photons.

For our application in calorimetry — where several thousand pixels can fire in a single event — a transimpedance amplifier is a better choice. The arrival time dispersion of the Cherenkov photons at the SiPM is typically around 1 ns, and the variance of this distribution — i.e., the width of the pulse — provides critical information on multi-particle pileup and, to a lesser degree, the incident angle of the incoming particle. The width variance of the input light is preserved in the output pulse shape. The amplifier does not shape the pulse beyond the limits of its bandwidth.

Practical factors also enter into the design considerations. In our application, the footprint of the printed circuit board (PCB) on which the SiPM is mounted must be less than $25 \times 25 \text{ mm}^2$. The system as a whole requires remote control of the variable gain amplifier, readback of an on-board probe, and local EEPROM storage of a detector identifier and gain settings. Owing to space restrictions, a single multi-component HDMI connector is used to convey $\pm 5 \text{ V}$, the bias voltage, and the communication protocol. Finally, to match the input requirement of the custom 800 MSPS digitizer, a differential signal with a 2.5 V DC common-mode voltage is needed.

4.1 Spice Model

Designing the summing circuit requires a model of the SiPM device itself. We started by using a SPICE model⁶ incorporating an electronic equivalent of the SiPM. Variants have been discussed in Refs. [9, 10]. Following Seifert, the parameters used to describe one of the 16 channels having 3584 microcells, each of $50 \mu\text{m}$ pitch, are listed in Table 1.

Two principle changes have been made to the current generation of Hamamatsu SiPM products compared to those modeled originally in the literature. They are fabricated using higher purity silicon wafers, and the quenching resistor now uses a Ni-based one rather than poly-crystalline silicon.

The 16-channel SiPM has a common bias line, but the $3 \times 3 \text{ mm}^2$ individual channels are isolated electrically; they represent 16 truly independent channels. The goal is to sum the 16 channels to create an effectively larger device, while preserving their timing and pulse-shape characteristics. Three options considered include: 1) a simple sum of all 16 channels in one stage, 2) an amplification of each of the 16 channels separately, followed by adding the outputs, and, 3) a multi-staged approach, which we adopted and will describe next. The selection of 3) reduces the capacitance four-fold compared to

---

⁶<http://bwrce.eecs.berkeley.edu/Courses/IcBook/SPICE/>

Table 1. Parameters used in the SPICE model to describe one of the 16 channels in the device

Description	Value
Capacitance of the reverse-biased diode pixel	$C_d = 55.8 \text{ fF}$,
Parasitic bulk capacitance	$C_g = 50 \text{ pF}$
Total measured capacitance of the channel	250 pF
Effective series resistance of the avalanche	$R_d = 1 \text{ k}\Omega$
Resistance of the Ni based quench resistor	$R_q = 100 \text{ k}\Omega$
Parasitic stray capacitance of the quench resistor	$C_q = 3.4 \text{ fF}$
Pixels fired in each of the 16 channels of the array	$N_c = 200$

1) and, therefore, reduces the RC time constant controlling the pulse width associated with the pixel recover time. Option 2 would be better, but it is difficult to implement because of the increased local heat generation and component space restrictions.

The SPICE model and board prototypes were developed in an iterative way with 12 versions of prototype boards in total. The performance of each version was evaluated with a variable width pulse generator driving an LED, and a narrow width laser. The pulse generator, Avtech AVPP-1A-B, provided a pulse width in the range from 0.5 ns to 3.0 ns. The pulsed laser, Picoquant LDH-P-C-405 driven by a PDL 200-B, was used to evaluate pulse-shape fluctuations and high-rate stability. The feedback from the prototype tests was used mainly to tune parasitic impedances of the first-stage components and the PCB material.

4.2 Final Design

The final schematic, shown in Fig. 1 is based on the concept of a multi-staged amplifier. At the inverting input of the first stage, current pulses from four SiPM channels are summed together and converted into voltage pulses in a fixed-gain transimpedance amplifier, which is operated at a constant gain of 300. This first stage is designed around a THS3202 2-GHz bandwidth current feedback op-amp. The impedance looking into the input of the op-amp is very low — less than $100 \text{ m}\Omega$ — at frequencies below 1 MHz. Above 1 MHz, the input impedance has a rising characteristic as a function of increasing frequency. The net input impedance approximates a 7 nH inductor in series with $50 \text{ m}\Omega$. If the SiPM channels were connected directly to the op-amp input, the input inductance would resonate with the SiPM's capacitance C , resulting in high-frequency oscillations. The resistor $R_S = 3.48 \text{ }\Omega$, which is placed in series between the common anode side of a group of four SiPMs and the op-amp input, dampens this resonant behavior. Another resonance exists at a lower frequency owing to the inductance of the printed circuit board traces and interconnections, and the SiPM capacitance. This resonance does not cause oscillation, but instead adds a damped sine wave to the exponential tail of the SiPM pulse. This resonance is dampened by resistor $R_L = 2.0 \text{ }\Omega$.

In the second stage, the four partial row sums are added together using a THS3201 op-amp operated at a voltage gain of 0.25.

The final stage is designed around a LMH6881 digitally controlled variable gain differential amplifier, which has differential inputs and outputs. The output of this stage drives a shielded twisted

Figure 1. Simplified schematic of the amplifier board. The 16 channels of the SiPM share a common bias voltage. The circuit proceeds from the top left by first summing four channels within a row and then summing the rows. A final output stage, which employs a variable gain amplifier, delivers the differential signal to external waveform digitizers. The dashed inductors are present only in the SPICE simulation and represent parasitic properties of the components and PCB.

pair connected to a waveform digitizer, which has a 100- Ω differential input impedance with a common mode voltage of 2.5 V. The amplifier's common-mode output voltage is set by a resistive voltage divider of the +5 V power supply connected to a bias set pin that is not shown in the simplified schematic. The inputs of the differential amplifier are AC-coupled through 0.1 μ F capacitors in series with 49.9 Ω resistors. The negative input is AC-coupled to ground while the positive input is AC-coupled to the output of the summing amplifier.

A pole-zero network is inserted in series with the positive input, consisting of a 50 pF capacitor in parallel with a 1 k Ω resistor. This reduces the width of the output pulse by canceling out the exponential tail of the output pulse, but it preserves the meaningful sensitivity to the pulse width that corresponds to the arrival time of the photons impinging on the SiPM active surface. The pole-zero network also removes the unwanted artifacts of AC coupling, serving as a weak baseline restorer; and it improves the overall stability of the signal chain. The SiPM chip inherently behaves as a capacitor itself, and the additional AC coupling between the amplifiers does not introduce a new problem. The solution preserves the pulse shape as accurately as practically possible.

Figure 2 shows the SPICE model predictions of the absolute value of the pulse shape after the first, second, and final stages of the amplifier board. We emphasize that these are the shapes the model produces; the actual pulse shape after the second stage is narrower, and the final output shape can be compared directly to what is seen in Fig. 6. The lower trace in Fig. 2 shows the current in the SiPM

Figure 2. Upper three traces: SPICE simulation of the SiPM pulse evolution as the pulse undergoes the three amplification stages on a SiPM board. The amplitudes of the pulses are arbitrarily scaled for an easier comparison of the pulse shapes. Lower trace: SiPM current in the model.

in the context of the SPICE model.

The gain of the differential output amplifier is set by a BeagleBone microcomputer over an SPI data bus. The gain is adjustable over a 20 dB range in 0.25 dB steps. A 1 V_{pp} amplitude at the output of the amplifier board corresponds to approximately 4000 pixels fired across all 16 channels.

The practical realization of this circuit is shown in Fig. 3, where both sides of the PCB are shown. The outer board dimensions are constrained to be within the $25 \times 25 \text{ mm}^2$ calorimeter crystals onto which the SiPM “front” side of the board is attached using a high-index-of-refraction optical epoxy. The rear side of the board has a 3-pin output for the differential signal out, and an HDMI socket into which a standard HDMI cable is inserted.

4.3 Operational point

Our desired overvoltage operational point is determined primarily by optimization of the gain stability, as noise levels in our overvoltage range of interest are small, at the per mille level, compared to the expected signal. Of greater concern is large gain changes associated with charge depletion following large pulses in the SiPM and related rate-dependent gain shifts. By defining gain G as the ratio of charge drawn through the SiPM to the number of incident photons⁷, a pulse of a known light level will be proportional to the gain. The gain is a function of the bias overvoltage, V . Following a pulse of a fixed light level, a SiPM will see a temporary decrease in V that is proportional to the size of the charge pulse and thus proportional to G . Therefore, to leading order, one would expect a fixed light

⁷Note that this definition of gain folds in the changes in photon detection efficiency with overvoltage and therefore is nonlinear in the overvoltage.

Figure 3. The left and middle panels show the final amplifier board with the surface-mount 16-channel SiPM. The three-pin connector on the bottom right of the middle panel provides the AC-coupled differential voltage signal out. The common bias voltage in, the board low-voltage power, and the SPI lines used to regulate gain and readout temperature are supplied through the HDMI connector. The right panel shows the completed board after the heat sink has been attached.

level pulse to effect a gain perturbation of the following form:

$$\frac{\delta G}{G_0} \propto -\frac{dG}{dV} r(t). \quad (4.1)$$

In the above equation, G_0 is the unperturbed gain, and $r(t)$ is a recovery function varying between 1 at the time of the pulse and 0 at $t = \infty$. The minus sign hints at the drop of gain following a pulse. This is a cumulative effect for multiple pulses and can therefore lead to rate-dependent gain shifts, which are a significant concern in the Muon $g - 2$ experiment. A key feature is that the relative gain drop is proportional to the derivative of gain with respect to overvoltage, a quantity that increases with increasing overvoltage for overvoltages reasonably close to zero.

Slow gain shifts can result from temperature changes. The SiPM breakdown voltage depends on temperature, so any uncompensated change in temperature will directly change the overvoltage and in turn the gain, which depends on overvoltage in a non-linear way. For infinitesimal changes and in the leading order of dV —in which the overvoltage is linear in temperature and the gain is linear in overvoltage, i.e., $G(V) = G(V_0) + \frac{dG}{dV} \delta V + \mathcal{O}(\delta V^2)$ —one finds that small drifts in temperature δT

lead to changes in gain as follows:

$$\frac{\delta G}{G_0} \propto -\frac{1}{G} \frac{dG}{dV} \delta T. \quad (4.2)$$

In the case of temperature-dependent gain changes, the relative size of the effect is proportional to the logarithmic derivative of gain with respect to overvoltage, a quantity that decreases with increasing overvoltage. The same is true of drifts in bias voltage not related to pulses.

The recommended operational voltage is prescribed by the manufacturer using their method to obtain the desired gain. The method utilizes the first-order derivative of the current-voltage characteristics, see Fig 5 of the relevant patent application [11]. The recommended voltage is prescribed as the breakdown voltage plus a fixed offset to achieve the desired multiplication factor. The breakdown voltage corresponds to the first inflection point on the current-voltage curve where most pixels turn into the Geiger mode. The fixed offset corresponds to the voltage difference between the first and second inflection points on the current-voltage curve. This choice of the offset balances the advantages of the high photodetection efficiency, and disadvantages of dark rate, crosstalk, and after-pulsing. The 16 channels of a SiPM array have different recommended operation points; however, for our purposes, the shapes of the current-voltage curves are similar enough for all of the channels of a SiPM array. Therefore, there is an effective current-voltage curve for a SiPM array on which the 16 voltage points recommended for the 16 channels lie. The 16 points are narrowly spread from each other in the linear region of the current-voltage curve. Owing to the uniformity of illumination, the 16 voltage points can be replaced by a single effective operational voltage, resulting in the effective gain of the SiPM array.

The manufacturer-recommended operational voltage is also our choice for the operational point in our application because it provides a sufficient gain, requires a temperature stability that can be met, and allows for the use of a single effective gain value for a whole SiPM array that scales with the bias voltage and temperature fluctuations in a predictable and near-linear manner.

4.4 Production and Testing

We contracted for 1400 SiPM amplifier boards to be fabricated and assembled commercially. It wasn't practically feasible to track the individual 16-channel SiPM arrays through the process of amplifier board assembly. The SiPM arrays were sorted into the groups of 10 SiPM arrays having similar average breakdown voltage V_{bd} specifications⁸. The groups were shipped to the assembly house such that the completed boards would accurately reflect which SiPM group was attached to each board.

Following assembly, a custom tailored heat sink was glued on the board, and then quality control (QC) tests were made. A dedicated QC station was constructed to test each of the 16 channels, plus the PCB communications and the temperature readback. As shown in the left inset of Fig. 4, a 16-hole Delrin collimator can be aligned precisely with respect to the 16 channels of the SiPM to isolate each channel optically from the others. An LED board having 16 SMD-type LEDs (right inset) is mounted to the rear of the collimator and is programmed to flash each of the 16 LEDs one by one

---

⁸Hamamatsu provides a measurement of V_{bd} for each of the 16 individual channels; the average value was used to sort the devices into groups of 10 for the purpose of amplifier board assembly.

Figure 4. Left: Elements of the QC station. The shown configuration features a fiber which illuminates the entire SiPM array from a laser light flash. The left inset shows the 16-channel collimator that is used together with the 16 LEDs in the right inset to configure the station to flash the 16 channels individually. Right: The normalized output amplitude for 10 SiPM arrays is plotted as a function of applied bias voltage. Three distinct groups of breakdown voltages exist in this group of 10 SiPM arrays.

during QC tests. The summed SiPM signal represents the output of only one channel at a time. These data provide the verification that each of the 16 channels is functioning properly.

In addition, we also tested the temperature sensor, programmable gain amplifier (PGA), and the flash memory (EEPROM) installed on the board. Approximately 4% of the boards failed the QC tests for various reasons, with most errors being recoverable by repairing or replacing various individual components, the external connector, or the seating of the heat sink, which accounted for 50% of the problems. As an example of the power of the QC program, one board had five of the 16 channels not responding (for unknown reasons). This defect would likely not have been noticed had we not flashed each channel individually.

In a second use of the QC station, the average breakdown voltage V_{bd} of the SiPM as a whole was measured with the aim of confirming the manufacturer's specification. The collimated LED system was replaced by a pulsed laser (red optical fiber in Fig. 4 left) that was positioned to illuminate the entire 16-channel SiPM at once. The bias voltage V_{bias} applied to the SiPM board was varied from 66 V to 69 V, and the output signal amplitude was plotted as a function of V_{bias} . Figure 4 right shows the results for 10 SiPMs belonging to the same group; i.e., they were specified to have similar V_{bd} values. The curve for each SiPM is normalized to the amplitude at 69 V. The SiPMs trend in three distinct groups as the bias voltage is lowered toward breakdown.

4.5 SiPM bias load tester

The transient response of the bias voltage power supply is critical in our application. The amount of light recorded by a SiPM decays exponentially with a half-life of 64 μs . For the purpose of testing and identifying a suitable power supply, a SiPM bias load tester (SBLT) was developed. It is a functional replacement of several SiPM boards from the power supply's point of view. The SBLT allows for a

We used the load tester, as well as subjecting an array of SiPMs to various light pulse sequences, to evaluate new commercial multi-channel bias supplies which had been developed for SiPMs. All of the new products failed to provide stable voltage in a variety of tests. However, the BK Precision 9124A single-channel unit was successful. In all cases, the challenge was to re-supply charge fast enough to sustain a high-rate drain, such as when the SiPMs are pulses at a sustained rate of several MHz. The failure modes for the multi-channel units included that the output voltage went into oscillation, or that it sagged significantly under the load. For those units, the advertised maximum current was no more than about 10 mA. The BK supply, with a full-scale maximum voltage of 72 V and a maximal output current of 1.2 A, on the other hand, was successful in the high-rate and other tests. The series resistance of the power supply was measured to be effectively $\approx 5\,\Omega$ for transient loads, and the measured transient response time was about $40\,\mu\text{s}$ for transient loads at the few mA level.

The intermediate conclusion is that SiPMs are in principle capable of operation at very high rates without significant gain drop but only if extra care is given to selection of a suitable bias voltage power supply. Further, it is important to achieve very low resistance of the complete power chain in series with the load. The desired quality of the power supply is its ability to maintain the voltage independent of the load. Examples of test results are shown in Section 5.3.

5. Performance and Tests

In this section, we report a number of results for the performance of the SiPM system described above. Characteristics of interest include the pulse shape, the time resolution, the calibration of fired pixels (or photo-electrons) to pulse height, the resolution of closely spaced pulses, the rate tolerance, and the gain stability as a function of rate. Measurements were carried out on the bench and using 3 GeV electrons together with the full calorimeter crystal array.

During the development phase, a fast 405 nm laser system was used to directly illuminate a SiPM over a range of rates and light level intensities. For dynamic gain stability tests, a background could be created by additionally firing a LED in known patterns with respect to the laser pulse. The LED driver was capable of creating a sharp pulse, whose SiPM response closely resembled the laser pulse in both time and amplitude.

The end-use configuration was tested using 3 GeV electrons provided by the SLAC End Station Test Beam Facility. The results reported here are based on the final implementation of a complete 54-element calorimeter station with individual SiPM readout of each PbF_2 crystal in the array. The signals from each of the channels were conveyed to custom 800 MSPS, 12-bit digitizers developed [7, 5] for the Muon $g - 2$ experiment. The complete laser-based calibration system [12, 13] was used to flash the front face of all crystals simultaneously at user-controlled rates and in specific pulse patterns. The stability of the laser light was sufficient to allow sequences of calibration runs to be made in which the laser light intensity could be used to reliably extract the photo-electron to pulse height ratio for each SiPM and thus calibrate the individual gains independent of the beam.

5.1 Pulse Shape and Calibration Procedure

A representative pulse shape response to the few-hundred ps wide laser pulse is shown in Fig. 6.

Figure 6. Templates following a large number of laser (dashed) and beam (solid) events. The narrower shape reflects the near instantaneous light arrival from the sub-ns laser photo-diode. The wider shape is from 3 GeV electrons striking normally in the center of a PbF₂ Cherenkov crystal and causing an electromagnetic shower.

The same figure also shows the response to a 3 GeV electron impinging centrally on the front face of one of the PbF₂ crystals. The pulse width from the beam is slightly wider, owing to the collection time of light from the electron shower. The FWHM is ~ 4.4 for the laser shape and is 5.2 ns for the beam shape. The smooth shapes are derived from a large number of individual pulses, which create “template” patterns that are used in subsequent individual pulse fits. A fit to a single digitized pulse requires only optimizing the time of the template because the amplitude is fixed from the integral of the pulse area. Importantly, the pulse shape is not a function of the number of pixels fired or of the rate. However, it is clearly a function of the light source, which in the application will be distinguished between beam or laser owing to external trigger tags.

Slight differences in pulse shape can be observed and are expected, owing to the arrival time of photons at the SiPM surface from showers initiated from electrons over a large angular range. In those cases, Monte Carlo anticipates the photon arrival profile, which is reflected in a slight increase in the pulse width. In our implementation, the crystals are wrapped in black Tedlar® foil to minimize the fraction of reflected — and therefore late-arriving — photons at the SiPM surface.

The calibration of photo-electrons per mean signal pulse integral (M) is made by the method described in Ref. [6]. For a Poissonian distribution of photons from the light source, the number of PE is obtained from the ratio M^2/σ^2 , where M is the mean of a distribution of pulses, and σ is the width of that distribution. The calibration procedure uses an automatic sequence of 10,000 pulse runs. Each run uses a unique setting of a neutral density filter wheel to control the light intensity in finite steps. Figure 7 shows a typical example of this process.

5.2 Timing resolution

Given the stable pulse shape and accurate template fitting procedure, the timing resolution can be

Figure 7. Plot of variance σ^2 versus pulse-integral mean M of a distribution of fitted laser pulses on a SiPM. The discrete mean values are based on using a multi-step filter wheel to attenuate the light. The sequence of runs begins and ends in the open position of the filter wheel, corresponding to the upper right most point. The inverse of the fitted slope corresponds to the PE/M and the good linearity implies that the variance entirely depends here on the statistics of the number of pixels fired on each event.

deduced from either the laser or beam events. The controlled response with the laser is more straightforward because nearly equal amplitude light pulses are injected simultaneously into all crystals. The difference between the fitted time of two SiPMs recorded within the same digitizer module is shown in Fig. 8. The width implies that the individual channels have a timing resolution of $26.7 \text{ ps} / \sqrt{2} = 19 \text{ ps}$. In this measurement, the light intensity in each crystal created approximately 2000 fired pixels.

5.3 High- and Variable-Rate Performance Tests

A critical test of either a SiPM or a conventional PMT in a high-rate experiment is the stability of the response to a high sustained average rate of light pulses and the short-timescale gain recovery following a significant pulse. It is well known that special care must be taken to maintain high voltage for PMTs in high-rate applications because the current drain in the capacitors of the base typically leads to pulse sagging or other pulse-shape distortions. How the gain might be affected by available charge has been discussed briefly in Section 4.5. The issue is related to the average number of pixels fired in an event as well as to the time scale to re-supply the charge that has been drained by the fired pixels on the leading pulse. There are several relevant time scales and pulse patterns that we tested.

Consider first the response to a high sustained pulse rate. We used the laser system described above to fire a train of pulses at the rate of 5 MHz into a single SiPM. Figure 9 left shows a trace recording of this sustained sequence of pulses, digitized at 800 MSpS. The first pulse of the pulse

Figure 8. Time difference in ps for two SiPMs each illuminated from the same laser pulse with approximately 2000 PE/pulse. The 19 ps time resolution of individual channels is a factor $\sqrt{2}$ smaller than the quoted width.

Figure 9. Response of the SiPM array to a direct uniform illumination of its surface using the Picoquant laser operating at the rate of 5 MHz. Left: a digitized trace is fitted using a template created from the first pulse of the train. Right: the first and last pulses of the pulse sequence are overlaid on top of each other. Most points are indistinguishable.

sequence was used to create a template for the fitter. All pulses fitted successfully using this template. The pulse shape is unchanged versus rate, showing no intrinsic signs of distortion, which is shown in Figure 9 right, where the first and last pulse of a pulse sequence are overlaid on top of each other. Critical to this test is a bias supply capable of delivering a high-enough current to maintain the voltage independent of the load.

The next example is related to the gain immediately following a significant pulse; i.e., on the time scale of several tens of ns, which is relevant for resolving two-pulse pileup events such as those shown in Fig. 10 where one needs to obtain correct times and energies of the pulses. The same time scale is typical for the SiPM chip to charge up, i.e., for the charge to transport from the local power supply on the SiPM board, formed by large capacitors, to the SiPM chip itself. When a SiPM pixel fires, it delivers charge, which is replenished afterwards. The delivered charge comes from all available sources, particularly the SiPM array itself. The SiPM array behaves as a charged capacitor;

Figure 10. Two 3-GeV electrons striking near the center of the same crystal with time separations of 4.55 ns (left) and 9.8 ns (right). Depending on the exact impact position, each pulse corresponds to approximately 2000 PE. The two pulses are easily resolved in both cases using the template fit. A small and predictable correction for the trailing pulse gain can be made based on the study of the short-term gain recovery.

i.e., the bias voltage experienced by the SiPM array drops momentarily before the charge is recovered. Figure 2, which plots the SPICE model pulse shapes throughout the circuit evolution, also shows the predicted voltage drop on the buffer capacitor as a function of time following a pulse. The scale is in arbitrary units; this voltage drop is tiny.

As an example, suppose that 1600 pixels fired simultaneously on the SiPM board, 100 pixels in each of the 16 channels. The charge delivered by a fired pixel is 1.5×10^6 electrons, for a total charge for each channel of 24 pC. This results in a voltage drop of 96 mV, taking into account the 250 pF capacitance of each of the 16 channels. The relative voltage drop is 4% compared to the typical operating overvoltage bias of 2.5 V.

The charge is replenished from the capacitors on the SiPM board. The only available path for the charge to the SiPM pixel goes through its quenching resistors. The product of the pixel capacitance, 55.8 fF, and the value of the quenching resistor, 100 k Ω , sets the expected timing constant for the charge recovery of about ~ 6 ns. To verify the charge recovery model and better understand the details of the charge recovery mechanism, the following exercise was performed.

In the study here, the SiPM board was fired using a LED pulse, and the process of charge recovery was probed using the laser. The LED pulse shape was tuned to match the laser pulse shape as closely as possible. The amplitude of both pulses corresponded to about fired 1600 pixels per SiPM board. The laser pulse was scanned from -40 ns to $+80$ ns with respect to the LED pulse. Each measurement consisted of two steps: 1) in the first step, only the laser is fired, 2) in the second step, both the LED and laser are fired. The pulse integrals were extracted by a custom two-source template fit, and the ratio of the laser pulse-integrals was calculated as the second step compared to the first step, i.e., the SiPM undergoing recovery compared to the unloaded SiPM. Assuming that the number of photons fired by the laser was the same in both steps — a good assumption — the ratio directly reveals the gain drop as a function of pulse-time separation. The measurement was repeated many thousand times, and the ratios were averaged to obtain a single point in Fig. 11. This measurement method is robust with respect to the fluctuations in laser intensity and temperature.

The results of this gain study indicate two components of the gain recovery time response. A

Figure 11. Relative gain of the laser pulse with respect to the reference LED pulse versus a fixed time interval between the two pulses.

representative fit is shown in Fig. 11 with 9- and 29-ns recovery time constants and relative amplitudes of 1.7 to 1, respectively. Note that this fit is merely a guide to the eye. The faster component is to be compared with the expected value of ~ 6 ns. The difference can be explained by the extra parasitic impedances of the PCB and SiPM chip, which are unaccounted for in our simple model. The slower component is typically attributed to the bulk properties of the SiPM chip, which is a satisfactory explanation for our purposes. In our application, the measured gain recovery allows us to correctly separate pileup events in the energy domain. The predictive drop measured here can also be used to correct the gains of trailing pulses in the reconstruction phase.

The next test is somewhat application specific, but it may be of general interest as an example that required tweaking some component values on the PCB to minimize gain variance. In the Muon $g - 2$ experiment, the storage ring is filled with muons, and observation of their subsequent decays occurs for $700 \mu\text{s}$, with a time-dilated muon lifetime of $64 \mu\text{s}$. Immediately after injection, the instantaneous rate on many of the detectors will exceed several MHz, but then it drops by more than four orders during the measuring period. From the SiPM perspective, an average physics event in a single crystal will fire 250 pixels of a SiPM, and there will be approximately 72 events such events during a measuring period, distributed exponentially in time. This kind of transient load will result in a gain sag, which must be minimized and also precisely determined in our application.

The local power supply on a SiPM board is formed by the capacitance of $5 \mu\text{F}$. The transient load of a SiPM array amounts to 5 nC for the typical SiPM gain of 1.5×10^6 electrons per pixel fired. Without any external bias voltage power supply, the delivered charge lowers the bias voltage by 1 mV at the end of the $700 \mu\text{s}$ period. For the typical bias over-voltage of 2.5 V , the relative voltage drop would be 4×10^{-4} , well within the specification on the acceptable gain drop. A detailed

Figure 12. Average gain as a function of time with an initial rate of 2.1 MHz, 1500 PE pulses, and a rate proportional to $\exp(-t/\tau)$ with time constant $\tau = 64 \mu\text{s}$. The integrated current drawn is higher than we expect in the Muon $g - 2$ experiment for the hottest crystal by a factor of 10. The slight droop and subsequent recovery of the gain is affected by a combination of the bias voltage supply and the buffer capacitance.

SPICE simulation, including a bias voltage power supply and power distribution lines, predicts the bias voltage drop on a SiPM array of $300 \mu\text{V}$. For an external bias voltage power supply to help with the transient load, its transient response needs to be faster than the load. The stability of a SiPM board also requires a 2Ω resistor in series with the power supply to decouple the cable. The measured transient response of the BK power supply was on the order of $40 \mu\text{s}$ in our setup.

To test the gain stability in this situation, the laser was programmed to fire in a pseudo-random pattern with an instantaneous rate that followed an exponential decay with time constant $\gamma\tau_\mu$. The spectrum of pulses was analyzed for pulse height stability under the assumption of a stable laser light level throughout the pattern. The validity of the assumption was verified by a neutral density filter which varied the load on the SiPM but preserved the load on laser. The observed gain drop scaled exactly as expected for a given filter value.

Figure 12 shows the results of a test with an instantaneous load approximately 10 times greater than expected in the experiment. The observed largest gain sag of 1.1% will be reduced proportionally in the application. Knowledge of the full gain curve will be made to sub-per-mille precision using a custom laser calibration system.

6. Conclusions

We have described our development of a custom amplifier board that supports a large-area, 16-channel SiPM from Hamamatsu. While the design was driven by the end-use calorimetry requirement for the

new $g - 2$ experiment at Fermilab, the desired performance characteristics are relatively common in many nuclear and particle physics applications. They include: operation in high magnetic fields, good (near) linear response for pulses of hundreds to thousands of photo-electrons, operation at MHz rates, short-duration “PMT-like” pulse shapes, and a high degree of gain stability. The final design PCB, together with careful attention to the external bias supply, achieves these goals. Notably, we measure ~ 5 ns pulse widths and achieve ~ 20 ps time resolution in the tests described.

For ease-of-use, the PCB also includes a variable gain amplifier, an on-board temperature read-out, and a built-in EEPROM for identification of the unit and storage of gain settings. The low voltage, bias voltage, and communication lines are all conveyed using commercial HDMI cables. The differential signal output is conveyed by a custom Samtec cable. The PCB components and the cabling are non-magnetic.⁹ At this time, we have completed the production and testing of 1400 SiPM / amplifier boards. Lengthy running of subsets of them over weeks of operation has shown no operational failures.

Acknowledgments

We thank A. Para for excellent advice on evaluating and choosing SiPM products and C. Hast for hosting us during the SLAC test beam measurements. We acknowledge earlier contributions to our SiPM testing and development from L. P. Alonzi, B. Kiburg, and P. Winter. This research was supported by the U.S. National Science Foundation MRI program (PHY-1337542), by the U.S. Department of Energy Office of Science, Office of Nuclear Physics under Award Number DE-FG02-97ER41020, by the Istituto Nazionale di Fisica Nucleare (Italy), and by the EU Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant Agreement No. 690835.

References

- [1] D. Renker. Geiger-mode avalanche photodiodes, history, properties and problems. *Nucl.Instrum.Meth.*, A567:48–56, 2006. doi:10.1016/j.nima.2006.05.060.
- [2] Elton S. Smith. Development of silicon photomultipliers and their applications to GlueX calorimetry. *AIP Conf. Proc.*, 1753:070006, 2016. doi:10.1063/1.4955369.
- [3] K. Ieki. Upgrade of the MEG liquid xenon calorimeter with VUV-light sensitive large area SiPMs. *Nucl. Instrum. Meth.*, A824:686–690, 2016. doi:10.1016/j.nima.2015.11.047.
- [4] G. T. Danby et al. The Brookhaven muon storage ring magnet. *Nucl. Instrum. Meth.*, A457:151–174, 2001. doi:10.1016/S0168-9002(00)00704-X.
- [5] J. Grange et al. Muon ($g-2$) Technical Design Report. 2015. arXiv:1501.06858.
- [6] A. T. Fienberg et al. Studies of an array of PbF₂ Cherenkov crystals with large-area SiPM readout. *Nucl. Instrum. Meth.*, A783:12–21, 2015. arXiv:1412.5525, doi:10.1016/j.nima.2015.02.028.

---

⁹Owing to the high sensitivity of any materials placed in or near the Muon $g - 2$ storage ring that could perturb the uniformity of the magnetic field, we had to remove by hand the HDMI ferromagnetic cable end clips. Additionally, we had to select non-magnetic capacitors for the PCB.

- [7] Antoine Chapelain. Development of the electromagnetic calorimeter waveform digitizers for the Fermilab Muon $g-2$ experiment. *PoS*, EPS-HEP2015:280, 2015.
- [8] N. Dinu, Z. Amara, C. Bazin, V. Chaumat, C. Cheikali, G. Guilhem, V. Puill, C. Sylvia, and J.F. Vagnucci. Electro-optical characterization of sipm: A comparative study. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 610(1):423 – 426, 2009.
- [9] F. Corsi, A. Dragone, C. Marzocca, P. Delizia, A. Del Guerra, et al. Modelling a silicon photomultiplier (SiPM) as a signal source for optimum front-end design. *Nucl.Instrum.Meth.*, A572:416–418, 2007. doi:10.1016/j.nima.2006.10.219.
- [10] S. Seifert, H. T. van Dam, J. Huizenga, R. Vinke, P. Dendooven, H. Lohner, and D. R. Schaart. Simulation of silicon photomultiplier signals. *IEEE Transactions on Nuclear Science*, 56(6):3726–3733, Dec 2009. doi:10.1109/TNS.2009.2030728.
- [11] K. Sato, S. Kamakura, S. Nakamura, T. Ohta, M. Hirayanagi, H. Suzuki, and S. Adachi. Photodiode array, method for determining reference voltage, and method for determining recommended operating voltage, January 10 2013. US Patent App. 13/247,192. URL: <https://www.google.com/patents/US20130009266>.
- [12] A. Anastasi et al. The calibration system of the new $g-2$ experiment at Fermilab. *Nucl. Instrum. Meth.*, A824:716–717, 2016. doi:10.1016/j.nima.2015.11.059.
- [13] A. Anastasi et al. Electron beam test of key elements of the laser-based calibration system for the muon $g-2$ experiment. 2016. arXiv:1610.03210.